Policy Research Working Paper 10545

Learning from Self and Learning from Others

Experimental Evidence from Bangladesh

Florence Kondylis
John Ashton Loeser
Mushfiq Mobarak
Maria Ruth Jones
Daniel Stein

WORLD BANK GROUP
Development Economics
Development Impact Evaluation Group
August 2023

Policy Research Working Paper 10545

Abstract

Can decentralizing demonstration accelerate learning about new technologies? This paper randomizes access to a fixed demonstration kit for new flood-saline-resilient seeds across villages in Bangladesh, with demonstration either by a single farmer or spread across many farmers. In the short run, higher learning from self and others under decentralization increases technology adoption. In the long run, the impacts of any demonstration persist, but the additional impacts of decentralization vanish. A Bayesian model of learning the returns to a new technology suggests belief dispersion caused noisy adoption along the learning path, and farmers' expected gains from demonstration are four times higher under decentralization.

This paper is a product of the Development Impact Evaluation Group, Development Economics. It is part of a larger effort by the World Bank to provide open access to its research and make a contribution to development policy discussions around the world. Policy Research Working Papers are also posted on the Web at http://www.worldbank.org/prwp. The authors may be contacted at fkondylis@worldbank.org and jloeser@worldbank.org.

Learning from Self and Learning from Others: Experimental Evidence from Bangladesh*

Florence Kondylis John Loeser Mushfiq Mobarak
Maria Jones Daniel Stein

JEL Classification Codes: D83, O12, Q12, Q16

Keywords: Agricultural Extension, Demonstration Plot, Technology Adoption, Learning

---

*Kondylis: Development Impact Evaluation, World Bank, fkondylis@worldbank.org; Loeser: Development Impact Evaluation, World Bank, jloeser@worldbank.org; Mobarak: Yale University, NBER, ahmed.mobarak@yale.edu; Jones: Development Impact Evaluation, World Bank, mjones5@worldbank.org; Stein: IDInsight, daniel.stein@idinsight.org. We thank Paul Christian, Kyle Emerick, Andrew Foster, Erin Kelley, Sylvie Lambert, Arianna Legovini, Karen Macours, Jeremy Magruder, and seminar audiences at IFPRI, Naples, Paris-Dauphine, Paris School of Economics, UC Berkeley, University of Virginia, and Wageningen University. We thank the Global Agriculture and Food Security Program, the World Bank Bangladesh Country Management Unit, the World Bank i2i, RSB, and SAFANSI Trust Funds for generous research funding. Kieran Byrne and Marcela Mello provided fabulous research assistance; Kristoffer Bjarkefur, Cynthia Sobieski, Christina Seybolt, Md. Munstasir, Mehrab Ali, and Ahmad Wahdat managed impressive field work. We thank Innovations for Poverty Action in Bangladesh for data collection work. Finally, we thank the staff at the IAPP project implementation unit for being outstanding research partners. The study protocols were approved by Yale University IRB (#1208010629) and IPA IRB (#00006083), and is registered in the AEA RCT Registry (AEARCTR-0010751). Computational reproducibility verified by DIME Analytics. The views expressed in this manuscript do not reflect the views of the World Bank. All errors are our own.

Contents

1 Introduction

Given the continued importance of agriculture to developing economies, accelerating agricultural technology adoption has the potential to substantially reduce global poverty (World Bank, 2007). However, information frictions can limit the diffusion of new technologies (De Janvry et al., 2017; Magruder, 2018). In response, agricultural extension services often set up *demonstration plots*, to inform farmers about the characteristics of new technologies and persuade them to adopt (Conley and Udry, 2010). While this model has shown some promise (Dar et al., 2020), it often falls short of triggering large learning gains (Kondylis et al., 2017; Duflo et al., 2022).

In this paper, we account for the role of learning from scale and decentralization in the optimal design of the demonstration plot, augmenting a recent literature that studies the diffusion of information through social networks. In addition to learning from both their own and others' demonstration, farmers may also learn from scale (the area under demonstration, holding fixed the number of farmers demonstrating) and from decentralization (the number of farmers, holding fixed the area under demonstration) (Foster and Rosenzweig, 1995).

While learning from scale and from others motivates traditional demonstration by a single well-connected farmer (Banerjee et al., 2013; Beaman et al., 2021), learning from decentralization or from oneself motivates expanding participation in demonstration to many farmers. The tradeoff between decentralization and scale is particularly stark under a fixed allocation of demonstration resources: as more and more farmers demonstrate, the demonstration plots get smaller and smaller. Accounting for these forces is essential to formulate even the simplest policy recommendations. For instance, when farmers learn primarily from themselves and learning from scale is weak, there may be low returns to selecting the best cultivator to run a large demonstration plot.

We design and implement a field experiment to estimate the impacts of approaches to demonstration suggested by competing theories of how farmers learn. We randomize three approaches to demonstration across 110 groups of 25 rural farmers called *farmer field groups*, ranging from traditional demonstration by a single farmer to fully decentralized demonstration in which many farmers gain first-hand experience with the new technology. We then combine our experimental results and predictions from a model of learning the returns of a new technology to separate and quantify the relative importance of learning from scale and learning from decentralization in the adoption of a new technology. We do so in climate change vulnerable coastal Barisal, in the context of the Government of Bangladesh's large-scale effort to promote new, higher-yielding flood-saline-resilient seeds for traditional pulses bypassed by the Green Revolution (Gollin et al., 2021). These new varieties of traditional pulses were developed not only with the goal of increasing productive resilience, but also to shift farmers away from profitable yet unsustainable dry season paddy cultivation.[1]

Our benchmark treatment is the standard (*regular*) demonstration protocol where the extension worker designates one farmer in the group to receive a demonstration package of seed and fertilizer sufficient to cultivate 0.4 hectares. This designated farmer is asked to run a one-season demonstration plot and share insights with others in the group.

In a second treatment (*shared*), the extension worker identifies up to four farmers who cultivated contiguous plots and are willing to set up a shared demonstration plot. The exact same demonstration package is divided across this group of four farmers. The shared arm is designed such that demonstration is socially decentralized, but not geographically, relative to the regular arm.

In a third set of randomly selected field groups, the extension worker splits the same demonstration package amongst up to twelve interested farmers, so that they can experiment with the new seed on their own plots. Their plots need not be contiguous, so this third (*decentralized*) arm triggers both social and geographic decentralization of demonstration.

We compare these three arms to a control arm in which the new seed is promoted by the extension agent to the farmer field groups, but no demonstration resources are provided to the village. The promoted seeds were locally multiplied and widely available in markets across all study villages.

As we experimentally induce many farmers to demonstrate in a village, we consider three main mechanisms that may drive differences in adoption of the new technology in the decentralized arm relative to other arms of our study. First, more farmers learn from self. Second, the amount of social learning also increases. Third, the size of each individual demonstration plot diminishes, which could affect the strength of learning

[1]New flood-saline-resilient seeds were promoted for five crops: lentils, mung, mustard, sesame, and wheat. Excluding wheat, this was the government's first effort to promote flood-saline-resilient varieties of these crops. While neither wheat nor sesame are pulses, our use of "pulses" is motivated by the fact that over 80% of cultivation of the five crops in our data is of the three pulses.

from each demonstration if scale matters. If growing conditions and the geographic proximity of plots are the main determinants of learning, engaging multiple farmers in the cultivation of one centralized demonstration plot will generate less learning than having the same number of farmers cultivate decentralized demonstration plots. Similarly, if there are high returns to the scale of demonstration, expanding participation holding fixed the area under demonstration will reduce learning.

Demonstration in any form induces learning, increasing the adoption rate for the new seeds by 145% (4.5 percentage points, or 35% adoption among pulses cultivators) in the two years *after* demonstration relative to control villages. The resulting adoption is sufficient to increase cultivation of the promoted pulses by 57% among farmers otherwise unlikely to cultivate pulses. Consistent with these positive impacts of demonstration on adoption, we find the new seeds generate positive returns: adoption eliminates the negative effects of saline-flood shocks on yields and profits over our study period. Despite these productive impacts, cultivating the new flood-saline-resilient seeds may not be optimal for all farmers: while the new seeds increase the returns to traditional pulses, the returns to pulses compared to paddy or other options will vary across farmers (Suri, 2011). This optimality of incomplete adoption is characteristic of the customized agricultural technologies essential for sustained productivity growth in a changing climate (Suri and Udry, 2022).

Socially and geographically decentralizing demonstration sharply increases adoption of the new seed in the first year after demonstration. Being assigned to the decentralized arm leads to an additional 6.5 percentage point increase in adoption relative to the regular treatment arm (significant at the 1% level), or a 258% increase in adoption over the control mean. In contrast, assigning multiple farmers to cultivate a demonstration plot in the shared arm increases adoption by 1.8 percentage points relative to the regular demonstration arm over our study period, although this effect is imprecisely estimated.

While the impacts of any demonstration on adoption persist two years after the end of the demonstration season, they fall, and the additional impacts of decentralization vanish over time. While demonstration increases adoption by 156% relative to the control group one year after the demonstration year, this effect is almost halved two years after demonstration. How can we reconcile that farmers may in fact learn faster under decentralized demonstration over the course of our experiment with the patterns of disadoption we observe?

We rationalize our results through the lens of a simple model in which Bayesian farmers learn the returns of a new technology (Besley and Case, 1994). Farmers start out with a prior belief about the distribution of the returns to potential new seeds; our experimental variation in demonstration induces adoption of the new seeds by providing a noisy signal of the returns to the new seeds, with which farmers update their beliefs about the realized new seeds. This model is commonly used as a microfoundation of technology adoption and diffusion in social networks (Young, 2009; Beaman et al., 2021). While other models could match our findings, we provide empirical evidence that our model captures how farmers learn—demonstration shifts farmers' beliefs and in turn their adoption decisions, and farmers' beliefs and adoption decisions respond to observed profits during demonstration. On the one hand, we observe long-run adoption that is persistently higher under any demonstration than adoption in the control group; in the model, this can be rationalized as farmers' expected returns converging to the actual returns, which are higher than their priors. On the other hand, this model allows adoption of a new technology to be nonmonotonic in learning the returns: noise generated along the learning path may induce farmers to make adoption mistakes, thus generating a "belief dispersion" effect that dominates the effect of convergence of average beliefs in our context. Hence, this model can reconcile patterns of disadoption ("inverted U-curve" of adoption) that closely matches our dynamics, while remaining general enough to predict an "S-curve" of adoption for alternative technology draws, as found by other studies (Griliches, 1957; Foster and Rosenzweig, 1995; Beaman et al., 2021). This is a crucial model feature to study technologies for which the adoption decision is discrete and reversible, such as improved seeds or irrigation.

We implement three sets of model-suggested tests of the mechanisms of learning, which relate back to our initial conceptual framework: we estimate the relative strength of learning from self and learning from others, and we test for learning from scale and from geographic heterogeneity. First, we estimate the relative strength of learning from self and learning from others from the relative impacts on adoption of participating in demonstration and of having one additional social connection to a demonstrator. We find learning from self is seven times stronger than learning from others; scaling this by average number of social connections, it implies that roughly half of learning is social. Second, we find that the impacts of decentralizing demonstration are fully explained by the impacts of participating in demonstration and of connections to demonstrators; as decentralizing demonstration reduces the scale of demonstration conditional on participation in demonstration and of connections to demonstrators, we fail to reject that there are no scale effects in learning. Third, and similarly, we fail to reject the absence of learning from geographic heterogeneity; we view this as consistent with observed convergence across demonstration modalities, as under learning from geographic heterogeneity, we would expect geographically decentralizing demonstration to generate persistently more information within the village, in sharp contrast to the prediction under social frictions.

If farmers are making so many adoption mistakes, is decentralizing demonstration always good? While decentralizing appears to speed up learning about a new technology, the costs generated by adoption mistakes on the learning path may swamp the gains from faster learning. To answer this question, we structurally estimate the model by matching the observed impacts of demonstration on farmers' adoption decisions; the estimated model allows us to quantify both the realized gains from decentralization and counterfactual gains under alternative technology draws, in addition to farmers' prior beliefs about the distribution of potential technology draws.

First, we estimate farmer's *ex-ante* gains from demonstration across potential draws of technologies from farmers' prior beliefs.[2] We find that, pooled across approaches to demonstration, farmers expect gains from demonstration slightly larger than the per farmer cost of demonstration kits. Importantly, the social externality from demonstration exceeds the cost of demonstration kits only under decentralized demonstration, as decentralizing demonstration quadruples farmers' ex-ante gains. In addition, while we found no evidence of scale effects in learning in our context, allowing for scale effects does not affect the conclusion that decentralizing demonstration increases the gains from demonstration—decentralizing demonstration also decentralizes observation, increasing gains under decreasing returns to observation of new technologies.

We conclude by estimating farmers' gains from demonstration under the technology draw realized in the context of our field experiment. We find that, even though the promoted flood-saline-resilient seed was more profitable than farmers' prior expectations, ex-ante rational noisy adoption generated ex-post *losses* from demonstration of

[2]In drawing technologies from farmers' prior beliefs, we take a "revealed preference" approach to external validity leveraging our structural estimates. This approach is distinct from, though complementary to, other approaches to external validity leveraging direct estimates of impacts across multiple draws (Rosenzweig and Udry, 2020).

similar magnitude to ex-ante gains. This is consistent with the notion that an aggressive promotion strategy (e.g., *"This is the Green Revolution of lentils!"*) may have increased the uncertainty of farmers' priors, thus driving high levels of noisy adoption. Considering a counterfactual where the agricultural extension service communicates more precisely with farmers reveals that, while ex-ante gains from demonstration fall (there is less left to learn), ex-post losses under the realized flood-saline-resilient seed turn to gains (farmers make fewer mistakes). This result corroborates the idea that improvements in the reliability of communication from agricultural extension services are important complements to the promotion of incrementally transformative agricultural technologies key to agricultural productivity growth (Chandrasekhar et al., 2022; Suri and Udry, 2022).

This paper is organized as follows. Section 2 describes the study context, and Section 3 our experimental design and data. Section 4 presents our experimental estimates of the impacts of demonstration and decentralizing demonstration on technology adoption. Section 5 develops a model of technology adoption and learning the returns to a new technology. Section 6 tests the model, and potential mechanisms underlying our experimental estimates. Section 7 estimates a structural model of technology adoption and learning that matches our results. Section 8 concludes.

2 Context

2.1 Setting

We study the adoption of new improved flood-saline-resilient seeds under the auspices of the Integrated Agricultural Productivity Project (IAPP). The project was implemented from August 2011 to December 2016 across eight districts by the Ministry of Agriculture of Bangladesh, funded by the Global Agricultural and Food Security Project and the Government of Bangladesh, with supervision from the World Bank. In Barisal district, where our study takes place, IAPP promoted dry season adoption of improved saline-flood resistant varieties for traditional crops bypassed by the Green Revolution: lentils, mung, mustard, sesame, and wheat ("pulses").[3]

The promotion of the new seeds was implemented through Bangladesh's national

[3]Although only three of these five crops are pulses—lentils, mung, and mustard—these three crops represent 83% of post-demonstration farmer-crop-survey wave observations with cultivation for the five crops; for concreteness, we refer to all five collectively as "pulses".

agricultural extension service. Promotion was administered by sub-Assistant Agricultural Officers, whom we refer to as extension agents. Each extension agent works in a single sub-district (*Union*), spanning on average 9 villages. To limit the potential for our demonstration interventions to displace extension agents' attention away from farmers in control and non-experimental villages, our experimental sample consists of 2 randomly selected villages in 55 Unions of the Barisal district, yielding a study sample of 110 villages across the district.[4]

One of the primary activities for extension agents under IAPP was to form farmer field groups of approximately 25 farmers in each of the project village. These farmers were self-designated during a village meeting with the extension agent; the only requirement was for these farmers to manifest a high level of interest in learning about new agricultural technologies and attend regular group meetings. The farmer field groups discussed a range of agricultural practices and technologies, including the adoption and properties of the new flood-saline-resilient seeds demonstrated in our experiment. They met monthly, on average, during the dry season.

2.2 Dry season agriculture in Barisal

Agriculture comprises two main growing seasons in Barisal: the monsoon season (*kharif* season) spans roughly from April to October, while the dry season (*rabi* season) runs roughly from November to March. Rice is the traditional staple crop in Bangladesh, and is well-suited to cultivation in the kharif season. While productive, rice cultivation in the rabi season (*boro* rice) requires heavy irrigation which is primarily provided by groundwater.[5] Heavy reliance on groundwater irrigation poses economic and environmental challenges: at the time of this study, heavy energy subsidies were in place to ensure the profitability of boro rice cultivation, and associated groundwater extraction was drawing down the water table.

We present basic descriptive statistics on dry season production of pulses and boro rice in Barisal in Appendix Table A5; our household survey data used to construct this table is described in Section 3.3. Consistent with the context above, the production of pulses, relative to boro rice, is associated with much less intensive irrigation but are also lower profits for farmers; these lower profits may be in part a consequence

[4] These villages are mapped in Appendix Figure A1.

[5] Most irrigation in Barisal originates from pump-operated shallow tube wells (67% in our study sample, compared to 5% using surface water).

of subsidies for groundwater irrigation. When cultivating boro rice is an option, farmers may therefore choose to do so rather than cultivate pulses. This decision may be influenced by the availability of more productive varieties of pulses promoted by extension agents through IAPP.

2.3 New Seeds, Same Old Crops

A direct consequence of climate change, the frequency of saline-flood events is intensifying in the Barisal district (Dasgupta et al., 2014). This particular manifestation of climate change is not unique to our study context, as the World Bank estimates that about 1.5 billion people are affected by floods, and a majority of them are in contexts where agriculture is a main source of livelihood (Dar et al., 2013; Lane, 2022). These tidal floods in Barisal reduce agricultural productivity not only by visibly damaging crops, but also as a shock to local groundwater salinity. Salinity brought by these tidal floods is compounded by rising groundwater salinity, as groundwater extraction for agriculture has depleted water tables and in turn enabled saltwater intrusion. These effects are strongest during the dry season, when the absence of rains prevents water table recharge and in turn salinity from leaving the superficial layers of the soil.

In this context of rising seas, falling water tables, and drying rivers, the adoption of flood-saline-resilient seeds holds potentially stark welfare implications; our experiment is embedded in the government's first large-scale effort to promote flood-saline-resilient varieties of pulses. These new flood-saline-resilient varieties of pulses were engineered by the Bangladesh Agriculture Research Institute (BARI) under the IAPP. These varieties were developed not only with the goal of increasing the productive resilience of pulses, but also to cause farmers to move away from cultivating boro rice, reducing the use of dry season irrigation and slowing groundwater depletion. Pulses were largely bypassed by the Green Revolution in Barisal; the delayed development of improved varieties of these crops not typically cultivated in richer countries has been an important determinant of slow agricultural productivity growth in lower income countries (Gollin et al., 2021). As part of this effort, extension services supported the multiplication of the new varieties by Union-level seed multipliers to ensure that farmers could access these new varieties on local markets throughout the duration of IAPP; this effort spans all our study years. This seed multiplication support ensured that seed supply did not limit adoption by farmers.[6]

In our context, we expect farmers will predominantly learn about the productivity of the improved seeds, rather than about the shape of or inputs to the production technology. Farmers in Barisal have been well exposed to the cultivation of the pulses our experiment targets, and the new seeds do not affect the recommended cultivation practices for these crops.[7] Instead, the improved seeds are both higher yielding and saline-flood-resilient—holding fixed cultivation practices and input use, farmers should expect higher yields, particularly during flooding, under cultivation of the improved seeds relative to traditional seeds.[8] While farmers are familiar with both saline-flood-resilient and high-yielding varieties of seeds in the case of rice, the increased yields associated with a given high-yielding variety of seeds will vary markedly both across varieties and contexts—for any given variety newly introduced to a context, farmers must learn just how impactful the variety is on yields (Laajaj et al., 2020).[9]

3 Experimentally Decentralizing Demonstration

3.1 Experimental design

We implement a demonstration experiment in our study sample of 110 villages of the saline-flood-prone Barisal district of Bangladesh. As Section 2 describes, our experiment was implemented under the auspices of IAPP, a large government program—specifically, our interventions were implemented through the farmer field group that was formed by IAPP in each village, and these group members make up our study sample.[10]

[6]The seed varieties promoted through IAPP are open-pollinated varieties, so in theory farmers could recycle seeds (that is, reproduce the improved seed from their harvest); in practice, we do not observe a meaningful degree of recycling of the improved seeds. Instead, the fact that these are open-pollinated varieties may have facilitated their supply on local markets.

[7]We empirically validate this by showing that the dynamics of adoption of the new seeds exactly tracks the decision to cultivate the promoted crops among novice pulses cultivators.

[8]Some of the improved seeds have shorter production cycles, potentially allowing them to be harvested more quickly.

[9]This does not imply farmers' input allocations will not adjust under the improved seeds: as the new seeds may both increase yields and reduce the downside risk of cultivating these crops, we may expect a "factor-deepening" effect as reported by Emerick et al. (2016).

[10]As there is exactly one farmer field group in each village and our study sample is comprised of group members only, we use farmer field group and village interchangeably when describing our

Timeline and experimental design We present the timeline of key activities associated with the IAPP project and the experiment in Figure 1. As the pulses are particularly suited for cultivation in the dry season, which starts in roughly November in Barisal, initial meetings with farmer groups were implemented a few months prior in July 2012. These initial farmer group meetings were used to collect our sampling frame for our baseline survey, which we describe in additional detail, along with other aspects of data collection, in Section 3.3. The baseline survey covered, among other things, agricultural production during the November 2011—March 2012 dry season, which we refer to as the "Pre-demo year".

Figure 1: Timeline

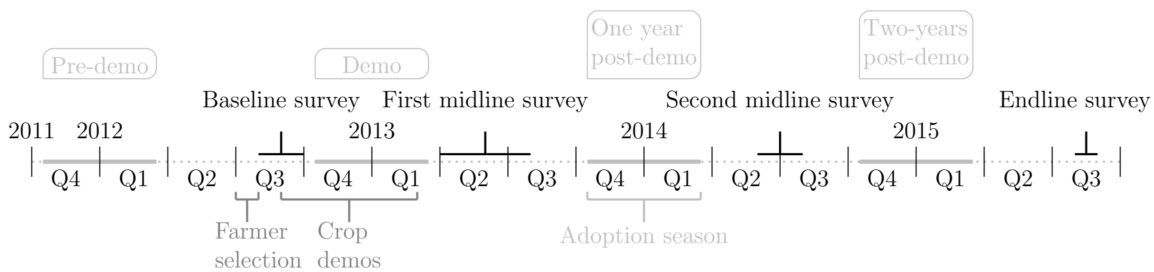

Notes: Dark gray solid lines indicate the boro season each year, denoted *Pre-demo*, *Demo*, *One year post-demo* and *Two years post-demo*, respectively. Survey periods are highlighted in black. Light gray brackets denote project interventions.

After initial farmer group meetings, we randomly assigned one fourth of the 110 villages in our study sample into one of three treatment arms plus a control group. Figure 2 summarizes the experimental design and sample sizes under each arm. We describe the activities implemented under each of the four experimental arms below.

Control In contrast to our three treatment arms, no demonstration resources were provided through farmer groups in control villages. However, as noted in Section 2, a number of activities associated with the promotion of the new seeds were implemented in control villages. Extension agents participate in regular meetings with the farmers groups to discuss and promote the new seeds, as well as other agricultural practices, regardless of treatment status. In addition, the new seeds were locally multiplied at the Union-level and available on all markets in the study area.

analysis.

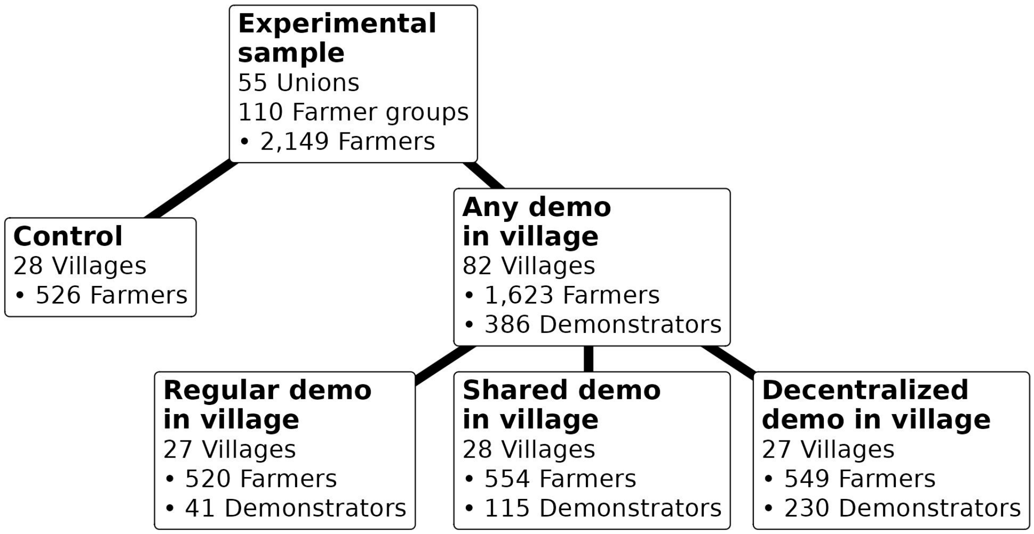

Figure 2: Experimental design

Notes: The assignment of villages across experimental arms is described in this figure. The number of farmers in our survey sample and, in villages assigned to demonstration, the number of demonstrators, associated with each experimental arm are presented in this figure.

Any demonstration In August and September of 2012, following initial farmer group meetings and treatment assignment, in each treatment village extension agents organized a meeting of the farmer group to introduce farmer-led demonstration of the new seeds. Demonstration was intended to complement the direct promotion of the new seeds by agricultural extension agents, and we discuss the mechanisms through which farmer-led demonstration may affect learning about and adoption of the new seeds by farmers in Section 3.2. The process of introducing demonstration of the new seeds was as follows.

First, extension agents selected up to three pulses with new seed varieties suitable for cultivation, and therefore for demonstration, in their Union during the November 2012 – March 2013 dry season ("Demo year"). This is important as it allows us to observe a demonstration crop for all villages in our study, even those not assigned to the demonstration treatment. The project then delivered one demonstration package for each selected crop per treated village, each of which contained sufficient quantities of fertilizer and new seeds to cultivate roughly 0.4 hectares, for provision to selected demonstration farmers.[11]

[11] In addition to the pulses, extension agents could also select improved varieties of dry season

Second, during the meeting to introduce farmer-led demonstration, the extension agent identified farmers who were willing to set up demonstration plots using the demonstration packages provided. Across treatment arms, each farmer could participate in demonstration of at most one crop.[12] Crucially, our three treatment arms generated variation in implementation at this stage, in how each demonstration package was divided across farmers; however, all three treatment arms hold fixed the total demonstration resources allocated across farmers.

Regular demonstration treatment arm In the *regular* demonstration arm, for each demonstration package, the extension agent designated a single farmer who would receive the full package to maintain a demonstration plot on their own land. Relative to control, the regular arm is designed to capture the impact of demonstration under the benchmark demonstration model (Kondylis et al., 2017; Beaman et al., 2021), where demonstration is led by a single farmer—the intent is that other farmers in the group learn from the demonstration, while the demonstration farmer also learns by demonstrating. Farmers selected under regular demonstration may differ from other farmers in their farmer groups; for instance, extension agents may select more central farmers, in the hope that this facilitates more learning by other farmers, or farmers with larger landholdings, who have a sufficiently large plot to cultivate the full demonstration package.

Shared demonstration treatment arm The *shared* demonstration arm, relative to the regular arm, decentralized demonstration socially but not geographically. For each demonstration package, the extension worker identified up to four farmers with contiguous plots to set up a shared demonstration plot. As noted above, these farmers received the same amount of demonstration resources as in the regular arm, thus holding fixed the size of demonstration within the village. Relative to the regular arm, the shared arm is designed to limit the geographic spread of the demonstration while increasing the number of farmers who demonstrate relative to the regular arm.

rice for demonstration. In contrast to the pulses, the development and promotion of improved varieties of rice is common in Barisal, and we found over 80 varieties available. We did not analyze these demonstrations, as the rice varieties made available for demonstration by IAPP were already familiar to farmers, so there was no scope for learning. Besides, rice demonstrations were unlikely to contribute to the IAPP objective of increasing cultivation of pulses to enable farmers to substitute away from dry season rice cultivation.

[12]Out of 386 demonstrators, we observe two who demonstrated multiple crops.

As a consequence, the scale of the demonstration and farmers' geographic proximity to it are expected to be similar across shared and regular arms.

Decentralized demonstration treatment arm The *decentralized* demonstration arm, relative to the regular arm, decentralized demonstration both socially and geographically. For each demonstration package, the extension worker allowed up to twelve farmer volunteers to split the demonstration package and experiment with the new seed on their own plots. This holds fixed the total size of demonstration in the village across all three treatment arms. In the decentralized arm, the resulting demonstration "minikits" remain sufficient to cultivate at least 0.03 hectares with the new seeds, which corresponds to a fairly typical plot size in our context. This approach to demonstration is in the spirit of a temporary subsidy (albeit one implemented by a quota floor) for experimenting with a new technology, as studied in Carter et al. (2021). Relative to the shared and regular arms, the decentralized arm allows many more farmers learn from their own demonstration, although the scale of each individual demonstration is smaller; in turn, a larger number of farmers may learn from observing these smaller demonstration plots, as farmers are more likely to be socially connected to a demonstrator and to be geographically proximate to a demonstration plot.

3.2 Conceptual Framework

We now discuss the mechanisms through which the demonstration modalities may interact with the way farmers learn about a new technology. The main rationale to provide demonstration resources is that demonstration increases both learning from self and learning from others: the farmer who leads the demonstration acquires first-hand experience with the technology, and other farmers in the village observe the demonstration and learn from the demonstrator. Crucial in designing the optimal allocation of demonstration resources is establishing the level of decentralization of demonstration that maximizes learning in the village. Our experimental design, and therefore our conceptual framework, relate to the identification of and mechanisms underlying this optimal level of decentralization—they focus on the case of a fixed allocation of demonstration resources within a village, adjusting the scale of each individual farmer's demonstration inversely proportionally to the number of farmers sharing the demonstration resources.[13]

As decentralization induces many farmers to engage in demonstration, opportunities for farmers to learn from their own experience increase, as the sub-panels in Figure 3 illustrate. At the same time, as more farmers experience the technology, there are more opportunities to learn from others, and social learning also increases. In the absence of other mechanisms, decentralization would cause learning to unambiguously increase through both channels. However, in the presence of "threshold effects" in social learning, a farmer may only decide to adopt once they have observed a minimum number of signals (Beaman et al., 2021). Conversely, above a certain number of signals, returns to decentralization may decrease (Foster and Rosenzweig, 1995; Conley and Udry, 2010). Depending on the relative strength of learning from self and others, decentralization may then only increase learning up to a point. Alternatively, traditional demonstration farmers may be more central than demonstration farmers under decentralized modalities, and decentralizing may reduce learning.

Similarly, as more farmers demonstrate the technology on their own plots, this increases farmers' proximity to a demonstration plot. If farmers learn most from observing the technology in similar growing conditions or proximity to their own plot (Conley and Udry, 2010), then decentralizing demonstration socially and geographically will also unambiguously increase learning through this proximity channel, relative to decentralizing socially only.

In contrast, decentralizing demonstration to many farmers under a fixed allocation of demonstration resources will mechanically reduce the scale of the demonstration farmers carry out and observe. If scale of the demonstration plays an important role in determining the precision of the signal a farmer extracts from their own and others' experimentation (Foster and Rosenzweig, 1995), then this mechanism will interact with the relative strength of learning from self and from others to predict an ambiguous effect of decentralizing demonstration activities on learning. When learning from scale is first order, the largest impacts are from providing all seed to one farmer, unless learning from self is very large relative from learning from

[13] While this conceptual framework discusses learning, we note that adoption of new technologies is commonly conceptualized as monotonic in learning, overcoming the empirical constraint that learning is often challenging to observe: it is assumed that increases in adoption correspond to increases in learning, and as more farmers learn, or as farmers learn more, about the technology, they become (perhaps weakly) more likely to adopt it (Griliches, 1957; Foster and Rosenzweig, 1995; Beaman et al., 2021).

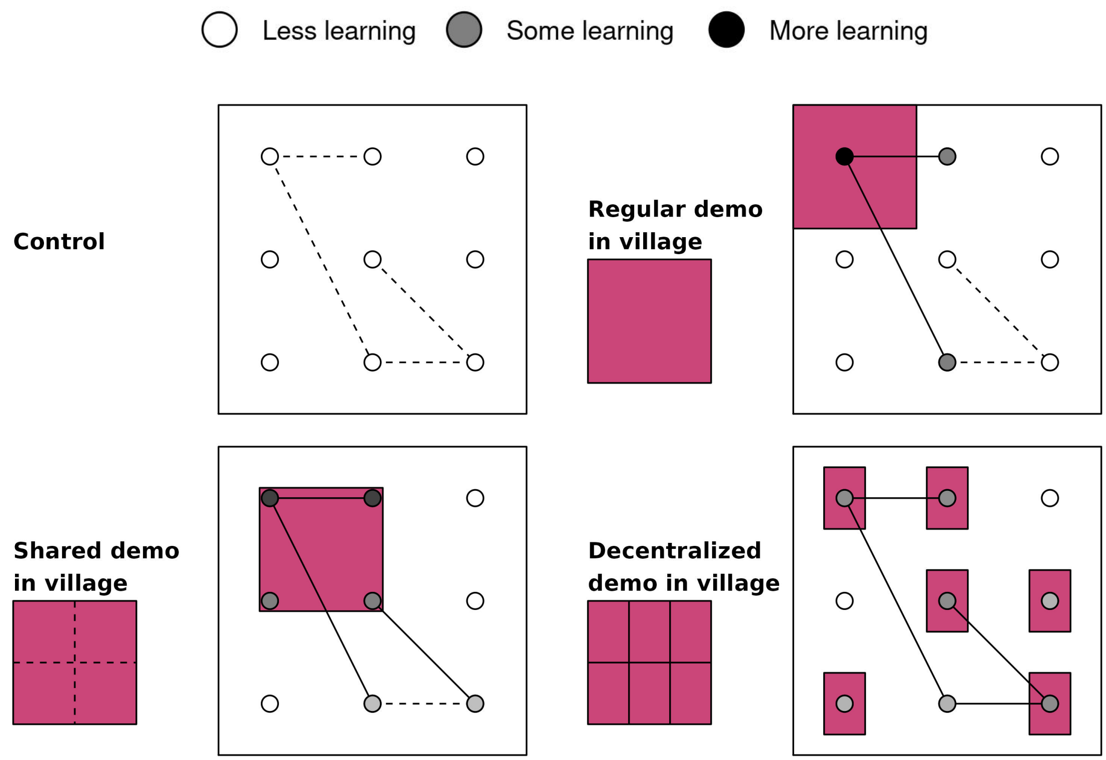

Figure 3: Linking Conceptual Framework and Experiment

Notes: A stylized representation of demonstration modalities across experimental arms, and associated learning predicted by our conceptual framework, is presented in this figure. Circles represent households. Solid and dotted lines between households represent social connections along which learning from demonstration does and does not occur, respectively. Circle color represents strength of learning from demonstration. Pink rectangles represent demonstration plots, with area proportional to allocated demonstration resources.

others. When learning from decentralization is first order, the largest impacts are from providing seeds to many farmers, unless learning from self is very large relative to learning from others.

By varying the degree of social and geographic decentralization of demonstration activities, our experimental design will allow us to gauge the relative importance of these mechanisms. However, as decentralization and scale interact with learning from self and others, there is no perfect experiment to isolate each of these forces. Hence, some structure will be needed to interpret the forces underlying our reduced form estimates in Section 4; a model in Section 5 will motivate additional empirical tests of the mechanisms of learning, implemented in Section 6, and allow us to estimate key parameters underlying learning in Section 7.

3.3 Data

We collected four rounds of detailed household surveys, complemented by monitoring data on demonstration plots, to support our analysis; we now discuss these data and their application to our analysis.

Household surveys Between 2012 and 2015, we administered four rounds of household surveys of farmer group members across our experimental sample. Each survey round was timed to immediately follow the harvest of the dry season; Figure 1 provides a detailed timeline of these survey rounds. Each survey round collected detailed records of agricultural production and technology adoption during the dry season, including adoption of the new flood-saline-resilient seeds promoted by the project. We leverage data from these surveys at the level of the village, household, household-survey wave, and household-crop-survey wave. Appendix A provides a detailed description of all variables used in our analysis and details their construction, and Appendix Table A6 shows how our sample size changes across survey waves and all sample restrictions used in our analysis.

Baseline survey The baseline survey was administered on a sample of 1,407 IAPP farmer group members within 78 randomly selected villages of the 110 villages in our experimental sample.[14] In addition to household characteristics, and the dry season agricultural data collected in each survey wave, the baseline survey included a module on social network linkages, in which we asked each household questions about their interactions with all other farmers in their farmer group (sampled or not). In our analysis, we consider a household to be in another household's social network if they report speaking to them, or were reported being spoken to by them, about farming in the last three months; to ensure consistent measurement across households, for all regressions leveraging social network relationships we restrict our sample to farmers and villages surveyed at baseline.

[14]The baseline survey covered 279 of the 386 eventual demonstrators, including 24 of 41 regular demonstrators, 81 of 115 shared demonstrators, and 174 of 230 decentralized demonstrators.

Follow-up surveys Our next three rounds of follow-up surveys preserved and expanded the baseline sample to add 496 farmer group members across the remaining 32 villages in our experimental sample, and an additional 246 farmer group members across the original 78 villages. Regressions that leverage baseline information such as social network relationships will not include these additional 32 villages and 742 group members; in Appendix Tables A19 and A20 we show that including or dropping these villages does not meaningfully affect our main results beyond the precision of our estimates.

In our first follow-up survey, just at the end of the demo year, we surveyed a much smaller sample of households and oversampled households in treatment arms with more demonstrators. We use the demo year household survey in two tables and one figure. In Table 5, we include demo year data for additional power when estimating the impacts of the new seeds on yields; these results are robust to excluding demo year data. In Appendix Table A1, we use this survey to construct measures of demonstration year profits for demonstrators. In Figure 4a, we show the dynamics of adoption of improved seeds across our four years of data.

In our second follow-up survey, one year post-demonstration, we additionally asked households about their knowledge of each of the new varieties of pulses, including characteristics of the new varieties. We construct an indicator that the household reported one of the new varieties for a given crop is a high-yielding variety; building on our conceptual framework, we interpret this as a binary proxy for households' crop-specific beliefs about the returns to the new seeds, and we incorporate this measure of beliefs into analysis of mechanisms in Appendix C.

Attrition Our main analysis focuses on impacts of demonstration one and two years post-demonstration. We attempted to revisit our full baseline sample in each of these waves; we reached between 81% and 89% of households across treatment arms during these two survey waves, and find no evidence of differential attrition across treatment arms (Appendix Tables A7 and A8).

Demonstration monitoring In addition to the household survey, we gathered operational data during the demonstration year (November 2012 – March 2013) through detailed communication with each of the extension agents. For each village, we identified the full set of crops for which demonstration packages were distributed, and which farmers demonstrated each crop. These administrative data align remarkably well with our self-reported household survey data; 5% and 15% of non-demonstrators (defined at the household-crop level) report adoption of improved seed and cultivation, respectively, for promoted crops during the demonstration year, in contrast to 80% and 94% of demonstrators.

Restriction to promoted crops Our main analysis leverages data at the farmer-crop-survey wave-level and restricts to what we refer to as "promoted crops". A key choice made by extension agents for treatment villages is the selection of pulses for demonstration of improved seeds; as we do not expect impacts of demonstration on non-demonstrated pulses, pooling demonstrated and non-demonstrated crops together in analysis would meaningfully reduce power. However, no crops were demonstrated in control villages. We therefore identify comparison crops in control villages based on the observation that extension agents were instructed to select crops for demonstration packages based on suitability in their Union (Section 3.1); specifically, we select the crops that were demonstrated in the other sampled village in their Union (and therefore selected by the same extension agent) as comparison crops.[15] We refer to comparison crops in control villages, and demonstrated crops in treatment villages, as "promoted crops", and non-promoted pulses as "placebo crops."

The restriction to promoted crops would result in bias if promoted crops were endogenous to treatment assignment; for instance, extension agents could have been aware of the village's treatment status (regular, shared, or decentralized) before selecting the demonstration package, and this information could have informed their choice of demonstrated crops.

We consider four tests for this bias, which we discuss and implement in Section 4.6. We compare, across treatment arms, the choice of promoted crops, baseline outcomes, and outcomes on placebo crops, and we estimate impacts controlling for crop-survey wave fixed effects. The results of these tests lead us to conclude that potential endogeneity of promoted crops to treatment assignment does not meaningfully bias our results.

[15] In 2 Unions, both villages were randomly selected to be control villages; these 4 villages are excluded from all analysis, as comparison crops are not defined.

3.4 Implementation fidelity

We leverage data from the monitoring survey to establish that our experimental design successfully generated different levels of social and geographic decentralization of demonstration activities across arms as Figure 3 illustrates. We present descriptive statistics on the implementation of our experimental design in Table 1. As discussed in Section 3.3, promoted crops are balanced across control and treatment villages and range in number from 1.61 to 1.88 crops. However, no demonstration farmers were designated in control villages, and no demonstration plots were established.

Table 1: Demonstration descriptive statistics

		Demo in village		
	Control	Regular	Shared	Decentralized
	(1)	(2)	(3)	(4)
Farmer group				
# of promoted crops	1.88	1.85	1.61	1.85
# of demo farmers	0.00	1.52	4.11	8.52
Farmer group-crop				
# of demos	0.00	0.82	2.58	4.62
# of demo plots	0.00	0.82	1.11	4.62
Farmer-crop				
Demo	0.000	0.042	0.129	0.223
# of connections to demo	0.000	0.463	0.952	2.683

Notes: Sample averages of implementation variables by treatment arm across farmer group, farmer group-crop, and farmer-crop observations are presented in this table.

In farmer groups assigned to the regular treatment arm, adhesion to treatment was high with on average 0.82 demonstration plots being set up for each promoted crop. Each farmer in a regular group had a 4.2% chance of demonstrating each promoted crop (roughly one farmer per 25 farmer group members), and has on average 0.46 connections to a demonstration farmer for a given promoted crop.

In farmer groups assigned to the shared treatment arm, implementation fidelity is also high with on average 2.58 demonstration farmers setting up just 1.11 demonstration plots for each promoted crop. The social decentralization of demonstration activities in the shared arm results in each farmer in a shared group having a 12.9%

chance of demonstrating and having 0.95 connections to a demonstration farmer, for a given promoted crop.

In farmer groups assigned to the decentralized treatment arm, the level of decentralization of demonstration activities is successfully higher than in the other treatment arms, with 4.62 demonstration farmers each setting up a demonstration plot for each promoted crop. Consequently, for each promoted crop, farmers in decentralized groups have a 22.3% chance of demonstrating and have on average 2.68 connections to a demonstrator.

4 Experimental results

We leverage the experimental variation produced by our RCT to produce three key results. First, we show that assigning demonstration resources to farmer groups increases group members' adoption of the promoted improved seed varieties (Figure 4a; Section 4.1). The resulting adoption fully offsets the negative impact of flood-saline events over the course of our study, as the improved seed increases crop resilience to saline flooding (Figure 4b; Section 4.2). Second, the social and geographical decentralization of demonstration to many farmers additionally increases technology adoption in the group (Figure 4c; Section 4.4). Third, these pooled estimates mask important dynamics (Section 4.3): in the long run, the impacts of demonstration on adoption persist but fall (Figure 4a), and the additional impacts of decentralizing demonstration on adoption vanish (Figure 4c). We discuss potential explanations of these dynamics in Section 4.5, and argue they are most consistent with increased learning (which we formalize in Section 5). We present analysis of robustness in Section 4.6.

4.1 Demonstration plots increase adoption

Empirical strategy We start by estimating farmer-crop level impacts of assigning demonstration resources to a farmer group, comparing outcomes in control villages and demonstration villages in the two dry seasons that follow the demonstration season.

As discussed in Section 3.3, we focus our analysis throughout the paper on promoted crops; we do so for power, as we do not expect effects of demonstration on

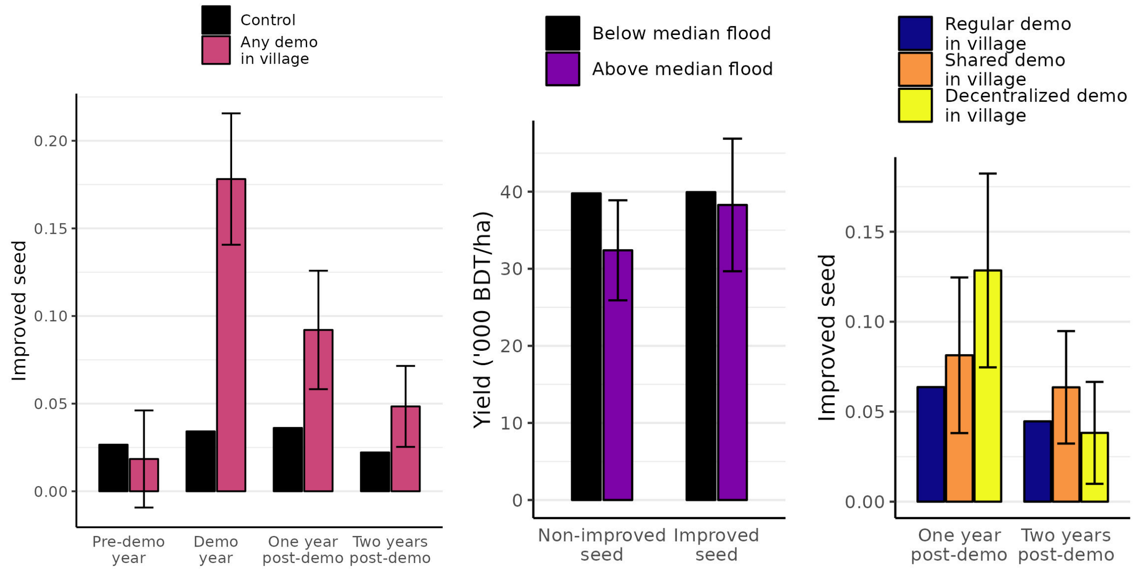

(a) Demonstration plots persistently increase adoption

(b) New improved seeds increase productive resilience

(c) Decentralizing demonstration increases adoption, but impacts vanish over time

Figure 4: Demonstration and decentralizing demonstration increase adoption of a productive technology

Notes: In Panel (a), control mean outcomes are presented in black, and treatment mean outcomes are presented in pink with 95% confidence intervals on treatment effects. Survey waves are on the horizontal axis. In Panel (b), non-improved seed, below median flood, mean outcomes, plus estimated differences from Equation 3, are presented. 95% confidence intervals are presented on effects of above median flood conditional on improved seed adoption. Improved seed adoption is on the horizontal axis. In Panel (c), regular demo in village mean outcomes are presented in blue, and shared demo in village and decentralized demo in village mean outcomes are presented in orange and yellow, respectively, with 95% confidence intervals on effects relative to regular demo in village. Survey waves are on the horizontal axis.

non-promoted crops. In Section 4.6, we present evidence that this restriction does not bias our results.

To estimate these impacts of assignment to demonstration on farmer groups, we estimate:

$$Y_{hcvt} = \alpha_t + \beta_t \text{Any demo in village}_v + \epsilon_{hcvt} \tag{1}$$

where Y_{hcvt} is outcome Y, for household h and promoted crop c, located in village v, in survey wave t. Any demo in village$_v$ is an indicator for village v being assigned demonstration resources under any treatment arm (regular, shared, or decentralized), and α_t are survey wave fixed effects. We impute zeros for outcome variables when the household does not cultivate the crop, with the exception of yields and fertilizer use; we therefore interpret estimated impacts on yields and fertilizer use as a combination of direct and selection effects.

Balance and attrition We establish that our village-level randomization achieved balance for key baseline farmer characteristics, both individually and jointly, across demonstration and control arms (Table 2). In Appendix Table A28, we show that inference in our tests of balance is not affected when we implement randomization inference, and we therefore instead present robust standard errors clustered at the village and associated p-values throughout our analysis.[16] In Appendix Table A7, we establish that attrition was low (12.5% to 18.9%) one and two years post-demonstration and not significantly affected by assignment to demonstration.

Results Pooling across the two years following the initial demonstration year, we find that assigning demonstration resources increases adoption of the new seed among farmer group members by 4.5 percentage points, or a 145% increase over the control mean (Column 1 of Table 3). This estimate conservatively suggests that adoption of the new seeds, conditional on cultivating the promoted crops, increased by 20 percentage points.[17]

Consistent with the objective of enabling substitution away from dry season paddy and towards pulses through the promotion of the new seeds, we find suggestive (but statistically insignificant) evidence of increased cultivation of promoted crops (Column 2 of Table 3) and decreased cultivation of paddy (Column 6 of Table 3). While neither of these coefficients are significant, we consider them further when we estimate heterogeneous effects with respect to experience with promoted crops.

We find no average impacts of being assigned to demonstration on yields or profit

[16] We present estimates of Equation 1 with randomization inference p-values included in Appendix Table A30. Inference is in general unaffected, although our estimates of the impacts of demonstration on fertilizer use fall marginally below conventional thresholds of statistical significance.

[17] We implement this calculation, and similar calculations throughout this Section 4, under the assumption that all demonstration-induced adoption of the new seeds is joint with demonstration-induced cultivation. Our estimated effect of demonstration on cultivation in Column 2 of Table 3 does not allow us to rule out either this possibility, or the possibility that all demonstration-induced adoption of the new seeds is *conditional* on cultivation. Under the latter possibility, our estimate would imply adoption of the new seeds, conditional on cultivating the promoted crops, increased by 35 percentage points.

Table 2: Balance across control and treatment and across treatment arms

	Mean (SD)				Coef. (SE) [p]		
		Any demo in village					
	Control	Regular demo in village	Shared demo in village	Decentralized demo in village	Difference		
	(1)	(2)	(3)	(4)	(2, 3, & 4) - (1)	(3) - (2)	(4) - (2)
Age of HHH (Years)	48.69	48.71	48.32	48.53	-0.18	-0.39	-0.18
	(11.93)	(11.91)	(11.85)	(12.56)	(0.83)	(0.97)	(1.14)
					[0.832]	[0.688]	[0.872]
Any pulses cultivation	0.203	0.240	0.246	0.160	0.012	0.006	-0.080
	(0.403)	(0.428)	(0.431)	(0.367)	(0.064)	(0.068)	(0.056)
					[0.850]	[0.934]	[0.153]
Veteran cultivator	0.117	0.100	0.126	0.117	-0.003	0.026	0.018
(Any promoted crop cultivation)	(0.323)	(0.300)	(0.332)	(0.322)	(0.061)	(0.046)	(0.041)
					[0.965]	[0.573]	[0.667]
Any pulses improved seed	0.050	0.076	0.052	0.040	0.006	-0.024	-0.036
	(0.218)	(0.266)	(0.223)	(0.196)	(0.027)	(0.028)	(0.027)
					[0.834]	[0.392]	[0.181]
Any irrigation	0.463	0.367	0.406	0.419	-0.065	0.039	0.052
	(0.499)	(0.483)	(0.492)	(0.494)	(0.111)	(0.118)	(0.119)
					[0.561]	[0.740]	[0.661]
Total plot area (ha)	0.202	0.163	0.198	0.171	-0.024	0.035	0.008
	(0.300)	(0.222)	(0.246)	(0.354)	(0.039)	(0.039)	(0.050)
					[0.537]	[0.369]	[0.873]
# of connections	9.05	10.75	7.19	10.50	0.37	-3.56	-0.25
	(6.65)	(6.43)	(5.04)	(5.76)	(1.62)	(1.57)	(1.75)
					[0.817]	[0.024]	[0.888]
# of observations (Farmer)	281	341	382	375			
# of clusters (Farmer group)	15	19	21	21			
Omnibus F-stat [p]					0.15	1.33	0.96
					[0.994]	[0.263]	[0.470]

Notes: Columns 1, 2, 3, and 4 present the baseline mean of the dependent variable and the baseline standard deviation of the dependent variable in parentheses, for farmers in control and regular, shared, and decentralized demo villages, respectively, and the total number of farmer and farmer group observations. Columns 5, 6, and 7 present differences, with robust standard errors clustered at the farmer group level in parentheses, and p-values in brackets.

for the promoted crops (Columns 3 and 5, Table 3). In Section 4.2, we use exogenous exposure to flood-saline shocks to present evidence that the new seeds increased productive resilience, and our point estimates imply the new seeds increase average yields by 2,900 BDT/ha, or roughly 7%. We note that our estimated impacts of demonstration on yields and profits cannot rule out effects consistent with these

Table 3: Demonstration increases adoption of improved seed

	Promoted crops, One and two years post-demo					
	Improved seed	Culti-vation	Yield ('000 BDT/ha)	Fertilizer ('000 BDT/ha)	Profit ('000 BDT/ha)	Paddy cultivation
	(1)	(2)	(3)	(4)	(5)	(6)
Any demo in village	0.045	0.047	-3.401	2.019	-0.572	-0.111
	(0.013)	(0.034)	(3.344)	(0.884)	(0.681)	(0.095)
	[0.000]	[0.163]	[0.310]	[0.023]	[0.401]	[0.245]
Control mean	0.031	0.130	42.746	5.149	1.874	0.438
Wave FE	X	X	X	X	X	X
# of observations (Farmer-crop-wave)	5,533	5,533	912	912	5,533	2,818
# of clusters (Farmer group)	96	96	81	81	96	96

Notes: Selected regression coefficients from Equation 1 are presented above, with included control variables listed below. Robust standard errors clustered at the farmer group level are in parentheses, and p-values in brackets.

impacts of the improved seed on yields.[18] We also revisit our average null effect on profits for promoted crops in Section 7.

Interestingly, we find a positive and significant impact on fertilizer expenditures (Column 4, Table 3)—we interpret this result as evidence of capital deepening in response to demonstration of the flood-saline resilient seed, as in Emerick et al. (2016). This effect is economically meaningful, and corresponds to a 39% increase over the control mean. We discuss this result further in Section 4.2 as we investigate the flood-saline resilient property of the new seeds.

Who adopts? One key determinant of the decision to adopt the new seeds is the joint decision to cultivate the associated promoted crop. As just 12% of farmers cultivated any of their village's promoted crops at baseline, we investigate baseline cultivation of promoted crops as a dimension of heterogeneity. This heterogeneity is particularly relevant as it captures baseline differences in determinants of crop choice, such as market access, experience, and plot characteristics. We refer to farmers who did and did not cultivate promoted crops at baseline as veteran and novice cultivators, respectively, following Conley and Udry (2010). We modify Equation 1 to allow for

[18] In particular, a 20 percentage point increase in use of improved seeds conditional on cultivation, times a 2,900 BDT/ha increase in yields, would result in a 600 BDT/ha increase in average yield, well within the 90% confidence interval for our estimated effect of demonstration on yield.

heterogeneity with respect to veteran status, and estimate:

$$Y_{hcvt} = \alpha_t + \beta_1 \text{Veteran cultivator}_{hv0} + \beta_2 \text{Any demo in village}_v + \beta_3 \text{Veteran cultivator}_{hv0} \times \text{Any demo in village}_v + \epsilon_{hcvt} \quad (2)$$

where Veteran cultivator$_{hv0}$ is an indicator variable that takes value 1 if for household h in village v cultivated any of the promoted crops in their village at baseline, and 0 if they did not ("novice cultivator"). We present estimates of Equation 2 in Table 4, pooling across the two years following the initial demonstration year.

Focusing on control villages, we find that veteran cultivators are 252% more likely to cultivate promoted crops; this result corroborates the notion that veteran cultivators are persistently different from novice cultivators in their propensity to cultivate the promoted crops. Veterans in control villages are also proportionately more likely to cultivate the new seeds, underscoring the link between decisions to cultivate promoted crops and adopt the new seeds. In addition, veteran cultivators are 51% less likely to cultivate dry season paddy.

Table 4: Novices and veterans

	Promoted crops, One and two years post-demo					
	Improved seed	Culti-vation	Yield ('000 BDT/ha)	Fertilizer ('000 BDT/ha)	Profit ('000 BDT/ha)	Paddy cultivation
	(1)	(2)	(3)	(4)	(5)	(6)
Veteran cultivator	0.054	0.252	-5.076	0.061	3.653	-0.231
	(0.019)	(0.059)	(6.825)	(1.374)	(2.437)	(0.096)
	[0.006]	[0.000]	[0.457]	[0.965]	[0.134]	[0.016]
Any demo in village	0.049	0.056	-5.090	2.216	-0.124	-0.129
	(0.011)	(0.031)	(4.134)	(1.495)	(0.687)	(0.118)
	[0.000]	[0.076]	[0.219]	[0.139]	[0.857]	[0.273]
Veteran cultivator ×	-0.008	-0.134	10.748	0.503	-0.855	0.118
Any demo in village	(0.024)	(0.062)	(8.032)	(2.143)	(2.666)	(0.111)
	[0.735]	[0.032]	[0.181]	[0.815]	[0.748]	[0.285]
[Any demo in village] +	0.041	-0.078	5.658	2.718	-0.979	-0.011
[Veteran cultivator ×	(0.026)	(0.063)	(7.539)	(1.505)	(2.584)	(0.103)
Any demo in village]	[0.120]	[0.218]	[0.456]	[0.076]	[0.706]	[0.917]
Control mean	0.020	0.099	44.983	5.038	1.632	0.454
Wave FE	X	X	X	X	X	X
# of observations (Farmer-crop-wave)	4,422	4,422	729	729	4,422	2,253
# of clusters (Farmer group)	70	70	63	63	70	70

Notes: Selected regression coefficients from Equation 2 are presented above, with selected sums of regression coefficients listed below, and included control variables also listed below. Robust standard errors clustered at the farmer group level are in parentheses, and p-values in brackets.

Allowing for the impact of demonstration to vary across veterans and novices does not reveal any significant heterogeneity in impacts on adoption of the new seeds, yield, fertilizer use, or profits (Table 4).

However, the similar resulting adoption of the new seeds caused by demonstration operates differentially across novice and veteran cultivators of these crops. Among novices, demonstration increased cultivation of the promoted crop by 57% relative to the control group; this increase in cultivation among novices is fully joint with increased adoption of the new seeds. The imprecisely estimated, large switch away from paddy as a result of demonstration observed in Table 3 is concentrated among novices, and the confidence interval does not allow us to reject that this decrease in paddy cultivation is fully accounted for by the increase in promoted crop cultivation (Columns 2 and 6, Table 4). While this estimate is also not statistically significant, it provides further suggestive evidence that demonstration enabled novice cultivators to cultivate pulses with new improved seeds, and in turn substitute away from paddy. In contrast, veteran cultivators do not respond to demonstration activities by altering their crop choice.

Taken together, these results suggest that returns to demonstration activities are homogeneous across novices and veterans: the promoted crops are traditionally grown across the study area, everyone in the group knows about the basic agronomy of these crops; instead, farmers are principally learning from demonstration about the new seeds.[19]

4.2 New improved seeds increase productive resilience

Empirical strategy Why would farmers choose to adopt the promoted varieties? The new seeds were developed to be resilient to flood-saline events which are relatively widespread across our study area (Appendix Figure A2), with an average of 25% of farmers reporting experiencing crop losses due to flooding in a given dry season.[20] While we found no average impacts of demonstration on promoted crops' yields or profits in Section 4.1, a simple treatment-on-the-treated back-of-envelope suggested that we may be under-powered to detect impacts given the levels of adoption triggered

[19] We revisit heterogeneity across novices and veterans in impacts of demonstration on adoption of the new seeds in Section 5.2, and present evidence consistent with the notion that novice cultivators were more uncertain about the productivity of the new seeds than veterans. One possible explanation is that novice cultivators were also learning about the productivity of the promoted crops, while veteran cultivators were only learning about the productivity of the new seeds.

[20] As noted above, while the incidence of tidal floods is relatively low during the dry season, their impact on soil salinity is strongest in this season.

by demonstration. The possibility that the estimated null effects on yields and profits may mask impacts sufficiently large to drive adoption is further corroborated by the economically meaningful impacts of demonstration on input use that we observe.

To increase our power to detect any yield impacts of using saline flood-resilient seeds, we restrict our sample to cultivators of the promoted crops and compare yields of farmers who do and do not adopt the improved seeds. Of course, these farmers may differ. We instead make the following identifying assumption: absent the adoption of improved saline flood-resilient seed, the impact of exogenous saline flood events on yields would have been the same across these two groups.

To capture the yield impacts of using the improved seeds over our study period, we estimate:

$$\begin{aligned} Y_{hcvt} = \alpha_{ct} + \beta_1 \text{Improved seed}_{hcvt} + \beta_2 \text{Above median flood}_{hvt} \\ + \beta_3 \text{Improved seed}_{hcvt} \times \text{Above median flood}_{hvt} + \epsilon_{hcvt} \end{aligned} \tag{3}$$

where Y_{hcvt} is household h's yield or profit per hectare for cultivated crop c in survey wave t. Improved seed$_{hcvt}$ indicates that household h adopted improved seeds to cultivate crop c in survey wave t; as the new flood saline-resilient varieties were not introduced until the demonstration year, we restrict our analysis to promoted crops during the demonstration year and one and two years post-demonstration. We use Above median flood$_{hvt}$ as our measure of flood-saline shocks; it is an indicator that takes a value 1 if village v's leave-household-h-out average flood exposure is above the median value in the data, and 0 otherwise. We plot variation in the modal value of above median flood exposure across villages in Appendix Figure A2, and we provide additional details on the construction of above median flood exposure in Appendix A. Given a flood realization, yields will be determined by crop-specific agronomic properties of the improved seed; hence, we control for crop-by-year fixed effects α_{ct}.

Balance In order for our estimates of Equation 3 to capture the impacts of saline floods with and without the adoption of improved seed, it must be the case that saline flood events are exogenous and do not affect farmers' decisions either to cultivate the promoted crops or to adopt the improved seeds. We present the following evidence consistent with these assumptions. First, in Appendix Table A24, we show that saline flood events do not affect farmers' decisions to cultivate promoted crops, adopt improved seed (either unconditional or conditional on cultivating), or use fertilizer. We note that this exogeneity of flooding is consistent with both limited variation within Barisal in probability of flood exposure, and the unexpected nature of flooding in any given year preventing any anticipatory effects of floods on adoption. We present additional evidence of the former in Appendix Table A25: while above median flood exposure strongly predicts own flood exposure, it does not predict flood exposure in the subsequent year.

Results We present estimates of Equation 3 in Table 5; we find that adopting saline-flood-resilient seeds eliminates the adverse impacts of saline floods.

First, we establish large crop damage in the presence of saline floods: for farmers who do not cultivate improved seeds, yields and profits per hectare fall by 19% and 63%, respectively.

Table 5: Saline-resistant improved seed eliminates negative impacts of saline-flood shocks on profits

	Promoted crops, Cultivation = 1, Demo year and one and two years post-demo			
	Yield ('000 BDT/ha)		Profit ('000 BDT/ha)	
	(1)	(2)	(3)	(4)
Improved seed	2.497	0.161	2.363	-0.399
	(1.987)	(2.165)	(2.096)	(2.416)
	[0.209]	[0.941]	[0.260]	[0.869]
Above median flood		-7.383		-6.111
		(3.311)		(3.330)
		[0.026]		[0.067]
Improved seed × Above median flood		5.723		6.874
		(3.558)		(3.952)
		[0.108]		[0.082]
Control mean	36.49	39.78	7.72	9.71
Crop-wave FE	X	X	X	X
# of observations (Farmer-crop-wave)	1,336	1,336	1,336	1,336
# of clusters (Farmer group)	90	90	90	90

Notes: Selected regression coefficients from Equation 3 are presented above, with included control variables listed below. Robust standard errors clustered at the farmer group level are in parentheses, and p-values in brackets.

Second, these adverse impacts of floods go to zero for farmers cultivating improved seeds. The difference in impacts of saline floods for farmers who do and do not cultivate improved seeds is statistically significant for profits, and on the threshold of statistical significance for yields. This reduction in the impacts of saline floods on yields and profits corresponds to an increase in yields and profits per hectare of 8% and 45%, respectively.

This increased productive resilience suggests that farmers may update more positively about the improved seeds in response to demonstration during a flood-saline shock. In Appendix Table A26, we estimate heterogeneity in impacts of village assignment to demonstration with respect to demo year flood-saline shocks, and we find exactly this. Demo year flood-saline shocks increase adoption of improved seeds one year post-demonstration, and this result is driven entirely by villages assigned to demonstration. In contrast, we find no effects on contemporaneous adoption (consistent with the exogeneity documented above) and no effects on adoption two years post-demo (dynamics we rationalize in Sections 4.5 and 6 as caused by the diffusion of information generated by demonstration through farmer groups).

While we presented evidence above that flood-saline shocks are exogenous, the decision to adopt the new seeds may be endogenous to other determinants of yields, which would prevent us from interpreting the increased flood-saline resilience under the new seeds as causal; we consider two possible forms of endogeneity, and we argue our estimate of the productivity impacts of the new seeds is likely conservative. First, the adoption of flood-saline-resilient seeds may be joint with other practices that help combat the negative effects of flood-saline shocks on yields; in this case, we interpret our estimates as the joint effect of these practices and the improved seeds. Second, it may be the case that farmers with higher or lower yields are more likely to adopt improved seeds; our estimates in Table 5 suggest that farmers who do and do not adopt improved seeds experience similar yields and profits absent flood-saline shocks. This alleviates concerns that farmers may be positively selected on yields into adopting the improved seed, and suggest these estimates will not overstate the flood resilient impact of the improved seeds.

4.3 Demonstration persistently increases adoption, but these effects fade over time

Empirical strategy We now study the dynamics of adoption in response to demonstration. We again estimate Equation 1, now allowing for the coefficient β_t on Any demo in village$_v$ to vary across the two post-demonstration years.

Results Rather than an S-curve of technology adoption, estimating the impacts of any demonstration activity over the two dry seasons following the demonstration season reveals an inverted U-curve of adoption of the new seeds. First, any demonstration results in a stark 5.6 percentage point increase in adoption one year post-demonstration. This corresponds to a 155% increase off of 3.6% adoption in the control group, or conservatively representing a 30 percentage point increase in adoption of the new seeds among cultivators of the promoted crops. Second, this effect fades to 2.6 percentage points two years after the demonstration season. This long-term level of adoption is still high relative to the control group: it represents a 118% increase off of a 2.2% adoption in control villages, and conservatively implies a 14 percentage point increase in adoption rate among cultivators of the promoted crops.[21]

These results suggest that adoption of the new seed varieties may be nonmonotonic in learning, as impacts on adoption remain high overall but fall over time. These dynamics are not driven by the control group catch-up, with control group adoption remaining stable and, if anything, falling over time, suggesting demonstration triggered the start of a learning process that had not yet begun in the control group two years post-demonstration. While other mechanisms might explain these results, such as forgetting after demonstration or falling market access to the new seeds, we discuss evidence that learning is responsible for these dynamics in Section 4.5. We further interpret these dynamics and return to the veteran-novice dimension of heterogeneity through the lens of a Bayesian model of learning the returns of a new technology in Section 5.2; we show these falling impacts are driven predominantly by novices, while impacts on adoption by veterans if anything grow over time.

[21] In Appendix Table A14, we show that the decrease in the effect of any demonstration on adoption from one to two years post-demonstration is statistically significant, while we fail to reject that control group adoption does not change from one to two years post-demonstration.

Table 6: Demonstration increases long-run adoption of improved seed, while decentralizing demonstration only increases medium-run adoption

	Promoted crops					
	One year post-demo		Two years post-demo		One and two years post-demo	
	Improved seed					
	(1)	(2)	(3)	(4)	(5)	(6)
Any demo in village	0.056	0.028	0.026	0.022	0.045	0.027
	(0.017)	(0.020)	(0.012)	(0.014)	(0.013)	(0.015)
	[0.001]	[0.167]	[0.026]	[0.106]	[0.000]	[0.068]
Shared demo in village		0.018		0.019		0.019
		(0.022)		(0.016)		(0.017)
		[0.422]		[0.236]		[0.278]
Decentralized demo in village		0.065		-0.006		0.035
		(0.027)		(0.014)		(0.018)
		[0.018]		[0.658]		[0.057]
Control mean	0.036	0.036	0.022	0.022	0.031	0.031
Wave FE	X	X	X	X	X	X
# of observations (Farmer-crop-wave)	3,289	3,289	2,244	2,244	5,533	5,533
# of clusters (Farmer group)	96	96	70	70	96	96

Notes: Selected regression coefficients from Equations 1 and 4 are presented above, with included control variables listed below. Robust standard errors clustered at the farmer group level are in parentheses, and p-values in brackets.

4.4 Decentralizing demonstration increases adoption, but these additional impacts vanish over time

Empirical strategy We now leverage our multi-arm experimental design to estimate the additional impact of shared and decentralized demonstration relative to regular demonstration. Keeping similar notations as above, we now estimate:

$$\begin{aligned}\text{Improved seed}_{hcvt} = \alpha_t + \beta_{1,t}\text{Any demo in village}_v + \beta_{2,t}\text{Shared demo in village}_v \\ + \beta_{3,t}\text{Decentralized demo in village}_v + \epsilon_{hcvt} \quad (4)\end{aligned}$$

where Shared demo in village$_v$ is a binary variable indicating that demonstration implemented in village v followed a shared modality, while Decentralized demo in village$_v$ indicates that a decentralized modality was followed in village v. As we control for Any demo in village$_v$, $\beta_{2,t}$ and $\beta_{3,t}$ capture the additional impact of shared and decentralized demonstration in excess of the impact of regular demonstration (the omitted demonstration category), respectively. We estimate this model both allowing for the impacts of demonstration modalities to vary across post-demonstration years (Columns 2 and 4, Table 6), and pooling across post-demonstration years (Column 6, Table 6).

Balance and attrition In Table 2, we verify balance in key baseline household characteristics for shared and decentralized demonstration relative to regular demonstration. We fail to reject identical sample means across these three treatment arms, with one exception: number of social network connections is significantly lower among farmers in the shared demonstration arm. As social network connections is an important household characteristic with respect to learning, we present estimates of the impact of decentralizing demonstration with number of connections as a control in Appendix Table A21, with no qualitative effects on our results. We also note that 1 of 21 sample mean comparisons across our two balance tables being significant at the 5% level is consistent with expectation.[22]

Results One year post-demonstration, when we found demonstration had the largest impacts on adoption, we find that decentralized demonstration doubles the impact of assigning demonstration on adoption of flood-saline-resilient seeds. Fully decentralized demonstration increases adoption by 6.5 percentage points in excess of a 2.8 percentage point effect of any demonstration, in total a 258% increase over 3.6% adoption in the control (Column 2, Table 6); this corresponds conservatively to a 42 percentage point increase in adoption of the new seeds conditional on cultivation of the promoted crops. In contrast, we fail to reject no additional impact of shared demonstration on adoption; however, the point estimate on shared demonstration corresponds to an economically significant increase in adoption of the new seeds (at 50% of control group adoption), and we revisit the interpretation of this magnitude in Section 6.

[22] As in Section 4.1, in Appendix Table A29 we show that inference in our test of balance is not affected when we implement randomization inference, and we therefore apply robust inference clustered at the village-level. Relatedly, we present estimates of Equation 4 with randomization inference p-values included in Appendix Table A31, and patterns of statistical significance are unaffected. In Appendix Table A8, we establish that attrition one and two years post-demonstration was not significantly affected by decentralizing demonstration.

These additional effects of decentralizing demonstration vanish two years post-demonstration, leading the pooled impact of any demonstration to also fall, as documented in Section 4.3. The impact of regular demonstration is fairly stable over the two post-demonstration years (2.8 and 2.2 percentage points one and two years post-demonstration, respectively; Column 4, Table 6). However, the impact of decentralized demonstration sharply converges to the impact of regular demonstration, with its additional impact fading from 6.5 percentage points in the first year to an insignificant -0.1 percentage points in the second year.

4.5 Discussion of experimental results and alternative mechanisms

In summary, we find that demonstration persistently increased adoption of a productive technology; decentralizing demonstration sharply increased adoption one year post-demonstration, but these additional effects faded two years post-demonstration. In Section 5, we rationalize these results in a model of learning the returns to a productive technology which predicts the possibility of nonmonotonic adoption as a function of learning—decentralized demonstration increased both learning and adoption relative to regular demonstration one year post-demonstration, while learning increased and adoption fell two years relative to one year post-demonstration. In Section 6, we present evidence consistent with this model.

Although our reduced form results one year post-demonstration indicate sharp differences in adoption of the new technology across the decentralized arm and the regular and shared arms, these do not allow us to directly speak to the underlying mechanics of learning for two main reasons (cf. Section 3.2). First, decentralizing demonstration may have increased adoption both by increasing learning from self and learning from others, as both increase in the decentralized arm relative to the regular and shared arms. Second, we cannot tell the role of scale and geographic heterogeneity from these results, as we cannot isolate them from variation in the strength of learning from self and from others across arms. In Section 5.3, we impose restrictions on our model of learning and generate model tests to speak to the relative importance of these mechanisms.

We consider two mechanisms besides learning that could explain the adoption dynamics we observe as a result of any demonstration and across demonstration modalities: seed recycling and supply, and forgetting.[23] We conclude these alternative mechanisms are unlikely to explain our results.

First, rather than learning about the new technology, farmers may have simply recycled seeds from demonstration year production; that is, demonstration acted as a supply shock, rather than triggering learning. In the absence of complementary seed purchase, this would mechanically lead to disadoption of the new seed variety over time. This is a possibility since the promoted varieties are open-pollinated. We leverage the fact that our survey records the sourcing method for each production input to document whether assignment to any demonstration or to different demonstration modalities predicts farmers' reported use of recycled improved seeds (Appendix Table A27). We find no statistically significant effect of assigning demonstration on the use of recycled improved seed, and point estimates are too small to explain the stark dynamics we observe; instead, farmers predominantly source the improved seed from seed multipliers or local markets. Alternatively, while demonstration may have triggered learning, falling impacts over time may have been caused by falling supply rather than learning. In Appendix Table A14, we show there were no significant dynamics of adoption in the control group, suggesting the supply of the new seeds remained stable over time.

Second, while impacts on adoption one year post-demonstration may have been caused by learning, falling impacts two years post-demonstration may have been caused by farmers forgetting about the new seeds. We present three pieces of evidence that forgetting did not cause these dynamics. First, even though the impacts of demonstration fade two years post-demonstration, the level of adoption remains high, with over 31% of farmers cultivating promoted crops in demonstration villages adopting improved seeds two years post-demonstration; forgetting would be more plausible in a context with lower rates of adoption and without farmer groups that regularly discuss the promoted crops. Second, the fall in impacts two years post-demonstration is concentrated in decentralized demonstration villages, so forgetting would have to be concentrated in these villages. Instead, in Table 6, we presented evidence that impacts converged across demonstration modalities, which we attribute to convergence towards similar and high levels of learning facilitated by farmer groups.

[23] In our model in Section 5, we consider other mechanisms closely related to learning—non-Bayesian learning, and forward-looking experimentation one year post-demonstration—and argue these do not affect our conclusions.

That is, learning catch-up, rather than differential forgetting. Consistent with this, in Appendix Table A32, we show that this convergence continues to hold conditional on veteran status.[24] Third, in Section 5.2 we show that selection into adoption of the new seeds improves over time, with veteran cultivators making up a larger share of total adoption two relative to one year post-demonstration. Consistent with these selection dynamics, in Section 7 we show that the impacts of demonstration on profits increased from one-to-two years post-demonstration, indicating increasing learning and at odds with a story in which farmers' knowledge fades over time (Foster and Rosenzweig, 1995).

4.6 Robustness of experimental results

Selection of promoted crops In Section 3.3, we noted that our analysis in Section 4 restricted to promoted crops for power; this would introduce bias if the selection of promoted crops was endogenous to treatment. In Appendix Tables A9 and A10, we show that promoted crops, and the number of promoted crops, are orthogonal to treatment assignment; between 1.61 and 1.88 crops on average were promoted across treatment arms. Additionally, in Appendix Tables A22 and A23, we estimate the impacts of treatment assignment with crop-wave fixed effects to control for the possibility that different treatment arms may have caused different promoted crops to be selected. As we found no evidence of this selection, including these fixed effects does not affect our results.

Placebo crop analysis In Appendix Tables A15 and A16, we estimate the impacts of treatment on placebo crops. Given the positive impacts we estimate in our main analysis, we would be particularly concerned if we observed negative "impacts" on placebo crops, as these would be suggestive of positive selection of promoted crops. Instead, we find some evidence of positive spillovers of demonstration onto adoption of improved seed for placebo crops (although of much smaller magnitudes than our main effects).

[24]Convergence also holds across variation in signals of the profitability of the new seeds—demonstrator profits (Appendix Table A1) and demo year saline-flood shocks (Appendix Table A26) increase adoption of the new seeds one year, but not two years, post-demonstration.

Pre-demo analysis Complementing our balance tests in Section 4, in Appendix Tables A17 and A18, we find no differences in baseline outcomes across treatment arms.

5 Model of learning and technology adoption

In the previous section we presented evidence of an "inverted U-curve" of adoption—demonstration increased medium run adoption, and impacts persisted but faded in the long run. These results are inconsistent with many models of learning and technology adoption: if we believe that farmers continue to learn about the technology over time, these results suggest that adoption itself may not be monotonic in learning.

We therefore lean on theory to answer two questions: first, should we expect adoption to correspond directly to learning? and second, if not, can we infer learning from adoption?

First, in Section 5.1, we model demonstration as improving the precision of farmers' signals of the returns to the demonstrated technology, in response to which farmers update their beliefs about these returns; our setup closely follows Besley and Case (1994).[25] When farmers adopt the demonstrated technology if they expect positive returns, we show that adoption may be non-monotonic in learning; initial learning introduces belief dispersion, which generates noisy adoption. We present an example of this phenomenon in Section 5.2; noisy adoption is concentrated among novice cultivators, consistent with novice cultivators holding less certain prior beliefs which in turn magnifies belief dispersion.

Second, in Section 5.3, we leverage the insight that adoption can still be used to infer the determinants of learning, at least during phases where adoption remains, on average, monotonic in learning. This insight generates tests of learning from scale, geography, and the set of demonstrators, and in turn explains the impacts of demonstration on adoption.

[25] In Appendix C, we present evidence that corroborates predictions from this model of learning: demonstration affects farmers' beliefs that the new seeds are high yielding, which in turn is associated with adoption of the new seeds, and higher experienced profits during the demonstration year affect both beliefs about and adoption of the new seeds.

5.1 Beliefs and technology adoption

Prior beliefs and the introduction of the new seeds Our model begins following the introduction of the new seeds to farmer groups by extension agents at the farmer group meetings. Extension agents can inform farmers about many aspects of the new seeds, such as their flood-saline-resilience, and farmers are already familiar with others, such as the suitability of the pulses to their plots and their household. However, as discussed in Section 2.3, farmers will remain uncertain about the productivity of the new seeds, A. Dissemination through meetings facilitates common initial beliefs across farmers about the seeds' productivity A, as the meetings provide farmers the opportunity to discuss the new seeds with extension agents and each other in advance of any learning from demonstration (Chandrasekhar et al., 2020). We therefore begin with the simplifying assumption that all farmers in the group hold a common normal prior about the returns to the new seeds.

$$A \sim \mathcal{N}(\mu_0, \sigma_0^2) \tag{5}$$

Signals and demonstration During the demonstration year, farmers observe an unbiased signal of the returns to the new seeds, Y. As in Besley and Case (1994) and Beaman et al. (2021), we model demonstration as affecting farmers' technology adoption decisions through the variance of this signal, σ_Y^2. This signal reflects farmers aggregating information about the productivity of the new seeds, gleaned from both their own demonstrations and the demonstrations of other farmers in their village.[26] As a consequence, the variance of the signal σ_Y^2 is affected by the demonstration modality assigned to the village, and may vary across farmers within village depending on who in the village demonstrates the new seeds.

$$Y|A = a \sim \mathcal{N}(a, \sigma_Y^2) \tag{6}$$

Updated beliefs and learning from demonstrations Following the demonstration year, and before the planting season one year post-demonstration, farmers update their beliefs about the returns to the new technology given their prior beliefs

[26]We assume farmers do not learn from other farmers' decisions to demonstrate, as these decisions were made on the basis of common knowledge introduced at the farmer group meetings and therefore do not contain additional information about the new seeds.

and the observed signal Y. We assume farmers update their beliefs following Bayes' rule.[27]

$$A|Y \sim \mathcal{N}(M, \sigma^2) \tag{7}$$

Farmers' updated expected returns to the new technology, $M \equiv \ell Y + (1-\ell)\mu_0$, are a weighted average of their prior expectations and the observed signal. The weight on the observed signal $\ell \equiv \sigma_0^2/(\sigma_0^2 + \sigma_Y^2)$ depends on the precision of the observed signal of the returns to the new seeds. We refer to ℓ as *learning*; consistent with this interpretation, farmers' updated uncertainty, $\sigma^2 \equiv (1-\ell)\sigma_0^2$, decreases with learning. Farmers' updated uncertainty is equal to prior uncertainty with no learning, and falls to 0 with full learning. When farmers do not observe any signal ("no learning"), that is $\sigma_Y^2 = \infty$, then learning $\ell = 0$. Alternatively, when farmers directly observe the returns to the new technology ("full learning"), that is $\sigma_Y^2 = 0$, then learning $\ell = 1$. As learning increases from no learning ($\ell = 0$) to full learning ($\ell = 1$), farmers observe an increasingly precise signal of the returns to the new technology and form increasingly precise updated beliefs.

Note that expected returns M vary across farmers, due to noise in the observed signal Y; as we will see below, this variation is a key determinant of the impacts of learning on adoption of a new technology.

Updated beliefs and adoption one year post-demonstration Not all farmers optimally adopt the new seeds (Suri, 2011); farmers face known idiosyncratic costs of adopting the new seeds $C \sim F_C$, and adopt the new seeds if they expect positive net returns $A - C$.[28] Although we refer to C as "costs", we interpret C as any

[27]There may be many reasons to expect deviations from Bayesian updating. Farmers may ignore that observed information may contain common sources, therefore overweighting positively correlated signals (Chandrasekhar et al., 2020). Alternatively, farmers may underweight signals from their social network (Conlon et al., 2022). In Appendix B, we consider deviations from Bayesian updating in the form of overreaction (underreaction) to the observed signal Y, and show these deviations generate equivalent updated beliefs as a function of the observed signal to reduced (increased) prior certainty.

[28]This model of adoption decisions implicitly imposes that farmers are myopic, that is they do not account for the impacts of their adoption on their future learning when deciding whether to adopt. These unaccounted "returns to experimentation" are largest when uncertainty about the technology is high, and therefore in general are decreasing in learning; in contrast, below we show that belief dispersion is increasing and then decreasing in learning, and therefore offers a parsimonious explanation of nonmonotonic adoption dynamics, the focus of our model.

known idiosyncratic factors that might drive decisions to adopt the new seeds, such as heterogeneous returns to cultivating pulses or to the flood-saline-resilience trait. Letting D indicate that the farmer adopts the technology,

$$\begin{aligned} D &= \mathbf{1}\{\mathbf{E}[A|Y] - C > 0\} \\ &= \mathbf{1}\{M - C > 0\} \end{aligned} \tag{8}$$

Learning and adoption one year post-demonstration Our primary outcome of interest in Section 4, as is common in the literature on technology adoption, is the fraction of farmers who adopt the new seeds, or "population adoption" $\mathbf{E}[D|A = a]$. Population adoption corresponds directly to observed adoption in our experiment; it is the fraction of farmers who adopt the new seeds, conditional on a particular realization of the returns to the new seeds $A = a$. We emphasize the latter point: while farmers' prior beliefs are over many potential realizations of the new seeds, it is one particular realization that is distributed to farmers in our experiment.

We make two additional comments on population adoption. First, population adoption depends on learning, as learning affects both farmers' expectation of the returns to technology given their signal and also the precision of the signal. Second, population adoption is a nonlinear transformation of the farmer's expected returns M, so the population average expected returns $\mathbf{E}[M|A = a]$ are insufficient to characterize population adoption.

Adoption of the new seeds depends on farmers' expected returns and their costs; as the distribution of costs C is fixed, we first characterize the population distribution of expected returns M in order to characterize population adoption.

$$M|A = a \sim \mathcal{N}(\ell a + (1 - \ell)\mu_0, \ell(1 - \ell)\sigma_0^2) \tag{9}$$

We interpret Equation 9 as follows. First, as learning increases from no learning to full learning, the population average expected returns $\mathbf{E}[M|A = a]$ shift monotonically from prior expectations μ_0 to the true returns a. Second, intermediate learning generates "belief dispersion": some farmers observe relatively positive signals, while others observe relatively negative signals. Belief dispersion, that is the population variance of expected returns, is *nonmonotonic*, and specifically is increasing and then decreasing, in learning.

Proposition 1. *Population adoption is increasing in population average expected returns, and is increasing (decreasing) in belief dispersion if marginally adopting farmers have higher (lower) average expected returns than average farmers.*

Corollary 1. *Population adoption as a function of learning may be increasing, decreasing, increasing then decreasing, or decreasing then increasing.*

Proof. See Appendix B. □

5.2 Novice and veteran adoption dynamics through the lens of the model

Key parameters underlying learning govern whether learning leads to increased adoption of a technology through an S-curve, with steadily growing adoption, or an inverted U-curve, with increasing and then decreasing adoption. We can see this in Equation 9. Population average expected returns will increase sharply with learning, generating the well-known "S-curve", when true returns a are much larger than prior expectations μ_0. On the other hand, belief dispersion will increase and then decrease sharply with learning, generating an "inverted U-curve", when prior uncertainty σ_0^2 is high.

We apply these predictions to our finding that veteran cultivators did not experience larger impacts of demonstration on adoption of the new seeds than novice cultivators, although veteran cultivators were more likely to adopt the new seeds with or without demonstration; should we have expected to find heterogeneous impacts? These results pooled across one and two years post-demonstration, and therefore pooled across periods during which learning was low (one year post-demonstration) and high (two years post-demonstration). We hypothesize that novice cultivators have greater prior uncertainty than veteran cultivators: as novice cultivators have less experience with dissemination of improved seeds targeting the promoted crops, we expect them to hold more uncertain prior beliefs. This additional prior uncertainty would lead to stronger inverted-U dynamics among novice cultivators, with relatively larger impacts one year post-demonstration than two years post-demonstration.

To test these predictions, we separately plot adoption of novice and veteran cultivators as a function of learning induced by treatment assignment: in control (no learning), in treatment one year post-demonstration (intermediate learning), and in treatment two years post-demonstration (full learning). We find evidence of stronger inverted-U dynamics among novice than veteran cultivators in Figure 5a, consistent with greater prior uncertainty.

We test for stronger inverted-U dynamics among novice than veteran cultivators formally by re-estimating Equation 2 in Figure 5b, which estimates heterogeneous impacts of demonstration with respect to veteran cultivator status, but with added interactions with two years post-demonstration. The change in the effect of demonstration on adoption, from one to two years post-demonstration, is significantly more negative for novices than for veterans. While average effects on adoption across novices and veterans are the same over our two post-demonstration years, the one year post-demonstration effect on adoption is driven primarily by novices, while the two year post-demonstration effect on adoption is driven primarily by veterans.

5.3 Learning from self and learning from others

We have established that technology adoption need not be monotonic in learning, even for Bayesian farmers; if so, what can technology adoption tell us about learning? We observe, building on Equation 9, that population technology adoption, conditional on prior beliefs, the true returns to the technology, and costs of adoption, is a smooth function of learning. As a consequence, small exogenous increases in learning cause proportional changes in technology adoption. Relative impacts on adoption are therefore informative of relative impacts on learning. Our experimental results one year post-demonstration present an opportunity to estimate such relative impacts, revealing initial learning generated by demonstration. We use these impacts to propose tests for three key mechanisms underpinning learning proposed in (Foster and Rosenzweig, 1995): learning from self and learning from others, learning from scale, and learning from geographic heterogeneity.

We generate tests for each of these three mechanisms of learning by considering testable restrictions they imply for the variance of the signal observed by farmers, σ_Y^2; we summarize predictions under each mechanism in Table 7. Across all three tests, and consistent with our conceptual framework in Section 3.2, we commonly impose that farmers optimally aggregate signals across their own demonstration and demonstrations in their social network.[29] This follows Foster and Rosenzweig (1995)

[29] While farmers may learn from demonstrations outside their social network, particularly once

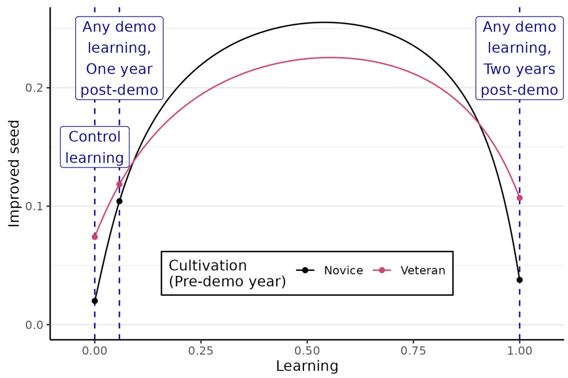

(a) Model implied learning and adoption

	Promoted crops, One and two years post-demo		
	Improved seed		
	(1)	(2)	(3)
Intercept	0.036	0.032	0.023
	(0.013)	(0.014)	(0.011)
	[0.005]	[0.022]	[0.047]
Veteran cultivator			0.067
			(0.029)
			[0.022]
Two years post-demo	-0.014	-0.009	-0.005
	(0.013)	(0.013)	(0.012)
	[0.295]	[0.483]	[0.673]
Veteran cultivator × Two years post-demo			-0.027
			(0.026)
			[0.303]
Any demo in village	0.056	0.071	0.077
	(0.017)	(0.020)	(0.018)
	[0.001]	[0.000]	[0.000]
Veteran cultivator × Any demo in village			-0.045
			(0.035)
			[0.195]
Two years post-demo × Any demo in village	-0.030	-0.046	-0.057
	(0.018)	(0.019)	(0.019)
	[0.093]	[0.015]	[0.003]
Veteran cultivator × Two years post-demo × Any demo in village			0.074
			(0.034)
			[0.031]
Pre-demo year sample only		X	X
# of observations (Farmer-crop-wave)	5,533	4,422	4,422
# of clusters (Farmer group)	96	70	70

(b) Adoption dynamics

Figure 5: Heterogeneous adoption dynamics suggest higher returns to adoption and more precise prior beliefs for veterans

Notes: In Panel (b), regression coefficients from Equation 2, with included interactions with an indicator for two years post-demonstration, are presented, with included control variables listed below. Robust standard errors clustered at the farmer group level are in parentheses, and p-values in brackets. In Panel (a), average adoption of improved seed conditional on veteran status, survey wave, and treatment, is graphed against average learning implied by the structural model estimated in Appendix D. Fit curves correspond to associated predicted adoption as a function of learning under alternative values of prior beliefs and true returns to the improved seed that match adoption conditional on veteran status.

and Beaman et al. (2021), among others, and additional demonstrations improve the precision of farmers' signals of the returns to the technology.

Learning from self and learning from others First, farmers may learn from themselves and from others. We suppose that farmers observe independent signals of

information has had sufficient time to diffuse beyond immediate connections, we build on a large literature on technology adoption in agriculture that provides evidence that learning from others is predominantly social and that social network connections provide a reasonable proxy for opportunities for social learning (e.g., Conley and Udry, 2010).

Table 7: Tests of determinants of learning

Learning from ...	Demo	# of connections to demo	Shared demo in village	Decentralized demo in village
self	+	·	·	·
others	·	+	·	·
scale	·	·	−	−
geographic heterogeneity	·	·	−	0

Notes: Predicted signs of coefficients in Equation 10, under different forms of learning, are presented above, with · when there is no prediction.

the returns to the new seeds from each demonstration; however, farmers may observe a more precise signal of the returns to the new seeds from their own demonstration. We consider two possible sets of explanations of this difference in precision. The first is technological—if returns are highly idiosyncratic across plots or heterogeneous farming practices, then the returns under one farmer's demonstration may not be informative for other farmers. The second is social—communication is inherently noisy and slow, and farmers may not immediately and fully learn from others' demonstrations.

In this case, the impact of a farmer demonstrating on adoption, relative to the impact of a demonstration in that farmer's social network, equals the relative strength of learning from self compared to learning from others.

Learning from scale Second, farmers may learn from the scale of demonstration. While farmers observe independent signals from each demonstration, the precision of each signal may grow with the size of the demonstration plot—farmers better observe the returns to the technology on larger demonstration plots.[30] In one extreme, learning would depend only on the scale of demonstration, rather than the number of demonstration plots; our results in Section 4.4 strongly reject this. We consider testing the other extreme, of any scale effects.

We leverage the observation that moving from regular to shared to decentralized demonstration reduces the scale of demonstration, conditional on whether a farmer demonstrates and the number of demonstrations in the farmer's social network. Un-

[30]In Foster and Rosenzweig (1995), this growing precision is microfounded by each unit of area of a demonstration plot generating a signal, and imperfect correlation of these signals within demonstration plot.

der learning from scale, shared and decentralized demonstration therefore *decrease* learning conditional on whether a farmer demonstrates and the number of demonstrations in the farmer's social network.

Learning from geographic heterogeneity Third, farmers may learn from geographic heterogeneity, that is observing the new seeds grown under different agronomic conditions. While each demonstration generates an independent signal of the returns to the new seeds under regular (as there is only one demonstration) or decentralized demonstration, all demonstrations occur under identical agronomic conditions under shared demonstration; shared demonstration only decentralizes socially, while decentralized demonstration equally decentralizes socially and geographically.

As a consequence, moving to shared demonstration, conditional on whether a farmer demonstrates and the number of demonstrations in the farmer's social network, reduces the geographic heterogeneity under which demonstrations occur, and therefore *decreases* learning.

6 Mechanisms underlying experimental results

In the previous section, we presented a model of learning and technology adoption that rationalized the "inverted U-curve" of adoption we observed, and concluded by deriving testable predictions (Table 7) of elements of the conceptual framework in Section 3.2 that motivated our experiment—the relative strength of learning from self and learning from others, learning from scale, and learning from geographic heterogeneity. In this section, we test those predictions—we find that just over half of learning is social, and fail to find evidence of learning from scale or learning from geographic heterogeneity.

Empirical strategy We implement the tests of the determinants of learning suggested by our model: the tests depend on the effect of demonstrating and of connections to demonstrators, and the effects of decentralizing demonstration conditional on demonstrating and connections to demonstrators. Since farmers converge to similar levels of learning across experimental arms two years post-demonstration, we estimate these effects in the first post-demonstration year before information could diffuse beyond demonstrators' networks; moreover, it is one year post-demonstration that we expect monotonic adoption with respect to learning at the initial learning generated by demonstration. In practice, we estimate augmented versions of Equation 4:

$$\begin{aligned}\text{Any improved seed}_{hcv2} = \alpha &+ \beta_1 \text{Any demo in village}_v \\ &+ \beta_2 \text{Shared demo in village}_v + \beta_3 \text{Decentralized demo in village}_v \\ &+ \gamma_1 \text{Demo}_{hcv} + \gamma_2 \text{\# of connections to demo}_{hcv} + \epsilon_{hcv} \quad (10)\end{aligned}$$

Demo_{hcv} in a binary variable that indicates whether household h demonstrated the improved seed for promoted crop c during the demonstration season, while $\text{\# of connections to demo}_{hcv}$ captures the number of baseline social connections household h had to households who demonstrated the improved seeds for promoted crop c during the demonstration season.

We interpret the coefficients on Demo_{hcv} and $\text{\# of connections to demo}_{hcv}$, γ_1 and γ_2, as the effects of demonstrating and of connections to demonstrators on adoption of the improved seed; as the choice to demonstrate, and as a consequence connections to demonstrators, are not random, we present evidence below that these coefficients do reflect causal impacts. Correspondingly, we interpret the coefficients on $\text{Shared demo in village}_v$ and $\text{Decentralized demo in village}_v$, β_2 and β_3, as the effects of decentralizing demonstration conditional on demonstration and connections to demonstrators.

We present estimates of Equation 10 in Column 1 of Table 8.

Learning from self and from others First, we find strong evidence of both learning from self and learning from others. Demonstrating increases the probability of adopting the new seeds by 12.7 percentage points, while one additional connection to a demonstrator increases the probability of adopting the new seeds by 1.7 percentage points; these effects are precisely estimated. As described in Section 5.3, the relative magnitudes of these effects approximates the relative strength of learning from self compared to learning from others; they imply learning from self is 7 times stronger than learning from an additional social connection. Scaling this average relative strength by demonstrators' average number of social connections (10.61) implies that 59% of learning is social.

Our estimates are comparable to existing estimates of the relative strength of learning from self and learning from others in the context of technology adoption in

Table 8: Impacts of decentralizing demonstration are fully explained by demonstration and connections to demonstrators

	One year post-demo				Two years post-demo
	Promoted crops			Placebo crops	Promoted crops
	Improved seed				
	(1)	(2)	(3)	(4)	(5)
Intercept	0.032	0.029		0.004	0.023
	(0.014)	(0.017)		(0.002)	(0.010)
	[0.023]	[0.093]		[0.045]	[0.024]
Any demo in village	0.025	0.024		0.031	0.020
	(0.026)	(0.025)		(0.013)	(0.014)
	[0.334]	[0.335]		[0.015]	[0.166]
Shared demo in village	0.000	0.000		-0.011	0.012
	(0.028)	(0.028)		(0.016)	(0.016)
	[0.994]	[0.996]		[0.508]	[0.455]
Decentralized demo in village	0.008	0.008		-0.016	-0.018
	(0.031)	(0.031)		(0.017)	(0.017)
	[0.804]	[0.792]		[0.330]	[0.285]
Demo	0.127	0.128	0.102	0.017	0.049
	(0.032)	(0.032)	(0.043)	(0.012)	(0.018)
	[0.000]	[0.000]	[0.020]	[0.152]	[0.007]
# of connections to demo	0.017	0.017	0.026	0.000	0.001
	(0.007)	(0.007)	(0.009)	(0.003)	(0.003)
	[0.013]	[0.015]	[0.006]	[0.901]	[0.753]
Veteran cultivator		0.029			
		(0.016)			
		[0.079]			
# of connections		0.000			
		(0.002)			
		[0.956]			
Farmer FE			X		
Pre-demo year sample only	X	X	X	X	X
# of observations (Farmer-crop-wave)	2,223	2,223	2,223	3,442	2,199
# of clusters (Farmer group)	70	70	70	70	70

Notes: Regression coefficients from Equation 10 are presented above in Column 1, with variants on Equation 10 in Columns 2 through 5, and with included control variables listed below. Robust standard errors clustered at the farmer group level are in parentheses, and p-values in brackets.

agriculture. Foster and Rosenzweig (1995) estimate the effect of others' experiences in one's village on profitability and initial adoption of improved seeds is between 1.2 and 2.7 times as large as the effect of own experience; these estimates imply between

54% and 73% of learning is social, qualitatively similar to our estimate of 59%. Beaman et al. (2021) estimate the effect of demonstration and number of connections to demonstrators on adoption of pit planting, and find demonstration increases adoption by between 12 and 19 times as much as connections to demonstrators; our estimate of 7 suggests social learning is somewhat stronger in our context, perhaps due to the social learning facilitated by farmer group participation in our context, or the greater importance of experiential learning for pit planting.

Differential selection into demonstration and social relationships Since household participation in demonstration was not randomly assigned, we perform four robustness checks to assess potential selection bias in our estimates of the impacts of demonstrating and of connections to demonstrators.

First, we fail to reject a joint test that demonstrators and non-demonstrators are observably similar, on average (Appendix Table A11). However, we do find that demonstrators are statistically significantly more likely to be veteran cultivators and to have more social connections; we include these characteristics as controls when estimating the effects of demonstrating and of connections to demonstrators in Column 2 of Table 8, and we find this does not affect our estimates.[31] We note that the average similarity of demonstrators and non-demonstrators is consistent with the relative homogeneity of farmers within farmer groups.

Second, we take advantage of the observation that multiple crops were demonstrated in each village, and each household was intended to demonstrate at most one variety, and control for household fixed effects in Column 3 of Table 8. This specification accounts for potential differences in unobserved characteristics across demonstrators and non-demonstrators, and across demonstrators across arms, that could bias our estimates of learning from self and from others. We find that the coefficients on demonstrating and on connections to demonstrators are not meaningfully affected by the inclusion of household fixed effects.

[31] While it may be surprising that demonstrators and non-demonstrators are observably similar, we note that this masks heterogeneity across demonstration modalities (Appendix Table A12). The broad coverage of decentralized demonstration, and the requirement that demonstrators have adjacent plots for shared demonstration, limited scope for selection. As a consequence, regular demonstrators appear significantly positively selected relative to shared and decentralized demonstrators. The average similarity of demonstrators and non-demonstrators is therefore a consequence of the much larger number of demonstrators under shared and decentralized demonstration, who also drive our estimates of the effect of demonstrating and of connections to demonstrators.

Third, since social network relationships were not randomly assigned, we leverage non-promoted placebo crops to perform a placebo test: if we are truly measuring the impact demonstrating and connections to demonstrators have on adoption of the promoted seeds through learning and net of any other unobserved confounder, then we should observe that they do not affect the adoption of improved seeds for non-promoted crops. Results from this placebo test in Column 4 of Table 8 confirm this notion.

Lastly, we compare the effects of demonstrating and of connections to demonstrations one year post-demonstration and two years post-demonstration; in Section 4.5, we argued that two years post-demonstration, information generated by demonstration had diffused through the farmer groups. As a consequence, we should not expect any persistent effects of demonstrating or of connections to demonstrators two years post-demonstration. Consistent with this, in Column 5 of Table 8, we find the effect of demonstrating falls by two thirds, and the effect of connections to demonstrators all but vanishes, two years post-demonstration.

Learning from scale and geographic heterogeneity In contrast to the strong evidence we find for both learning from self and from others, we fail to find evidence of learning from scale or geographic heterogeneity in our context. In Column 1 of Table 8, we find no evidence that shared or decentralized demonstration affect adoption of the new seeds, after accounting for the average effects of demonstrating and of connections to demonstrators. Interpreting these null effects through the lens of our model predictions (Table 7) leads us to fail to reject the absence of learning from scale and from geographic heterogeneity in our context, as they imply that decreasing the scale of demonstration and reducing the geographic heterogeneity under which demonstrations occur do not significantly reduce learning.

Learning from scale and geographic heterogeneity may vary across technologies and approaches to demonstration. In fact, aspects of the demonstrated technology we study, and the context we study it in, may have resulted in weaker learning from scale and geographic heterogeneity than would manifest for other technologies and contexts.

Learning from scale, for instance, may only be of consequence when demonstration plots are meaningfully smaller than the typical scale of cultivation. Even under fully decentralized demonstration, demonstration “minikits” provided sufficient seed and fertilizer to cultivate 0.03 hectares—this would enable cultivating an area roughly 17 meters by 17 meters, almost one third of the average area cultivated for pulses conditional on cultivation and comparable to the size of a typical plot. If demonstration "minikits" had been much smaller than the inputs required for traditional scales of cultivation, learning from scale may have been first order—we consider this possibility through counterfactuals in a structurally estimated model of experiential and social learning in Section 7.

Learning from geographic heterogeneity will depend crucially on whether observing the new seeds under geographically heterogeneous conditions is informative about their returns. For instance, if farmers were uncertain about the flood saline-resilience of the new seeds (something we argued is unlikely to be the case in Section 2.3), and there was substantial local heterogeneity in exposure to flood saline-shocks, we would expect farmers would learn from observing the new seeds both with and without flood saline-shocks. In contrast, if farmers' experiences with the new seeds, rather than direct observation of cultivation of the new seeds, are most informative about the relative productivity of the new varieties, then geographically decentralizing demonstration will not affect learning. In this case, communication between farmers about their experiences will drive learning (i.e., wisdom of the crowd); our dynamic results, that demonstrating and connections to demonstrators have large effects on adoption of the new seeds one but not two years post-demonstration, are consistent with communication, rather than observation, driving learning from others.

7 Valuing learning, demonstration, and decentralization

In the previous section, we established that our results are consistent with a model of experiential and social learning, in which the impacts of learning on adoption reflect both the convergence of average beliefs to the true returns and belief dispersion on the learning path. This finding implies that adoption and learning are not equivalent in our context, motivating us to revisit our estimates of the impacts of demonstration and decentralizing demonstration in Section 4. In Section 7.1, we define the gains from demonstration. In Appendix D, we describe and implement an approach to estimate a structural model of experiential and social learning. In Section 7.2, we apply our estimated structural model to these definitions to answer three questions: what are the gains from demonstration, both ex-post (for the realized technology) and ex-ante? what are the gains from decentralizing demonstration? and how would counterfactual learning environments impact the gains from decentralization?

7.1 Ex-ante and ex-post profits

Learning and the gains from demonstration We distinguish between two notions of the gains from demonstration, the ex-ante and ex-post gains from demonstration. The ex-post gains from demonstration correspond to the observed impacts of demonstration on farmer profits, which depend on the promoted new seeds. We define expected ex-post profits to be

$$\pi^{\text{ex-post}} = \mathbf{E}[D(A-C)|A=a] \tag{11}$$

Ex-post profits condition on the "realized" new seeds $A = a$, and are equal to the true returns to the technology net of costs for adopters.

The ex-post gains from demonstration are equal to the difference in ex-post profits across treatment arms. The ex-post gains from demonstration can be estimated as effect of treatment on profits, that is the "observed" gains from demonstration; our estimates of the impacts of any demonstration on profits in Table 3 are the observed gains from demonstration, averaged over one and two years post-demonstration.

Alternatively, the ex-ante gains from demonstration correspond to farmers' expected gains from demonstration. The ex-ante gains from demonstration are determined by farmers' prior beliefs; they are the equal to expected ex-post gains from demonstration, with expectation taken with respect to farmers prior beliefs. They are equal to the difference in ex-ante profits across treatment arms, that is farmers' expected profits under their prior beliefs as a function of learning. We define ex-ante profits to be

$$\pi^{\text{ex-ante}} = \mathbf{E}[D(A-C)] \tag{12}$$

In contrast to ex-post profits, ex-ante profits do not depend on the "realized" new seeds, but on farmers' prior beliefs.[32]

[32] In Appendix D.4, we derive expressions for ex-ante and ex-post profits under our parametric assumptions in Appendix D.1; these expressions are applied in Section 7.2 to calculate ex-ante and ex-post gains from demonstration.

Farmers' adoption decisions maximize ex-ante profits, rather than ex-post profits—as a consequence, while ex-ante profits are strictly increasing in learning, ex-post profits may actually fall with learning under particular realizations of the technology.

7.2 Estimated gains from demonstration and decentralization

We report estimates of the gains from demonstration in Appendix Table A13, focusing on three comparisons: the gains from any demonstration relative to control, the gains from shared demonstration relative to regular demonstration, and the gains from decentralized demonstration relative to regular demonstration. For each comparison, we calculate the ex-post gains from demonstration, which we compare to observed gains from demonstration, and the ex-ante gains from demonstration, both one and two years post-demonstration. In addition, we consider gains from demonstration one year post-demonstration under three counterfactual learning environments—no social learning, maximal scale effects in learning, and initializing learning at 0.5.

Ex-ante gains from demonstration We begin by calculating the ex-ante gains from demonstration. One year post-demonstration, averaged across demonstration modalities, farmers expect average gains of 2,280 BDT/ha (approximately 5% of average pulses yields). These gains are large, as a consequence of the high uncertainty of farmers' prior beliefs—farmers ascribe a low likelihood to transformative productivity gains from the new seeds, but this possibility generates large expected gains. Scaled by the total number of farmers per farmer group (25) and the average area cultivated conditional on cultivation (0.135ha), they correspond to 7,700 BDT expected gains from demonstration one year post-demonstration, or roughly 1.8 times the cost of the demonstration kit. These gains vary meaningfully across treatment arms—they range from 0.7 times cost under regular demonstration, to 1.6 times cost under shared demonstration, to 3.2 times cost under decentralized demonstration.

Two years post-demonstration, our estimate of farmers' high prior uncertainty about the returns to the new seeds leads to very large gains from demonstration under full learning. Concretely: farmers expect, with modest probability, new seeds that will enable them to fully replace their dry season paddy cultivation with pulses—a "*Green Revolution of lentils*."

To conservatively evaluate the efficiency of demonstration, we focus on the ex-ante *demonstration externality* one year post-demonstration. In our conceptual framework, we noted that demonstration could, somewhat narrowly, be interpreted as a subsidy implemented as quota; we therefore evaluate its efficiency as a Pigouvian subsidy by comparing the learning externality generated by demonstration to the cost of the demonstration kit. We calculate this externality by comparing the ex-ante gains from demonstration to the gains in a model without social learning ($\omega = 0$). On average, 50% of the gains from demonstration are the learning externality, similar to our calculation in Section 6 that 59% of learning one year post-demonstration is social. As a result, the demonstration externality is smaller than the cost of the demonstration kit under regular or shared demonstration, but is 70% larger than the cost of the demonstration kit under decentralized demonstration.

Ex-post and observed gains from demonstration We next calculate the ex-post gains from demonstration. One year post-demonstration, averaged across demonstration modalities, farmers experienced average *losses* of 1,670 BDT/ha. Farmers update their beliefs substantially in response to observed noisy signals of the productivity of the new seeds as a consequence of the high uncertainty of their prior beliefs; the resulting "noisy adoption" generates losses that are of comparable magnitude to the ex-ante gains from demonstration. Noisy adoption is part and parcel of learning for the realized new seeds given farmers' highly uncertain priors; as a consequence, decentralizing demonstration generates correspondingly larger ex-post losses from demonstration.

Two years post-demonstration, farmers have learned the true returns to the new seeds, and ex-post gains from demonstration are positive. These returns are sufficiently higher than farmers' prior expectations that adoption of the new seeds is significantly higher two years post-demonstration than under farmers' prior expectations. However, the farmers induced to adopt the new seeds through learning were marginal adopters at prior expected returns; as a consequence, the ex-post gains from demonstration are small two years post-demonstration for the realized new seeds.

We note a key distinction between the ex-ante and ex-post gains from demonstration.

The ex-post gains from demonstration are descriptive and internally valid. They correspond to observed impacts on profits of demonstration of the realized new seeds.

Consistent with this, observed dynamic impacts of demonstration and decentralizing demonstration on profits are in general of similar magnitudes and signs to the estimated ex-post gains from demonstration. This match is particularly striking, as dynamic impacts of demonstration on profits were not targeted in our structural estimation. Moreover, it implies that noisy adoption along the learning path can quantitatively explain how demonstration reduced profits one year post-demonstration, but instead increased profits two year post-demonstration, with the difference statistically significant.

In contrast, the ex-ante gains from demonstration are normative and externally valid. The ex-ante gains from demonstration are externally valid, in that they do not depend on the realized new seeds and are applicable to the distribution of alternative technologies that farmers expect to be introduced to them by the agricultural extension service. The ex-ante gains from demonstration are normative, in that they correspond to farmers' valuation of demonstration.

The ex-ante and ex-post gains from demonstration are complementary measures; while the ex-post gains are appropriate for interpreting the impacts of the realized new seeds, the ex-ante gains are appropriate for valuing demonstration as a policy.

Gains from demonstration and decentralizing demonstration under counterfactual learning environments We conclude our analysis of our structural model by evaluating the gains from demonstration under two counterfactual learning environments: when learning is only from scale, and under improved communication from the extension service.

First, we calculate the gains from decentralizing demonstration when learning is only from scale. Specifically, we impose that farmers learn exactly the same when they observe all demonstrations, regardless of the number of demonstrations.[33] This contrasts with our assumed structural model, which imposes that farmers learn exactly the same when they observe one demonstration, regardless of the scale of that demonstration, consistent with our reduced form results in Section 6.[34] Even at this

[33]We implement this by dividing the precision of the signal generated by each demonstration (η) by the number of demonstrators in each village-crop, reducing the strength of learning when there are multiple demonstrators proportionally to the number of demonstrators.

[34]Each of these approaches corresponds to an extreme parameter value in Foster and Rosenzweig (1995). The counterfactual of learning only from scale can be microfounded by no within-demonstration plot correlation in signals generated by segments of the demonstration plot, while our structural model can be microfounded with full within-demonstration correlation in signals.

counterfactual of learning only from scale, we find that fully decentralized demonstration increases the ex-ante gains from demonstration by 23%. While this counterfactual holds fixed the amount of information generated by regular and decentralized demonstration, our parameter estimates imply decreasing returns to learning. As decentralizing demonstration also decentralizes learning, it increases the average returns to learning, and therefore the gains from demonstration.[35]

Second, we calculate the gains from demonstration and decentralizing demonstration under improved communication from the agricultural extension service. We consider a counterfactual in which agricultural extension services nuance their communication with farmers, differentiating between seed varieties that incrementally improve on existing varieties and should induce marginal farmers to adopt ("Customized lentils"), and seed varieties that transform production under a particular crop and should induce large scale transformation of cropping systems ("This is the Green Revolution of lentils"). We arbitrarily impose this communication is sufficiently precise that learning starts at 0.5; this is equivalent to initial communication from the agricultural extension service generating roughly as much learning as two demonstrations by oneself. Under this counterfactual, the gains from demonstration fall by 77%, as a consequence of decreasing returns to learning (there is less left to learn). In addition, only the ex-ante gains from decentralized demonstration remain close to the cost of the demonstration kit, at 0.7 times the cost. However, the ex-post gains from demonstration one year post-demonstration become positive, as belief dispersion is decreasing in learning at higher initial learning (farmers make fewer mistakes). Improved communication from the agricultural extension service reduces noisy adoption under the incrementally transformative agricultural technologies key to agricultural productivity growth, and as a consequence is complementary to decentralized approaches to demonstration that accelerate learning.

8 Conclusion

In this paper, we present experimental evidence that decentralizing demonstration at no additional cost accelerates learning about new technologies. Traditional approaches provide a large amount of demonstration resources to carefully selected

[35] As a consequence, this result could empirically differ in contexts where the returns to learning are convex, as in Beaman et al. (2021).

farmers, and rely on the information generated by their demonstrations to diffuse through their communities. We show that, as an alternative to traditional demonstration, offering the same total demonstration resources to many interested farmers accelerates learning. As more farmers learn by demonstrating, experiential learning increases; as farmers are connected with more demonstrating farmers, information spreads further.

Decentralizing demonstration leads to a proliferation of demonstrations while holding fixed the aggregate scale of demonstration—each individual demonstration plot gets smaller and smaller. Model-suggested tests dismiss a role of the scale of individual demonstrations. Complementary counterfactual simulations suggest that decentralizing demonstration may be optimal even in the limit where shrinking demonstrations proportionately reduces the information they generate. Taken together, these results lend support to the notion that learning in agriculture is all about the experience generated by demonstration and communication, regardless of the scale of these individual demonstration efforts.

Our finding that decentralizing demonstration accelerates learning complements and somewhat reconciles recent empirical and theoretical work on learning in networks. In one class of models, increasing the number of demonstrations is more effective at triggering learning than carefully choosing the right demonstrators (Akbarpour et al., 2020). In other classes, choosing the wrong demonstrators can instead lead to persistently lower levels of adoption (Beaman et al., 2021), especially when shame can slow learning (Banerjee et al., 2021). We present experimental evidence that simply letting anyone interested gain first-hand experience with a new technology increases learning without requiring operationally costly network-based targeting.

Finally, in our context of farmers learning the returns to new flood-saline-resilient seeds, learning and mistakes go together hand-in-hand. While the new varieties increase productive resilience, it is not optimal for all farmers to cultivate these new varieties. This context leads to an inverted U-curve of technology adoption, rather than the oft-cited S-curve, as adoption mistakes are increasing and then decreasing in learning. More precise communication strategies from the agricultural extension service would reshape these adoption dynamics by reducing adoption mistakes on the learning path, and may be worthwhile complements to providing farmers with first-hand experience of new technologies. These results and their implications are amplified by global climate change and the increased uncertainty it generates, as both the challenge of and returns to developing and disseminating customized agricultural technologies grow larger and larger.

References

Agness, Daniel, Travis Baseler, Sylvain Chassang, Pascaline Dupas, and Erik Snowberg, "Valuing the Time of the Self-Employed," 2020.

Akbarpour, Mohammad, Suraj Malladi, and Amin Saberi, "Just a few seeds more: value of network information for diffusion," 2020.

Andrews, Isaiah, Matthew Gentzkow, and Jesse M Shapiro, "Measuring the Sensitivity of Parameter Estimates to Estimation Moments," *The Quarterly Journal of Economics*, 2017, p. qjx023.

_ , _ , and _ , "On the informativeness of descriptive statistics for structural estimates," *Econometrica*, 2020, *88* (6), 2231–2258.

Banerjee, Abhijit, Arun G Chandrasekhar, Esther Duflo, and Matthew O Jackson, "The diffusion of microfinance," *Science*, 2013, *341* (6144), 1236498.

_ , Emily Breza, Arun G Chandrasekhar, and Benjamin Golub, "When less is more: Experimental evidence on information delivery during India's demonetization," 2021.

Beaman, Lori, Ariel BenYishay, Jeremy Magruder, and Ahmed Mushfiq Mobarak, "Can network theory-based targeting increase technology adoption?," *American Economic Review*, 2021, *111* (6), 1918–43.

Besley, T. and A. Case, "Diffusion as a Learning Process: Evidence from HYV Cotton," Papers 174, Princeton, Woodrow Wilson School - Development Studies 1994.

Carter, Michael, Rachid Laajaj, and Dean Yang, "Subsidies and the african green revolution: Direct effects and social network spillovers of randomized input subsidies in mozambique," *American Economic Journal: Applied Economics*, 2021, *13* (2), 206–229.

Chandrasekhar, Arun G, Esther Duflo, Michael Kremer, João F. Pugliese, Jonathan Robinson, and Frank Schilbach, "Blue Spoons: Sparking Communication About Appropriate Technology Use," Working Paper 30423, National Bureau of Economic Research September 2022.

_ **, Horacio Larreguy, and Juan Pablo Xandri**, "Testing models of social learning on networks: Evidence from two experiments," *Econometrica*, 2020, *88* (1), 1–32.

Conley, Timothy G and Christopher R Udry, "Learning about a new technology: Pineapple in Ghana," *The American Economic Review*, 2010, pp. 35–69.

Conlon, John J, Malavika Mani, Gautam Rao, Matthew W Ridley, and Frank Schilbach, "Not Learning from Others," 2022.

Dar, Manzoor H, Alain De Janvry, Kyle Emerick, David Raitzer, and Elisabeth Sadoulet, "Flood-tolerant rice reduces yield variability and raises expected yield, differentially benefitting socially disadvantaged groups," *Scientific reports*, 2013, *3* (1), 3315.

_ **,** _ **,** _ **, Erin M Kelley, and Elisabeth Sadoulet**, "Casting a Wider Net: Sharing Information Beyond Social Networks," 2020.

Dasgupta, Susmita, Md. Moqbul Hossain, Mainul Huq, and David Wheeler, "Climate change, groundwater salinization and road maintenance costs in coastal Bangladesh," *Policy Research working paper*, 2014, *WPS 7147.*

Duflo, Esther, Daniel Keniston, Tavneet Suri, and Céline Zipfel, "Chat Over Coffee? Diffusion of Agronomic Practices Through Information Networks in Rwanda," 2022.

Emerick, Kyle, Alain de Janvry, Elisabeth Sadoulet, and Manzoor H Dar, "Technological innovations, downside risk, and the modernization of agriculture," *American Economic Review*, 2016, *106* (6), 1537–61.

Foster, A. and M. Rosenzweig, "Learning by doing and learning from others: Human capital and technical change in agriculture," *Journal of political Economy*, 1995, pp. 1176–1209.

Gollin, Douglas, Casper Worm Hansen, and Asger Mose Wingender, "Two blades of grass: The impact of the green revolution," *Journal of Political Economy*, 2021, *129* (8), 2344–2384.

Griliches, Zvi, "Hybrid corn: An exploration in the economics of technological change," *Econometrica, Journal of the Econometric Society*, 1957, pp. 501–522.

Janvry, Alain De, Elisabeth Sadoulet, and Tavneet Suri, "Field experiments in developing country agriculture," in "Handbook of Economic Field Experiments," Vol. 2, Elsevier, 2017, pp. 427–466.

Jones, Maria, Florence Kondylis, John Loeser, and Jeremy Magruder, "Factor market failures and the adoption of irrigation in rwanda," *American Economic Review*, 2022, *112* (7), 2316–52.

Kondylis, Florence, Valerie Mueller, and Jessica Zhu, "Seeing is believing? Evidence from an extension network experiment," *Journal of Development Economics*, 2017, *125*, 1–20.

Laajaj, Rachid, Karen Macours, Cargele Masso, Moses Thuita, and Bernard Vanlauwe, "Reconciling yield gains in agronomic trials with returns under African smallholder conditions," *Scientific reports*, 2020, *10* (1), 14286.

Lane, Gregory, "Adapting to Floods with Guaranteed Credit: Evidence from Bangladesh," 2022.

Magruder, Jeremy R, "An assessment of experimental evidence on agricultural technology adoption in developing countries," *Annual Review of Resource Economics*, 2018, *10*, 299–316.

Rosenzweig, Mark R and Christopher Udry, "External validity in a stochastic world: Evidence from low-income countries," *The Review of Economic Studies*, 2020, *87* (1), 343–381.

Suri, Tavneet, "Selection and comparative advantage in technology adoption," *Econometrica*, 2011, *79* (1), 159–209.

_ **and Christopher Udry**, "Agricultural technology in africa," *Journal of Economic Perspectives*, 2022, *36* (1), 33–56.

World Bank, "World Development Report 2008: Agriculture for Development," 2007.

Young, H Peyton, "Innovation diffusion in heterogeneous populations: Contagion, social influence, and social learning," *American economic review*, 2009, *99* (5), 1899–1924.

A Data appendix

In this section, we define key variables used in our analysis; as appropriate, we follow Jones et al. (2022) and many of our definitions are identical to theirs. All agricultural variables are constructed using dry season data for the household's two most important dry season plots with respect to cultivation of pulses.

Household variables: All household variables are constructed from the baseline.

- *Age of HHH*: Age of the household head.
- *Any pulses cultivation*: Any cultivation of pulses.
- *Veteran cultivator*: Any cultivation of promoted crops.
- *Any pulses improved seed*: Any use of improved seed for pulses; set to 0 when no pulses cultivation.
- *Any irrigation*: Any use of irrigation for paddy or pulses; set to 0 when no paddy or pulses cultivation.
- *Total plot area*: Area cultivated in hectares for paddy and pulses; set to 0 when no paddy or pulses cultivation.

Household-crop-survey wave variables:

- *Cultivation*: Indicator for any cultivation.
- *Improved seed*: Indicator for any use of improved seed; set to 0 when no cultivation.
- *Improved seed (Recycled)*: Indicator for any use of recycled improved seed; set to 0 when no improved seed use.
- *Yield ('000 BDT/ha)*: Harvested production, valued at median crop-survey wave sale prices, divided by area cultivated; winsorized at the 99th percentile and set to missing when no cultivation.
- *Fertilizer ('000 BDT/ha)*: Expenditures on fertilizer, divided by area cultivated; winsorized at the 99th percentile and set to missing when no cultivation. Only observed one year post-demo and two years post-demo.
- *Profit ('000 BDT/ha)*: Yield minus expenditures on hired labor per hectare minus household labor per hectare valued at 60% of the median agricultural

wage (0.6×236 BDT/day) minus fertilizer expenditures per hectare; winsorized at the 99th percentile and then set to 0 when no cultivation. The choice of 60% follows a recommendation from Agness et al. (2020).

  - *Profit (Demo year) ('000 BDT/ha)*: Demonstration year profits; set to missing when no demonstration or no reported cultivation. Expenditures on fertilizer are not subtracted as fertilizer was provided as part of the demonstration kit.

- *Area cultivated*: Area cultivated; set to missing when no cultivation.
- *Beliefs (High yielding)*: Indicator for reporting that one of the IAPP-promoted varieties is high yielding. Only observed one year post-demo.

Household-survey wave variables:

- *Paddy cultivation*: Any paddy cultivation.
- *Above median flood*: Indicator that leave-out flood exposure is above median. To construct leave-out flood exposure, we first construct a household's self-reported flood exposure as the fraction of crops cultivated by the household for which production was affected by flooding. We construct a household's associated leave-out flood exposure as the fraction of household-by-cultivated crop observations in that household's farmer group *excluding that household* for which production was affected by flooding.

Experimental variables:

- *Any demo in village*: Indicator that the household's farmer group was assigned to either *regular* demonstration, *shared* demonstration, or *decentralized* demonstration arms.
- *Shared demo in village*: Indicator that the household's farmer group was assigned to the *shared* demonstration arm.
- *Decentralized demo in village*: Indicator that the household's farmer group was assigned to the *decentralized* demonstration arm.
- *Demo*: Household-crop demonstration indicator. In analysis of impacts on placebo crops, this variable is the fraction of promoted crops that the household demonstrated.

Social network variables: All social network variables are constructed from the baseline. A household i and household j in the same farmer group are defined as connected if either household i with household j or household j with household i reported discussing farming techniques in the last 3 months.

- *# of connections*: Number of social network connections, defined at the household level.
- *# of connections to demo*: Number of social network connections to demonstrators defined at the household-crop level. In analysis of impacts on placebo crops, this variable is the average number of social connections to demonstrators (defined at the household-crop level) across promoted crops.

B Model appendix

Proof. **Proposition 1 and Corollary 1.** First, we apply the law of iterated expectations to population adoption $\mathbf{E}[D|A]$, conditioning D on expected returns M.

$$\begin{aligned}
\mathbf{E}[D|A] &= \mathbf{E}[\mathbf{E}[D|A,M]|A] && \text{LIE} && \\
&= \mathbf{E}[\mathbf{E}[D|M]|A] && D \perp A|M && \text{(Equation 8)} \\
&= \mathbf{E}[F_C(M)|A] && D = \mathbf{1}\{C < M\} && \text{(Equation 8)}
\end{aligned}$$

Intuitively, population adoption is the expected fraction of farmers with costs less than their expected returns.

Second, we express population adoption as a function of population average expected returns and the population variance of expected returns, which we denote μ_M and σ_M^2. Noting that expected returns $M|A$ are normally distributed (Equation 9), we parametrize $M|A \equiv \mu_M + \sigma_M Z|A$, where $\mu_M \equiv \ell A + (1-\ell)\mu_0$, $\sigma_M^2 \equiv \ell(1-\ell)\sigma_0^2$, and $Z \sim \mathcal{N}(0,1)$ is a standard normal random variable. We write

$$\begin{aligned}
\mathbf{E}[D|A] &= \mathbf{E}[F_C(M)|A] \\
&= \mathbf{E}[F_C(\mu_M + \sigma_M Z)|A]
\end{aligned}$$

Proposition 1 Third, we differentiate population adoption with respect to population average expected returns and the population variance of expected returns. We let f_C denote the density of costs; differentiating yields

$$\begin{aligned}\frac{d}{d\mu_M}\mathbf{E}[D|A] &= \frac{d}{d\mu_M}\mathbf{E}[F_C(\mu_M+\sigma_M Z)|A] \\ &= \mathbf{E}\left[\frac{d}{d\mu_M}F_C(\mu_M+\sigma_M Z)|A\right] \\ &= \mathbf{E}\left[f_C(\mu_M+\sigma_M Z)|A\right] \\ &= \mathbf{E}[f_C(M)|A]\end{aligned}$$

Population adoption is increasing in population average expected returns, as $\frac{d}{d\mu_M}\mathbf{E}[D|A] = \mathbf{E}[f_C(M)|A]$, and the density of costs $f_C(\cdot)$ is positive. The intuition underlying the sign of this derivative is clear—when farmers expect higher returns, they are more likely to adopt. That this derivative is the expectation of the density of costs evaluated at expected returns also has a clear intuition; it is the density of farmers for whom expected returns are equal to costs, that is the density of marginal adopters. An increase in population average expected returns therefore increases adoption by the density of marginal adopters.

$$\begin{aligned}\frac{d}{d\sigma_M}\mathbf{E}[D|A] &= \frac{d}{d\sigma_M}\mathbf{E}[F_C(\mu_M+\sigma_M Z)|A] \\ &= \mathbf{E}\left[\frac{d}{d\sigma_M}F_C(\mu_M+\sigma_M Z)|A\right] \\ &= \mathbf{E}\left[Zf_C(\mu_M+\sigma_M Z)|A\right] \\ &= \mathbf{E}\left[\frac{M-\mu_M}{\sigma_M}f_C(M)|A\right] \\ &= \frac{\mathbf{E}\left[\frac{M-\mu_M}{\sigma_M}f_C(M)|A\right]}{\mathbf{E}[f_C(M)|A]}\mathbf{E}[f_C(M)|A]\end{aligned}$$

The derivative of population adoption with respect to the population variance of expected returns is equal to the product of two terms: the average normalized distance between costs and population average expected returns among marginal adopters, and the density of marginal adopters. The intuition is as follows—increased belief dispersion "pushes" the distribution of beliefs away from population average expected returns, so adoption increases (decreases) when marginal farmers have relatively high (relatively low) expected returns compared to population average expected returns.

Population adoption is therefore increasing (decreasing) in the population variance of expected returns if marginal farmers have higher (lower) average expected returns than average farmers.

Corollary 1 Lastly, we differentiate population adoption with respect to ℓ, applying our expressions for population average expected returns μ_M and the population variance of expected returns σ_M^2.

$$\begin{aligned}\frac{d\mathbf{E}[D|A]}{d\ell} &= \frac{d\mu_M}{d\ell}\frac{d\mathbf{E}[D|A]}{d\mu_M} + \frac{d\sigma_M}{d\ell}\frac{d\mathbf{E}[D|A]}{d\sigma_M} \\ &= (A-\mu_0)\frac{d\mathbf{E}[D|A]}{d\mu_M} + \frac{\sigma_0^2}{\sigma_M}\left(\frac{1}{2}-\ell\right)\frac{d\mathbf{E}[D|A]}{d\sigma_M}\end{aligned}$$

As learning ℓ increases from 0 to 1, population average expected returns increase linearly from prior expected returns μ_0 to the true returns A—this force leads to adoption increasing (decreasing) with respect to learning when true returns A are above (below) prior expected returns μ_0.

As learning ℓ increases from 0 to 1, the population variance of expected returns increases and then decreases; when marginal farmers have relatively high (relatively low) expected returns compared to population average expected returns, this force causes adoption to increase and then decrease (decrease and then increase) with respect to learning.

Whether the effects of learning ℓ on population average expected returns or on the population variance of expected returns will dominate will depend on the underlying parameters. As a consequence, any of the following are possible: increasing adoption in learning ("S-curve"), decreasing adoption in learning, increasing then decreasing adoption in learning ("inverted U-curve"), or decreasing then increasing adoption in learning ("U-curve"). □

Overreaction (underreaction) to signals generates identical adoption to a less (more) precise prior We proceed in two steps. First, we show that the distribution of expected returns is identical in a model in which farmers form posterior expected returns as if the observed signal has variance n times lower than its true variance and in a model in which farmers' prior variance is n times higher. Second, we note that population adoption depends on beliefs only through the distribution of expected returns; therefore population adoption is identical under overreaction (underreaction) and less precise (more precise) priors.

Let $\widehat{\ell} \equiv \frac{\sigma_0^2}{\sigma_0^2+\sigma_Y^2/n} = \frac{n\sigma_0^2}{n\sigma_0^2+\sigma_Y^2}$ denote learning under the perceived signal variance, σ_Y^2/n. Note that learning under the perceived signal variance is equal to learning in a model without misperceptions, but where the prior variance equals $n\sigma_0^2$.

Let $\widehat{M} \equiv \widehat{\ell}Y + (1-\widehat{\ell})\mu_0$ denote expected returns under misperception of the signal variance. Note that when $n > 1$, $\widehat{\ell}$ places too much weight (relative to the variance minimizing weight) on the observed signal Y, and farmers' beliefs overreact to signals. Calculating the expectation and the variance of this expression, conditional on the technology, yields

$$\widehat{M}|A \sim \mathcal{N}(\widehat{\ell}A + (1-\widehat{\ell})\mu_0, \widehat{\ell}(1-\widehat{\ell})n\sigma_0^2)$$

Expected returns $\widehat{M}$ have the same distribution as in a model where the prior variance equals $n\sigma_0^2$ and there are no misperceptions.

C Learning the returns to a new technology explains medium run and long run impacts of demonstration

Our model of learning and technology adoption in Section 5, closely following Besley and Case (1994), assumed the following causal chain: demonstrations generate signals of the profitability of the new seeds, farmers update their beliefs about the profitability of the new seeds in response to observed signals, and decide to adopt the new seeds on the basis of these beliefs. We test this causal chain: does demonstration affect beliefs? are beliefs associated with adoption? do profits affect beliefs? and adoption?

Empirical strategy We implement these tests of the causal chain underlying our model using versions of the following empirical specification:

$$\begin{aligned} Y_{hcvt} = \alpha_{ct} &+ \beta_{1,t}\text{Any demo in village}_v + \beta_{2,t}\text{Shared demo in village}_v \\ &+ \beta_{3,t}\text{Decentralized demo in village}_v + \gamma_t W_{hcv} + \epsilon_{hcvt} \quad \text{(A1)} \end{aligned}$$

We report estimates of Equation A1 in Appendix Table A1, which we use to implement each of these four tests: demonstration affects beliefs (Column 1), beliefs are associated with adoption (Columns 2 and 3), profits affect beliefs (Column 4), and profits affect adoption (Columns 5 and 6).

Equation A1 extends Equation 4, used to estimate the impacts of demonstration and decentralizing demonstration on adoption, in three ways. First, it includes an additional covariate W_{hcv}: either household h's belief that crop c is high yielding, measured one year post-demonstration (Columns 2 and 3), or household h's demonstration year profits per hectare for crop c conditional on demonstrating (Columns 4, 5, and 6). Second, it varies the outcome Y_{hcvt}: either beliefs (Columns 1 and 4), or adoption of the new seeds one year post demonstration (Columns 2 and 5) or two years post demonstration (Columns 3 and 6). Third, it includes additional controls: crop fixed effects in Columns 1, 2, and 3, to control for across-crop variation in beliefs that may correlate with propensity to adopt the new seed, and crop-by-treatment assignment fixed effects (which absorb controls for treatment assignment) in Columns 4, 5, and 6, to control for across-crop variation in profits generated by treatment assignment through the scale of demonstration.

Results We find that the way farmers learn in our data corroborates predictions from our model of learning (Appendix Table A1).

First, village assignment to demonstration positively affects farmers' beliefs that the promoted seeds are high-yielding one year post-demonstration (Column 1). Consistent with our results in Table 6, our point estimates imply that decentralized demonstration increases farmers' beliefs, albeit not statistically significantly.

Second, farmers' beliefs one year post-demonstration are strongly and precisely associated with increased adoption of the new seeds one year post demonstration (Column 2). However, this association between beliefs one year post-demonstration and adoption fades somewhat two years post-demonstration, in line with the notion that information about the new seeds had diffused and farmers converged towards learning the true returns and an adoption steady state (Column 3).

Third, signals of the profitability of the new seeds (i.e., measured profits) generated by demonstration affect beliefs. Leveraging variation in profits conditional on crop and demonstration scale (i.e., crop-by-treatment assignment fixed effects), our point estimate implies a one standard deviation increase in demonstration year profits increases the probability the demonstrator believes the new seeds are high yielding by 7.4 percentage points, or a 26% increase (Column 4).

Lastly, signals of the profitability of the new seeds generated by demonstration affect adoption of the new seeds. A one standard deviation increase in demonstration year profits increases the probability the demonstration adopts the new seeds one year post-demonstration by 4.8 percentage points, 38% of the effect of demonstrating on adoption (Column 5). In Appendix D.1, we establish a close link between this estimate and the learning generated by demonstration, and interpret the magnitude of this effect through this lens: the larger the effect of demonstration profits on adoption, the more weight farmers put on the signal generated by demonstration, which implies more learning (as defined in our model) from demonstration. As with the association of adoption with beliefs, the effect of demonstration year profits on adoption fades two years post-demonstration (Column 6); farmers put less weight on information from their own demonstration as information about the new seeds diffuses.

D Structural model of learning and technology adoption

We proceed with estimating a structural model of experiential and social learning in three steps. In Appendix D.1, we impose additional restrictions on the model in Section 5 for identification and for tractability. We estimate the model in Appendix D.2, and present and discuss parameter estimates in Appendix D.3. To apply our estimated structural model to counterfactual analysis, we derive expressions for the gains from demonstration under our assumed restrictions in Appendix D.4. Lastly, we conduct two sets of analysis of robustness. We analyze sensitivity and informativeness for our estimation, and discuss consistency with our identification argument, in Appendix D.5. In Appendix D.6, we estimate the model under alternative assumptions for identification.

D.1 Additional assumptions

The model in Section 5 contains the following key primitives which we must estimate or calibrate: prior beliefs (μ_0, σ_0), the returns to the new seeds a, the distribution

Table A1: Observed profits during demonstration shape beliefs, and in turn beliefs determine medium run, but not long run, adoption decisions

	Promoted crops					
	One year post-demo		Two years post-demo	One year post-demo		Two years post-demo
	Beliefs (High yielding)	Improved seed		Beliefs (High yielding)	Improved seed	
	(1)	(2)	(3)	(4)	(5)	(6)
Any demo in village	0.084	0.027	0.017			
	(0.027)	(0.020)	(0.014)			
	[0.002]	[0.177]	[0.228]			
Shared demo in village	-0.005	-0.002	0.019			
	(0.029)	(0.022)	(0.019)			
	[0.861]	[0.942]	[0.317]			
Decentralized demo in village	0.044	0.055	-0.003			
	(0.042)	(0.027)	(0.015)			
	[0.297]	[0.042]	[0.850]			
Beliefs (High yielding)		0.067	0.024			
		(0.021)	(0.013)			
		[0.001]	[0.061]			
Profit (Demo year) ('000 BDT/ha)				0.0026	0.0017	-6e-04
				(0.0012)	(8e-04)	(8e-04)
				[0.023]	[0.038]	[0.461]
Demo only				X	X	X
Crop FE	X	X	X			
Crop-Treatment assignment FE				X	X	X
# of observations (Farmer-crop-wave)	2,934	2,934	2,037	170	184	123
# of clusters (Farmer group)	91	91	67	55	59	42

of costs F_C, and the variance of the observed signal σ_Y^2. We impose six additional assumptions below in order to estimate or calibrate these primitives. The first four assumptions relate to the identification of key parameters necessary for the identification of learning: new seed adoption at prior beliefs $F_C(\mu_0)$, new seed adoption after learning the true returns $F_C(a)$, the true returns to the new seeds relative to prior expectations $a - \mu_0$, and the variance of the observed signal σ_Y^2, respectively. The last two assumptions are functional form assumptions, necessary for tractability, on the precision of the signal σ_Y^2 and on the distribution of costs F_C, respectively.

Assumption 1. *Control group adoption of new seeds is determined by prior beliefs.*

First, we assume that adoption in the control group is determined by prior beliefs, and is therefore equal to the fraction of farmers whose costs of adoption are less than prior expected returns to the new seeds, $F_C(\mu_0)$. This is consistent with our finding in Figure 4a that adoption of the new seeds in the control group is stable over time, suggesting control group beliefs remain relatively static. This assumption is further motivated by the context—as farmer group meetings in which the new seeds were introduced happened in both control and treatment villages, control village farmers should have similar awareness of the technology and beliefs as treatment village farmers who did not demonstrate and were not connected to any demonstrators, consistent with our finding in Table 8 that demonstration and number of connections to demonstrators fully explain the impacts of treatment one year post-demonstration.

Assumption 2. *Demonstration villages learn the true returns to the new seeds two years post-demonstration.*

Second, we assume that farmers in treatment villages learn the true returns to the new seeds two years post-demonstration, and adoption in treatment villages two years post-demonstration is therefore equal to the fraction of farmers whose costs of adoption are less than the true returns to the new seeds, $F_C(a)$. This is consistent with our finding of convergence in adoption of the new seeds two years post-demonstration across treatment arms in Table 6.[A1] It is also consistent with our finding that signals from demonstration have no impact on adoption two years post-demonstration in Table 8 and Appendix Table A1, as should be the case if farmers have learned the true returns. This full learning is consistent with farmers sharing information outside their social networks two years, but not one year, post-demonstration. The active engagement of agricultural extension workers in the promotion of these new seeds and the organization of farmer groups in our experimental sample support this notion.

Assumption 3. *Impacts of the new seeds on profits inform true returns relative to prior expected returns.*

Third, we use estimates of the impact of the technology on profits based on Table 5 to calibrate the true returns to the technology relative to prior expectations $a - \mu_0$.[A2] To do so, we assume that farmers' prior expectation is that the new seeds have no impact on profits; in this case, the impact of the new seeds on profits is equal to $a - \mu_0$. We note that farmers may instead expect positive profits, in which case the impact of the new seeds on profits may be larger than $a - \mu_0$. Alternatively, farmers

[A1]In Appendix Table A32, we implement a more stringent test of conditional convergence; convergence two years post-demonstration continues to hold conditional on veteran cultivator status.

[A2]Specifically, we assume the new seeds increase production by reducing downside risk, consistent with our estimates in Table 5. Therefore, we assume that impacts of the new seeds on profits are equal to one half of the difference in impacts of above median flood on profits between improved and non-improved seeds.

may place a value on the technology greater than expected profits, given that the technology effectively eliminates downside risk from saline-flood shocks, leading the impact of the new seeds on profits to be smaller than $a - \mu_0$. In Appendix D.6, we therefore also estimate the model, and report counterfactuals, under scenarios where the impact of the technology on profits is half as large, equal to, or twice as large as $a - \mu_0$.

Assumption 4. *Impacts of demonstration year profits on adoption of the new seeds correspond to impacts of the signal from own demonstration on adoption.*

Fourth, we use our estimates of the impact of demonstration year profits on adoption of the new seeds one year post-demonstration in Appendix Table A1 as a proxy for the impacts of the signal from own demonstration on adoption of the new seeds. We allow that the signal observed by the demonstrator is not observed by us, as it depends on factors such as idiosyncratic growing conditions. We instead assume that a 1 BDT/ha increase in demonstration year profits, conditional on crop and treatment assignment, is associated with a 1 BDT/ha increase in the signal observed by a demonstrator; we analyze robustness to deviations from this assumption in Appendix D.6. This estimated impact allows identification of unobserved learning; a heuristic argument is as follows. First, by definition, the weight farmers put on their observed signal (the aggregation of own and connections' demonstrations) is equal to learning. Second, by assumption, impacts on adoption two years post-demonstration are equal to the impacts of observing the true returns to the new seeds, which is equal to the impact of observing a signal, on which full weight is placed, equal to the true returns. Therefore, the impacts on adoption of an increase in the signal observed by a demonstrator, from prior expectations to the true returns to the new seeds, divided by the impacts on adoption of observing the true returns to the new seeds, are approximately equal to learning.

Assumption 5. *Parametrization of learning from self and from others.*

Fifth, we impose additional assumptions on the signals of the returns to the new seeds observed by farmers. These assumptions are inspired by our findings in Section 6—farmers do not learn from scale or from geographic heterogeneity, farmers learn more from their own demonstration than from a connection's demonstration, and own demonstration and number of connections to demonstrators fully explain effects of demonstration on adoption one year post-demonstration. To tractably match these findings, we assume that farmers observe independent and normally distributed signals of the returns to the technology from their own and connections' demonstrations. First, signals from own demonstration are distributed $\mathcal{N}(a, \sigma_0^2/\eta)$, where η parametrizes the relative precision of the signal compared to farmers' prior beliefs. And second, signals from connections' demonstrations are distributed $\mathcal{N}(a, \sigma_0^2/(\omega\eta))$, where ω parametrizes the relative precision of learning from others compared to learning from self; our theory in Section 5.3 suggests ω is approximately the ratio of the effect of number of connections to demonstrators on adoption to the effect of own demonstration on adoption. Farmers efficiently aggregate signals, implying the variance of the signal for household h for crop c, and associated learning, one year post-demonstration is equal to

$$\sigma^2_{Y,hcv1} = \frac{\sigma_0^2}{\eta \text{Demo}_{hc} + \omega\eta \# \text{ of connections to demo}_{hc}} \tag{A2}$$

$$\ell_{hcv1} = \frac{\eta \text{Demo}_{hc} + \omega\eta \# \text{ of connections to demo}_{hc}}{1 + \eta \text{Demo}_{hc} + \omega\eta \# \text{ of connections to demo}_{hc}} \tag{A3}$$

Assumption 6. *Parametrization of distribution of costs of new seed adoption.*

Sixth, we assume that the costs of adopting the new seeds are normally distributed. As average costs is not separately identified from farmers' prior expected returns and the true returns to the new seeds, we normalize costs to be mean 0 and therefore assume $C \sim \mathcal{N}(0, \sigma_C^2)$. Letting Φ be the standard normal CDF, this implies that $F_C(y) = \Phi(y/\sigma_C)$.

D.2 Estimation

Structural parameters The assumptions imposed in Section D.1 on the model in Section 5 result in a model with 6 parameters. Farmers have prior expectation of the returns to the technology μ_0 with prior variance σ_0^2, observe signals of the true returns to the technology a which have relative precision η (when observed from own demonstration) and $\omega\eta$ (when observed from connections' demonstration), and make adoption decisions based on normally distributed costs with mean 0 and variance σ_C^2.

Moments and informal analysis of identification We estimate 6 parameters $(\mu_0/\sigma_C, \sigma_0/\sigma_C, a/\sigma_C, \eta, \omega, \sigma_C)$ by simulated method of moments, matching 6 moments suggested by our assumptions in Section D.1. First, we match control group adoption from Table 3, which by Assumptions 1 and 6 is equal to $\Phi(\mu_0/\sigma_C)$. Second, we match treatment group adoption two years post demonstration from Table 6, which by Assumptions 2 and 6 is equal to $\Phi(a/\sigma_C)$. Third, we match the differential impact of improved seed on profits when saline flooding is above median from Table 5, which by Assumption 3 is equal to $a - \mu_0$ divided by the probability of an above median saline flood (0.5). These first three moments identify (μ_0, a, σ_C). Fourth, we match the impact of demonstrator profits on their adoption one year post demonstration from Appendix Table A1, which by Assumption 4 informs the impact of a more positive signal on adoption, and as discussed enables identification of learning. We simulate this moment with random draws of signals, accounting for the fact that while demonstrator profits corresponds only the demonstrator's signal, demonstrators also observe signals from their connections. Fifth and sixth, we match the impact of demonstration on adoption and the impact of connections to demonstrators on adoption in Table 8. Along with Assumption 5, these latter three moments jointly identify (σ_0, η, ω).[A3]

Implementation of generalized method of moments As the structural model is just identified, we perfectly match the 6 moments above. To conduct inference, we estimate the cross-equation variance of our moments clustering standard errors at the village level.

D.3 Parameter estimates

We report the estimates of the parameters of the structural model in Appendix Table A2a (and matched moments in Appendix Table A2b), and we briefly discuss the interpretation of the parameter estimates.

Learning from self and from others (η, ω) First, we estimate a strength of learning from self (η) of 0.44; this value implies that a demonstrator with no connections reaches learning of 0.31. Second, we estimate a relative learning from others (ω) of 0.10; consistent with our theory in Section 5.3, this is closely approximated by

[A3]We present discussion of the sensitivity (Andrews et al., 2017) and informativeness (Andrews et al., 2020) of our parameter estimates and matched moments in Appendix D.5, and we find they correspond closely with this informal analysis of identification.

Table A2: Structural estimates and moments

(a) Parameter estimates

Parameter	Estimate
η: learning from self	0.441
	(0.486)
ω: relative learning from others	0.095
	(0.074)
a/σ_C: returns to new seeds	-1.660
	(0.065)
μ_0/σ_C: prior expected returns	-1.863
	(0.144)
σ_0/σ_C: prior standard deviation	5.046
	(1.613)
σ_C: costs standard deviation ('000 BDT/ha)	16.950
	(16.695)

(b) Moments

Moment	Estimate	Model fit
Demo	0.127	0.127
Table 8, Column 1, Row 5	(0.032)	
# of connections to demo	0.017	0.017
Table 8, Column 1, Row 6	(0.007)	
Any demonstration adoption two years post-demonstration	0.048	0.048
Table 6, Column 3, Row 1 + Control mean	(0.007)	
Control adoption	0.031	0.031
Table 3, Column 1, Control mean	(0.010)	
Profit ('000 BDT/ha) impact on adoption	0.002	0.002
Appendix Table A1, Column 5, Row 5	(0.001)	
Technology impact on above median flood profit ('000 BDT/ha)	6.874	6.874
Table 5, Column 4, Row 3	(3.937)	

the ratio of the impact of number of connections to demonstrators to the impact of demonstrating on adoption of the new seeds (0.13). The value implies that a demonstrator with five connections to demonstrators reaches learning of 0.39. In Appendix D.1, in order to estimate the structural model, we assumed that farmers fully learned the returns to the new seeds two years post-demonstration in demonstration villages; our estimate of η implies a farmer demonstrating six times reaches learning of 0.73, suggesting the additional year of communication enabling learning from others to become as strong as learning from self, facilitated by farmer group meetings, could plausibly enable close to full learning.

True returns and prior expectations $(a/\sigma_C, \mu_0/\sigma_C, \sigma_0/\sigma_C)$ Third, we estimate the returns to the technology a/σ_C, in units of standard deviations of costs, of -1.66. As costs are assumed to be normally distributed, the estimate can be interpreted as a z-score associated with the probability of adoption of the new seeds under full learning (assumed to be adoption two years post-demonstration in demonstration villages). Fourth, and similarly, we estimate farmers' prior expected returns μ_0/σ_C, again in units of standard deviations of costs, of -1.86. The true returns to the new seeds are higher than farmers' prior expectations, which we recover from the positive effect of demonstration on adoption two years post-demonstration. Fifth, we estimate farmers' prior standard deviation of 5.05. Farmers expect that heterogeneity in the returns to newly introduced technologies is substantially larger than heterogeneity in their idiosyncratic costs of adoption: many new technologies are adopted by almost no farmers, while many new technologies are adopted by almost all farmers.

Standard deviation of costs σ_C Lastly, we estimate a standard deviation of costs of 17,000 BDT/ha. These costs can be broadly interpreted as unobserved heterogeneity that drives farmers to jointly adopt improved seed and cultivate associated pulses. This standard deviation of costs is equal to roughly one third of the difference in yields between pulses and paddy in Appendix Table A5; this relative magnitude is sufficiently large that it is plausible that some farmers adopt the new seeds while others do not. This standard deviation is roughly one half of the standard deviation of observed profits among pulses cultivators; this relative magnitude is sufficiently small to reflect plausible heterogeneity in the returns to adopting improved seed across farmers.

D.4 Gains from demonstration

In Section 7.1, we defined the ex-post and ex-ante gains from demonstration as functions of ex-post and ex-ante profits, respectively. To calculate the ex-post and ex-ante gains from demonstration in our estimated structural model of learning and technology adoption, below we derive expressions for these gains as functions of estimated parameters.

We first note the following fact about jointly normally distributed random variables. Suppose X and Y are jointly normally distributed, with means μ_X and μ_Y, variances σ_X^2 and σ_Y^2, and covariance σ_{XY}.

$$\mathbf{E}[\mathbf{1}\{X > 0\}Y] = \Phi\left(\frac{\mu_X}{\sigma_X}\right)\mu_Y + \phi\left(\frac{\mu_X}{\sigma_X}\right)\frac{\sigma_{XY}}{\sigma_X} \tag{A4}$$

Using the fact that adoption $D = \mathbf{1}\{M - C > 0\}$, both ex-post profits (Equation 11) and ex-ante profits (Equation 12) take this form under our assumptions in Section D.1.

Second, we calculate

$$\left(\begin{array}{c} M - C \\ A - C \end{array}\right)\Bigg| A \sim \mathcal{N}\left(\left(\begin{array}{c} \mu_0 + \ell(A - \mu_0) \\ A \end{array}\right), \left(\begin{array}{cc} \ell(1-\ell)\sigma_0^2 + \sigma_C^2 & \sigma_C^2 \\ \sigma_C^2 & \sigma_C^2 \end{array}\right)\right) \tag{A5}$$

$$\left(\begin{array}{c} M - C \\ A - C \end{array}\right) \sim \mathcal{N}\left(\left(\begin{array}{c} \mu_0 \\ \mu_0 \end{array}\right), \left(\begin{array}{cc} \ell\sigma_0^2 + \sigma_C^2 & \ell\sigma_0^2 + \sigma_C^2 \\ \ell\sigma_0^2 + \sigma_C^2 & \sigma_0^2 + \sigma_C^2 \end{array}\right)\right) \tag{A6}$$

Lastly, we apply Equation A4 to Equations A5 and A6 to calculate ex-post profits (Equation 11) and ex-ante profits (Equation 12).

$$\begin{aligned}\pi_{hcvt}^{\text{ex-post}} = \Phi&\left(\frac{\mu_0 + \ell_{hcvt}(a - \mu_0)}{\sqrt{\ell_{hcvt}(1 - \ell_{hcvt})\sigma_0^2 + \sigma_C^2}}\right) a \\ &+ \phi\left(\frac{\mu_0 + \ell_{hcvt}(a - \mu_0)}{\sqrt{\ell_{hcvt}(1 - \ell_{hcvt})\sigma_0^2 + \sigma_C^2}}\right)\frac{\sigma_C^2}{\sqrt{\ell_{hcvt}(1 - \ell_{hcvt})\sigma_0^2 + \sigma_C^2}}\end{aligned} \tag{A7}$$

$$\pi_{hcvt}^{\text{ex-ante}} = \Phi\left(\frac{\mu_0}{\sqrt{\ell_{hcvt}\sigma_0^2 + \sigma_C^2}}\right)\mu_0 + \phi\left(\frac{\mu_0}{\sqrt{\ell_{hcvt}\sigma_0^2 + \sigma_C^2}}\right)\sqrt{\ell_{hcvt}\sigma_0^2 + \sigma_C^2} \tag{A8}$$

In Appendix Table A13, we apply Equations A7 and A8 to calculate ex-post profits and ex-ante profits across treatment arms, and to simulate them under counterfactuals.

D.5 Informativeness and sensitivity

In Appendix Table A3, we report sensitivity of (informativeness for) each of the estimated parameters reported in Appendix Table A2a to (of) the moments used to estimate them in Appendix Table A2b. We note the following, consistent with our informal analysis of identification in Appendix D.2:

- Estimated learning from self η, which corresponds closely to learning after demonstrating $\eta/(1+\eta)$, is larger when the impact of learning the true returns to the new seeds on adoption is smaller, increases in demonstration profits have larger impacts on adoption of the new seeds, and when the new seeds have larger impacts on profits. This is consistent with our approximation of learning: learning is approximately the impact of an increase in demonstration profits equal to the impact of the new seeds on profits, divided by the impact of learning the true returns to the new seeds on adoption.

- Estimated relative learning from others ω is larger when the impacts of connections to demonstrations on adoption are larger and when the impact of demonstrating on adoption are smaller, as ω is approximately equal to the ratio of these impacts.

- The estimated returns to the new seeds a/σ_C are, by assumption, a transformation of demonstration village adoption of the new seeds two years post-demonstration.

- Estimated prior expectations μ_0/σ_C are, by assumption, a transformation of control village adoption of the new seeds.

- Estimated prior standard deviation σ_0/σ_C is large when the impacts of belief dispersion on adoption are large, as is true when the impacts of demonstration on adoption are larger than would be expected given learning and the impacts of learning the true returns to the new seeds on adoption. As a consequence, the estimated prior standard deviation is increasing in the impacts of demonstrating, or of having additional connections to demonstrations, on adoption of the new seeds.

- Estimated costs standard deviation σ_C is increasing in the impacts of the new seeds on profits, and decreasing in the impacts of learning the true returns to the new seeds on adoption. This is consistent with a first order approximation of learning the true returns to the new seeds on adoption given our assumption of normally distributed costs of adoption.

Table A3: Informativeness and sensitivity

	Sensitivity {Informativeness}					
	η	ω	a/σ_C	μ_0/σ_C	σ_0/σ_C	σ_C
Demo	1.250	-1.594	0.000	0.000	23.128	0.000
Table 8, Column 1, Row 5	{0.000}	{0.338}	{0.005}	{0.000}	{0.317}	{0.023}
# of connections to demo	14.009	4.881	0.000	0.000	184.815	0.000
Table 8, Column 1, Row 6	{0.050}	{0.142}	{0.023}	{0.000}	{0.667}	{0.001}
Any demonstration adoption two years post-demonstration	-23.472	2.535	9.947	0.000	12.776	-831.553
Table 6, Column 3, Row 1 + Control mean	{0.039}	{0.037}	{1.000}	{0.000}	{0.030}	{0.069}
Control adoption	31.546	-3.698	0.000	14.218	-32.806	1188.586
Table 3, Column 1, Control mean	{0.498}	{0.286}	{0.000}	{1.000}	{0.050}	{0.599}
Profit ('000 BDT/ha) impact on adoption	273.978	-26.413	0.000	0.000	-313.910	0.000
Appendix Table A1, Column 5, Row 5	{0.139}	{0.061}	{0.003}	{0.000}	{0.000}	{0.009}
Technology impact on above median flood profit ('000 BDT/ha)	0.066	-0.006	0.000	0.000	-0.076	2.466
Table 5, Column 4, Row 3	{0.246}	{0.024}	{0.011}	{0.008}	{0.062}	{0.372}

D.6 Structural model robustness

To complement our analysis of sensitivity and informativeness in Appendix D.5, we replicate Section 7, estimating the structural model and implied gains from demonstration, under alternative assumptions made for identification in Section D.1.

First, we consider robustness to Assumption 3, that the impact of the new seeds on profits is equal to the difference between the true returns to the new seeds and farmers' prior expected returns. We consider the cases where our estimate of the impact of the new seeds on profits is biased by a factor of 0.5 (e.g., farmers learn that the value the new seed more than its impact on profits as it reduces downside risk) or by a factor of 2.0 (e.g., farmers held prior expectations that the new seed is profitable) for this difference.

Second, we consider robustness to Assumption 4, that the impact of demonstration year profits on adoption is equal to the impact of the signal from own demonstration on adoption. We consider cases where our estimate of the impact of demonstration year profits on adoption is biased by a factor of 0.5 (e.g., demonstration year profits are measured with error) or by a factor of 2.0 (e.g., demonstration year profits are correlated with other determinants of adoption of the new seeds) for the impact of the signal.

In Appendix Table A4, we report parameter estimates and gains from demonstration under all permutations of these alternative assumed bias factors (with 1.0 corresponding to no bias) for Assumptions 3 and 4.

Table A4: Structural estimates and gains from demonstration under bias in estimates of impacts of floods and demonstration year profits

	Base	Robustness							
Assumed bias									
Technology impact on profit	1	0.5	0.5	0.5	1.0	1.0	2.0	2.0	2.0
Demonstration profit impact on adoption	1	0.5	1.0	2.0	0.5	2.0	0.5	1.0	2.0
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Parameters									
η: learning from self	0.441	4.308	0.984	0.441	0.984	0.234	0.441	0.234	0.127
	(0.486)	(31.840)	(1.361)	(0.486)	(1.361)	(0.224)	(0.486)	(0.224)	(0.106)
ω: relative learning from others	0.095	0.007	0.055	0.095	0.055	0.121	0.095	0.121	0.124
	(0.074)	(0.093)	(0.079)	(0.074)	(0.079)	(0.075)	(0.074)	(0.075)	(0.076)
a/σ_C: returns to new seeds	-1.660	-1.660	-1.660	-1.660	-1.660	-1.660	-1.660	-1.660	-1.660
	(0.065)	(0.065)	(0.065)	(0.065)	(0.065)	(0.065)	(0.065)	(0.065)	(0.065)
μ_0/σ_C: prior expected returns	-1.863	-1.863	-1.863	-1.863	-1.863	-1.863	-1.863	-1.863	-1.863
	(0.144)	(0.144)	(0.144)	(0.144)	(0.144)	(0.144)	(0.144)	(0.144)	(0.144)
σ_0/σ_C: prior standard deviation	5.046	5.805	4.908	5.046	4.908	5.992	5.046	5.992	6.567
	(1.613)	(10.528)	(1.721)	(1.613)	(1.721)	(2.035)	(1.613)	(2.035)	(2.384)
σ_C: costs standard deviation ('000 BDT/ha)	16.950	33.901	33.901	33.901	16.950	16.950	8.475	8.475	8.475
	(16.695)	(33.390)	(33.390)	(33.390)	(16.695)	(16.695)	(8.347)	(8.347)	(8.347)
Gains from demonstration ('000 BDT/ha)									
One year post-demonstration, ex-ante, Any demo	2.28	8.44	5.90	4.56	2.95	2.18	1.14	1.09	0.78
One year post-demonstration, ex-ante, Decentralized demo	3.14	11.30	7.85	6.28	3.92	3.09	1.57	1.55	1.15
One year post-demonstration, ex-post, Any demo	-1.67	-3.37	-3.72	-3.35	-1.86	-1.71	-0.84	-0.85	-0.67
One year post-demonstration, ex-post, Decentralized demo	-2.06	-4.14	-4.33	-4.12	-2.17	-2.18	-1.03	-1.09	-0.92
Two years post-demonstration, ex-ante, Any demo	21.04	51.62	40.38	42.09	20.19	27.00	10.52	13.50	15.34
One year post-demonstration, ex-ante, no social learning, Any demo	1.14	6.40	3.33	2.27	1.66	1.00	0.57	0.50	0.36

Estimates that are robust to bias in Assumptions 3 and 4 While some of our estimates are affected by potential bias in Assumptions 3 and 4, many of our key estimates, which are linked to most of the qualitative conclusions drawn from our structural model in Section 7, are not affected by this potential bias.

- The impact of the new seeds on profits and the impact of demonstration year profits on adoption are the only two matched moments with financial units: '000 BDT/ha, and ('000 BDT/ha)$^{-1}$, respectively. Upward bias in the impact of the new seeds on profits, and inversely proportional downward bias in the impact of demonstration year profits on adoption, is therefore equivalent to changing financial units: the standard deviation of costs σ_C is biased proportionally upward, as are the gains from demonstration, while all other parameters are unaffected.

- Relative learning from others ω, and the fraction of learning that is social, are approximately proportional to the relative impact of connections to demonstrators and demonstrating on adoption of the new seeds, and are therefore only weakly affected by Assumptions 3 and 4.

- Gains from decentralizing demonstration relative to gains from regular demonstration, are approximately proportional to their relative impacts on learning, which are approximately equal to their relative impacts on adoption of the new seeds, and are therefore only weakly affected by Assumptions 3 and 4. Similarly, prior expected returns and the true returns to the new seeds, μ_0/σ_C and a/σ_C, are estimated irrespective of these assumptions.

- Ex-post gains from demonstration relative to ex-ante gains from demonstration are very closely linked to adoption of the new seeds under no learning, under demonstration, and under full learning, as these levels of adoption are strongly informative of "noisy adoption". These relative gains, and relatedly prior uncertainty σ_0/σ_C, are therefore only weakly affected by Assumptions 3 and 4.

The case of upward bias in impact of new seeds on profits (Assumption 3) When the impact of the new seeds on profits is biased upward for the difference between the true returns to the new seeds and farmers' prior expectations, the following occur:

- Consistent with our discussion in Section D, the standard deviation of costs σ_C is biased proportionately upward and learning from self η is biased proportionately upward.

- Consequentially, gains from demonstration are biased proportionately upward (in magnitude).

The case of upward bias in impact of demonstration year profits on adoption (Assumption 4) When the impact of demonstration year profits on adoption is biased upward for the impact of the signal from own demonstration on adoption, the following occur:

- Consistent with our discussion in Section D, learning from self η is biased proportionately upward.
- Consequentially, gains from demonstration are biased somewhat upward (in magnitude).

Non-targeted moments and plausible biases Certain forms of bias are inconsistent with moments not targeted by the structural estimation. In general, biases that result in larger gains from demonstration result in ex-post gains from demonstration that are inconsistent with the observed gains from demonstration (i.e., impacts on profits). Other biases, such as upward bias in either the impact of the new seeds on profits or of demonstration year profits on adoption, are more difficult to rule out—under these other biases evaluated in Appendix Table A4, it is possible that the learning generated by self-demonstration, the standard deviation of costs, and the gains from demonstration are proportionally smaller than those we estimate.

E Appendix Tables

Table A5: Agricultural descriptive statistics

	"Pulses"					
	Lentil	Mung	Mustard	Sesame	Wheat	Boro
	(1)	(2)	(3)	(4)	(5)	(6)
Yield ('000 BDT/ha)	36.89	39.36	39.74	26.91	46.34	88.09
Hired labor expenditures ('000 BDT/ha)	2.50	2.80	3.80	3.54	5.42	12.17
Household labor ('000 BDT/ha)	17.72	19.47	19.67	19.01	19.92	13.40
Fertilizer expenditures ('000 BDT/ha)	5.99	4.90	8.36	3.53	11.24	22.18
Seed expenditures ('000 BDT/ha)	3.97	2.62	1.57	1.71	2.90	3.66
Irrigation expenditures ('000 BDT/ha)	0.24	0.25	1.80	0.19	6.23	9.55
Profit ('000 BDT/ha)	6.48	9.32	4.54	-1.07	0.63	27.13
Any improved seed	0.106	0.399	0.306	0.218	0.601	0.599
Share irrigated	0.049	0.064	0.340	0.077	0.803	0.895
Area cultivated (ha)	0.118	0.161	0.109	0.135	0.101	0.297
Cultivated	0.165	0.236	0.046	0.025	0.065	0.388
# of observations (Farmer-crop-wave)	511	734	144	78	203	1,205

Notes: Sample averages of outcomes by crop across farmer-crop-survey wave observations are presented in this table. We restrict to observations one year and two years post-demo, in which input expenditures and household labor were collected disaggregated by crop. Profits in this table are calculated as yields net of expenditures on hired labor, fertilizer, seed, and irrigation, and household labor valued at 60% of the median wage of 236 BDT/day (choice of 60% following Agness et al., 2020). For reference, the median wage in our data is 236 BDT/person-day and the exchange rate was approximately 75 BDT = 1 USD over our study period.

Table A6: Sample descriptive statistics

	Sample								
							Pre-demo year sample		
	Pre-demo year	Demo year	One year post-demo	Two years post-demo	(3) + (4)	(2) + (3) + (4)	One year post-demo	Two years post-demo	(8) + (9)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Full sample	7,035	4,750	9,380	6,150	15,530	20,280	6,105	6,015	12,120
	<1,407>	<950>	<1,876>	<1,230>	<3,106>	<4,056>	<1,221>	<1,203>	<2,424>
	{78}	{109}	{109}	{77}	{109}	{109}	{77}	{77}	{77}
Household analysis sample	6,895	4,600	9,075	6,020	15,095	19,695	5,985	5,895	11,880
(Drop 2 Unions	<1,379>	<920>	<1,815>	<1,204>	<3,019>	<3,939>	<1,197>	<1,179>	<2,376>
without demonstration)	{76}	{105}	{105}	{75}	{105}	{105}	{75}	{75}	{75}
	2								
Promoted crops only	2,552	1,642	3,289	2,244	5,533	7,175	2,223	2,199	4,422
	<1,308>	<842>	<1,676>	<1,142>	<2,818>	<3,660>	<1,133>	<1,120>	<2,253>
	{71}	{96}	{96}	{70}	{96}	{96}	{70}	{70}	{70}
			6(1,2), A1(1,2)	6(3,4), A1(3)	3(1,2,5,6), 6(5,6), 5b(1)		8(1,2,3)	8(5)	4(1,2,5,6), 5b(2,3)
Demo	280	215	355	262	617	832	259	254	513
	<279>	<215>	<353>	<261>	<614>	<829>	<258>	<253>	<511>
	{53}	{62}	{72}	{53}	{72}	{72}	{53}	{53}	{53}
			A1(4,5)	A1(6)					
Cultivated	164	424	610	302	912	1,336	433	296	729
	<159>	<371>	<556>	<282>	<838>	<1,209>	<398>	<277>	<675>
	{39}	{84}	{76}	{55}	{81}	{90}	{58}	{55}	{63}
					3(3,4)	5			4(3,4)
Placebo crops only	4,343	2,958	5,786	3,776	9,562	12,520	3,762	3,696	7,458
	<1,379>	<920>	<1,815>	<1,204>	<3,019>	<3,939>	<1,197>	<1,179>	<2,376>
	{76}	{105}	{105}	{75}	{105}	{105}	{75}	{75}	{75}
							8(4)		

Notes: This table presents number of farmer-crop-wave observations, number of farmer-wave observations in angle brackets, and number of farmer groups in curly brackets. Tables that use each sample for analysis are presented below these reported numbers of observations, with associated columns when not all columns use that sample.

Table A7: Attrition across control and treatment

	Mean (SD)		Coef. (SE) [p]
	Control	Any demo in village	Difference
	(1)	(2)	(2) - (1)
Surveyed (Demo year)	0.324	0.374	0.050
	(0.469)	(0.484)	(0.030)
			[0.098]
Surveyed (One year post-demo)	0.840	0.875	0.035
	(0.367)	(0.331)	(0.053)
			[0.503]
Surveyed (Two years post-demo)	0.811	0.866	0.055
	(0.392)	(0.341)	(0.052)
			[0.292]
# of observations (Farmer)	281	1098	
# of clusters (Farmer group)	15	61	
Omnibus F-stat [p]			1.22
			[0.307]

Table A8: Attrition across treatment arms

	Mean (SD)			Coef. (SE) [p]	
	Regular demo in village	Shared demo in village	Decentral-ized demo in village	Difference	
	(1)	(2)	(3)	(2) - (1)	(3) - (1)
Surveyed (Demo year)	0.343	0.372	0.405	0.029	0.062
	(0.475)	(0.484)	(0.492)	(0.025)	(0.031)
				[0.262]	[0.042]
Surveyed (One year post-demo)	0.865	0.887	0.872	0.022	0.007
	(0.342)	(0.316)	(0.335)	(0.036)	(0.036)
				[0.540]	[0.848]
Surveyed (Two years post-demo)	0.850	0.885	0.861	0.034	0.011
	(0.357)	(0.320)	(0.346)	(0.038)	(0.035)
				[0.369]	[0.757]
# of observations (Farmer)	341	382	375		
# of clusters (Farmer group)	19	21	21		
Omnibus F-stat [p]				1.40	1.95
				[0.258]	[0.137]

Table A9: Promoted crop balance across control and treatment

	Mean		Coef. (SE) [p]
	Control	Any demo in village	Difference
	(1)	(2)	(2) - (1)
Wheat promoted	0.542	0.537	-0.005 (0.117) [0.965]
Lentil promoted	0.375	0.329	-0.046 (0.113) [0.686]
Mung promoted	0.375	0.390	0.015 (0.114) [0.894]
Sesame promoted	0.167	0.146	-0.020 (0.086) [0.814]
Mustard promoted	0.417	0.366	-0.051 (0.115) [0.659]
# of observations (Farmer group)	24	82	
# of clusters (Farmer group)	24	82	
Omnibus F-stat [p]			0.06 [0.998]

Table A10: Promoted crop balance across treatment arms

	Mean			Coef. (SE) [p]	
	Regular demo in village	Shared demo in village	Decentralized demo in village	Difference	
	(1)	(2)	(3)	(2) - (1)	(3) - (1)
Wheat promoted	0.593	0.571	0.444	-0.021	-0.148
				(0.135)	(0.137)
				[0.876]	[0.283]
Lentil promoted	0.333	0.250	0.407	-0.083	0.074
				(0.124)	(0.134)
				[0.505]	[0.581]
Mung promoted	0.333	0.393	0.444	0.060	0.111
				(0.132)	(0.134)
				[0.653]	[0.411]
Sesame promoted	0.222	0.071	0.148	-0.151	-0.074
				(0.095)	(0.107)
				[0.118]	[0.492]
Mustard promoted	0.370	0.321	0.407	-0.049	0.037
				(0.131)	(0.135)
				[0.709]	[0.785]
# of observations (Farmer group)	27	28	27		
# of clusters (Farmer group)	27	28	27		
Omnibus F-stat [p]				0.61	0.62
				[0.690]	[0.682]

Table A11: Demonstrator characteristics

	Mean (SD)		Coef. (SE) [p]
	Non-demo	Demo	
	(1)	(2)	(2) - (1)
Age of HHH (Years)	48.50	48.59	0.09
	(12.04)	(12.22)	(1.05)
			[0.931]
Any pulses cultivation	0.220	0.240	0.020
	(0.415)	(0.428)	(0.032)
			[0.543]
Veteran cultivator	0.112	0.147	0.035
(Any promoted crop cultivation)	(0.315)	(0.355)	(0.021)
			[0.087]
Any pulses improved seed	0.054	0.072	0.018
	(0.226)	(0.258)	(0.020)
			[0.367]
Any irrigation	0.390	0.319	-0.071
	(0.488)	(0.467)	(0.058)
			[0.223]
Total plot area (ha)	0.171	0.145	-0.026
	(0.286)	(0.231)	(0.024)
			[0.285]
# of connections	9.06	10.61	1.54
	(5.94)	(6.17)	(0.90)
			[0.088]
# of observations (Farmer)	762	279	
# of clusters (Farmer group)	57	53	
Omnibus F-stat [p]			1.09
			[0.379]

Table A12: Demonstrator characteristics across treatment arms

	Mean (SD)			Coef. (SE) [p]	
	Regular demo	Shared demo	Decentralized demo	Difference	
	(1)	(2)	(3)	(2) - (1)	(3) - (1)
Age of HHH (Years)	44.17	46.58	50.13	2.41	5.97
	(11.28)	(11.77)	(12.33)	(2.45)	(2.47)
				[0.326]	[0.017]
Any pulses cultivation	0.500	0.247	0.201	-0.253	-0.299
	(0.511)	(0.434)	(0.402)	(0.124)	(0.113)
				[0.042]	[0.009]
Veteran cultivator	0.250	0.111	0.149	-0.139	-0.101
(Any promoted crop cultivation)	(0.442)	(0.316)	(0.358)	(0.100)	(0.098)
				[0.166]	[0.303]
Any pulses improved seed	0.208	0.062	0.057	-0.147	-0.151
	(0.415)	(0.242)	(0.233)	(0.091)	(0.087)
				[0.110]	[0.084]
Any irrigation	0.250	0.321	0.328	0.071	0.078
	(0.442)	(0.470)	(0.471)	(0.162)	(0.150)
				[0.662]	[0.605]
Total plot area (ha)	0.147	0.179	0.129	0.032	-0.018
	(0.174)	(0.262)	(0.221)	(0.058)	(0.048)
				[0.588]	[0.712]
# of connections	12.58	7.60	11.73	-4.98	-0.85
	(7.10)	(4.85)	(6.13)	(2.22)	(2.40)
				[0.025]	[0.722]
# of observations (Farmer)	24	81	174		
# of clusters (Farmer group)	13	19	21		
Omnibus F-stat [p]				3.06	2.75
				[0.014]	[0.023]

Table A13: Gains from demonstration and decentralization

	Gains from demonstration ('000 BDT/ha)		
	Any demo in village − Control	Shared demo in village − Regular demo in village	Decentralized demo in village − Regular demo in village
	(1)	(2)	(3)
One year post-demonstration			
ex-ante	2.28	1.13	3.14
ex-post	-1.67	-0.79	-2.06
observed	-1.66	1.11	-1.14
	(0.68)	(0.92)	(1.09)
Two years post-demonstration			
ex-ante	21.04	0	0
ex-post	0.03	0	0
observed	1.32	0.24	0.02
	(1.27)	(1.24)	(1.45)
One year post-demonstration, counterfactual learning environment			
ex-ante, no social learning	1.14	0.72	1.54
ex-ante, learning only from scale	0.81	0.00	0.17
ex-ante, initial learning = 0.5	0.53	0.25	0.72
ex-post, initial learning = 0.5	0.01	0.01	0.02

Notes: Gains from demonstration and decentralizing demonstration are presented in this table, at one and two years post-demonstration and under different counterfactuals. Ex-ante and ex-post gains from demonstration are differences in ex-ante and ex-post profits (Equations 12 and 11), respectively, across treatment arms; details of the structural model used to calculate these values and its estimation are presented in Appendix D. Observed gains from demonstration are observed differences in average profits, with robust standard errors clustered at the farmer group level in parentheses.

Table A14: Dynamics of impacts of any demonstration are statistically significant, while falling adoption in control is not

	Promoted crops	
	One and two years post-demo	
	Improved seed	
	(1)	(2)
Intercept	0.036	
	(0.013)	
	[0.005]	
Two years post-demo	-0.014	-0.010
	(0.013)	(0.013)
	[0.295]	[0.438]
Any demo in village	0.056	
	(0.017)	
	[0.001]	
Two years post-demo × Any demo in village	-0.030	-0.043
	(0.018)	(0.019)
	[0.093]	[0.025]
Farmer FE		X
# of observations (Farmer-crop-wave)	5,533	5,533
# of clusters (Farmer group)	96	96

Table A15: Limited impacts of demonstration on placebo crops

	Placebo crops, One and two years post-demo				
	Improved seed	Culti-vation	Yield ('000 BDT/ha)	Fertilizer ('000 BDT/ha)	Profit ('000 BDT/ha)
	(1)	(2)	(3)	(4)	(5)
Any demo in village	0.014	0.025	0.862	0.172	0.334
	(0.005)	(0.018)	(3.704)	(1.058)	(0.293)
	[0.005]	[0.173]	[0.816]	[0.871]	[0.253]
Control mean	0.006	0.057	35.851	5.371	0.371
Wave FE	X	X	X	X	X
# of observations (Farmer-crop-wave)	9,562	9,562	727	727	9,562
# of clusters (Farmer group)	105	105	88	88	105

Table A16: No impacts of decentralizing demonstration on placebo crops

	Placebo crops					
	One year post-demo		Two years post-demo		One and two years post-demo	
	Improved seed					
	(1)	(2)	(3)	(4)	(5)	(6)
Any demo in village	0.017	0.023	0.009	0.011	0.014	0.018
	(0.004)	(0.008)	(0.008)	(0.011)	(0.005)	(0.008)
	[0.000]	[0.004]	[0.270]	[0.314]	[0.005]	[0.016]
Shared demo in village		-0.006		0.001		-0.003
		(0.011)		(0.012)		(0.010)
		[0.591]		[0.915]		[0.774]
Decentralized demo in village		-0.012		-0.007		-0.010
		(0.009)		(0.009)		(0.008)
		[0.188]		[0.396]		[0.181]
Control mean	0.004	0.004	0.009	0.009	0.006	0.006
Wave FE	X	X	X	X	X	X
# of observations (Farmer-crop-wave)	5,786	5,786	3,776	3,776	9,562	9,562
# of clusters (Farmer group)	105	105	75	75	105	105

Table A17: Pre-demonstration outcomes are uncorrelated with demonstration

	Promoted crops, Pre-demo year					
	Improved seed	Culti-vation	Yield ('000 BDT/ha)	Fertilizer ('000 BDT/ha)	Profit ('000 BDT/ha)	Paddy cultivation
	(1)	(2)	(3)	(4)	(5)	(6)
Any demo in village	-0.008	-0.004	0.143	-0.199	-0.104	-0.069
	(0.014)	(0.036)	(11.107)	(1.485)	(0.876)	(0.113)
	[0.564]	[0.911]	[0.990]	[0.893]	[0.905]	[0.544]
Control mean	0.027	0.067	41.175	4.470	2.090	0.457
Wave FE	X	X	X	X	X	X
# of observations (Farmer-crop-wave)	2,552	2,552	164	164	2,552	1,308
# of clusters (Farmer group)	71	71	39	39	71	71

Table A18: Pre-demonstration outcomes are uncorrelated with decentralizing demonstration

	Promoted crops, Pre-demo year	
	Improved seed	
	(1)	(2)
Any demo in village	-0.008	-0.003
	(0.014)	(0.016)
	[0.564]	[0.837]
Shared demo in village		-0.006
		(0.011)
		[0.565]
Decentralized demo in village		-0.008
		(0.010)
		[0.443]
Control mean	0.027	0.027
Wave FE	X	X
# of observations (Farmer-crop-wave)	2,552	2,552
# of clusters (Farmer group)	71	71

Table A19: Impacts of demonstration are robust to restricting to pre-demo year sample

	Promoted crops, One and two years post-demo					
	Improved seed	Cultivation	Yield ('000 BDT/ha)	Fertilizer ('000 BDT/ha)	Profit ('000 BDT/ha)	Paddy cultivation
	(1)	(2)	(3)	(4)	(5)	(6)
Any demo in village	0.049	0.040	-1.936	2.328	-0.194	-0.113
	(0.013)	(0.043)	(4.136)	(1.080)	(0.838)	(0.113)
	[0.000]	[0.359]	[0.640]	[0.032]	[0.817]	[0.319]
Estimate with full sample	0.045	0.047	-3.401	2.019	-0.572	-0.111
Control mean	0.027	0.133	43.528	5.038	2.118	0.423
Wave FE	X	X	X	X	X	X
# of observations (Farmer-crop-wave)	4,422	4,422	729	729	4,422	2,253
# of clusters (Farmer group)	70	70	63	63	70	70

Table A20: Impacts of decentralizing demonstration are robust to restricting to pre-demo year sample

	Promoted crops					
	One year post-demo		Two years post-demo		One and two years post-demo	
	Improved seed					
	(1)	(2)	(3)	(4)	(5)	(6)
Any demo in village	0.071	0.038	0.025	0.022	0.049	0.031
	(0.020)	(0.025)	(0.012)	(0.014)	(0.013)	(0.016)
	[0.000]	[0.129]	[0.035]	[0.116]	[0.000]	[0.061]
Shared demo in village		0.020		0.017		0.018
		(0.028)		(0.016)		(0.019)
		[0.475]		[0.290]		[0.338]
Decentralized demo in village		0.071		-0.006		0.033
		(0.033)		(0.015)		(0.019)
		[0.031]		[0.675]		[0.092]
Estimates with full sample	0.056	0.028	0.026	0.022	0.045	0.027
		0.018		0.019		0.019
		0.065		-0.006		0.035
Control mean	0.032	0.032	0.023	0.023	0.027	0.027
Wave FE	X	X	X	X	X	X
# of observations (Farmer-crop-wave)	2,223	2,223	2,199	2,199	4,422	4,422
# of clusters (Farmer group)	70	70	70	70	70	70

Table A21: Impacts of decentralizing demonstration are robust to controlling for number of connections

	Promoted crops					
	One year post-demo		Two years post-demo		One and two years post-demo	
	Improved seed					
	(1)	(2)	(3)	(4)	(5)	(6)
Any demo in village	0.071	0.035	0.025	0.021	0.049	0.028
	(0.019)	(0.024)	(0.012)	(0.014)	(0.013)	(0.016)
	[0.000]	[0.147]	[0.030]	[0.139]	[0.000]	[0.072]
Shared demo in village		0.029		0.021		0.025
		(0.027)		(0.017)		(0.020)
		[0.284]		[0.229]		[0.204]
Decentralized demo in village		0.071		-0.006		0.033
		(0.033)		(0.015)		(0.020)
		[0.029]		[0.695]		[0.092]
# of connections	0.002	0.002	0.001	0.001	0.001	0.002
	(0.002)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
	[0.131]	[0.118]	[0.468]	[0.242]	[0.137]	[0.094]
Estimates with full sample, no control	0.056	0.028	0.026	0.022	0.045	0.027
		0.018		0.019		0.019
		0.065		-0.006		0.035
Control mean	0.032	0.032	0.023	0.023	0.027	0.027
Wave FE	X	X	X	X	X	X
# of observations (Farmer-crop-wave)	2,223	2,223	2,199	2,199	4,422	4,422
# of clusters (Farmer group)	70	70	70	70	70	70

Table A22: Impacts of demonstration are robust to crop-wave fixed effects

	Promoted crops, One and two years post-demo					
	Improved seed	Culti-vation	Yield ('000 BDT/ha)	Fertilizer ('000 BDT/ha)	Profit ('000 BDT/ha)	Paddy cultivation
	(1)	(2)	(3)	(4)	(5)	(6)
Any demo in village	0.044	0.046	-4.448	1.258	-0.613	-0.111
	(0.011)	(0.027)	(3.037)	(0.690)	(0.585)	(0.095)
	[0.000]	[0.090]	[0.143]	[0.068]	[0.295]	[0.245]
Estimate with Wave FE	0.045	0.047	-3.401	2.019	-0.572	-0.111
Control mean	0.031	0.130	42.746	5.149	1.874	0.438
Crop-wave FE	X	X	X	X	X	X
# of observations (Farmer-crop-wave)	5,533	5,533	912	912	5,533	2,818
# of clusters (Farmer group)	96	96	81	81	96	96

Table A23: Impacts of decentralizing demonstration are robust to crop-wave fixed effects

	Promoted crops					
	One year post-demo		Two years post-demo		One and two years post-demo	
	Improved seed					
	(1)	(2)	(3)	(4)	(5)	(6)
Any demo in village	0.055	0.034	0.027	0.027	0.044	0.032
	(0.016)	(0.017)	(0.011)	(0.014)	(0.011)	(0.012)
	[0.001]	[0.053]	[0.019]	[0.056]	[0.000]	[0.010]
Shared demo in village		0.003		0.009		0.006
		(0.018)		(0.017)		(0.015)
		[0.863]		[0.582]		[0.691]
Decentralized demo in village		0.057		-0.008		0.029
		(0.025)		(0.015)		(0.016)
		[0.022]		[0.598]		[0.076]
Estimates with Wave FE	0.056	0.028	0.026	0.022	0.045	0.027
		0.018		0.019		0.019
		0.065		-0.006		0.035
Control mean	0.036	0.036	0.022	0.022	0.031	0.031
Crop-wave FE	X	X	X	X	X	X
# of observations (Farmer-crop-wave)	3,289	3,289	2,244	2,244	5,533	5,533
# of clusters (Farmer group)	96	96	70	70	96	96

Table A24: Flood exogeneity

	Promoted crops, Demo year and one and two years post-demo			
	Cultivation	Improved seed	Improved seed | Cultivation = 1	Fertilizer ('000 BDT/ha) | Cultivation = 1
	(1)	(2)	(3)	(4)
Above median flood	0.014	0.000	-0.032	0.366
	(0.020)	(0.010)	(0.049)	(0.716)
	[0.486]	[0.985]	[0.509]	[0.610]
Control mean	0.171	0.071	0.418	6.563
Crop-wave FE	X	X	X	X
# of observations (Farmer-crop-wave)	7,161	7,161	1,336	1,336
# of clusters (Farmer group)	96	96	90	90

Table A25: Flood first stage and serial correlation

	Promoted crops, Demo year and one and two years post-demo	
	Reported flood exposure	Above median flood
	(1)	(2)
Above median flood	0.051	
	(0.010)	
	[0.000]	
Above median flood (1 year lag)		0.059
		(0.071)
		[0.403]
Control mean	0.012	0.326
Crop-wave FE	X	X
# of observations (Farmer-crop-wave)	6,753	4,596
# of clusters (Farmer group)	96	96

Table A26: Demonstration year floods increase adoption of improved seed one year post-demonstration, but not contemporaneously or two years post-demonstration

	Promoted crops					
	Demo year		One year post-demo		Two years post-demo	
	Improved seed					
	(1)	(2)	(3)	(4)	(5)	(6)
Above median flood (Demo year)	0.041	0.010	0.042	0.008	0.006	-0.008
	(0.034)	(0.019)	(0.016)	(0.021)	(0.018)	(0.016)
	[0.217]	[0.594]	[0.008]	[0.704]	[0.725]	[0.600]
Any demo in village		0.127		0.025		0.031
		(0.034)		(0.019)		(0.019)
		[0.000]		[0.190]		[0.106]
Above median flood (Demo year) × Any demo in village		0.014		0.038		0.009
		(0.041)		(0.031)		(0.026)
		[0.731]		[0.222]		[0.745]
[Above median flood (Demo year)] + [Above median flood (Demo year) × Any demo in village]		0.024		0.046		0.000
		(0.037)		(0.022)		(0.022)
		[0.522]		[0.037]		[0.984]
Control mean	0.098	0.000	0.033	0.000	0.035	0.000
Crop FE	X	X	X	X	X	X
# of observations (Farmer-crop-wave)	1,642	1,642	1,429	1,429	836	836
# of clusters (Farmer group)	96	96	96	96	70	70

Table A27: Impacts of demonstration and decentralizing demonstration are not explained by seed recycling

	Promoted crops					
	One year post-demo		Two years post-demo		One and two years post-demo	
	Improved seed (Recycled)					
	(1)	(2)	(3)	(4)	(5)	(6)
Any demo in village	0.004	-0.003	0.004	0.001	0.004	-0.001
	(0.009)	(0.009)	(0.003)	(0.003)	(0.006)	(0.006)
	[0.620]	[0.757]	[0.256]	[0.740]	[0.483]	[0.836]
Shared demo in village		0.012		0.005		0.009
		(0.011)		(0.005)		(0.008)
		[0.286]		[0.320]		[0.254]
Decentralized demo in village		0.009		0.002		0.007
		(0.008)		(0.004)		(0.005)
		[0.242]		[0.575]		[0.209]
Control mean	0.013	0.013	0.002	0.002	0.010	0.010
Wave FE	X	X	X	X	X	X
# of observations (Farmer-crop-wave)	3,289	3,289	2,244	2,244	5,533	5,533
# of clusters (Farmer group)	96	96	70	70	96	96

Table A28: Balance with respect to any demonstration is robust to randomization inference

	Mean (SD)		Coef. (SE) [p]
	Control	Any demo in village	Difference
	(1)	(2)	(2) - (1)
Age of HHH (Years)	48.69	48.51	-0.18
	(11.93)	(12.10)	(0.83)
			[0.832]
			{0.854}
Any pulses cultivation	0.203	0.215	0.012
	(0.403)	(0.411)	(0.064)
			[0.850]
			{0.840}
Veteran cultivator	0.117	0.115	-0.003
(Any promoted crop cultivation)	(0.323)	(0.319)	(0.061)
			[0.965]
			{0.959}
Any pulses improved seed	0.050	0.056	0.006
	(0.218)	(0.229)	(0.027)
			[0.834]
			{0.814}
Any irrigation	0.463	0.398	-0.065
	(0.499)	(0.490)	(0.111)
			[0.561]
			{0.552}
Total plot area (ha)	0.202	0.178	-0.024
	(0.300)	(0.282)	(0.039)
			[0.537]
			{0.569}
# of connections	9.05	9.42	0.37
	(6.65)	(5.97)	(1.62)
			[0.817]
			{0.815}
# of observations (Farmer)	281	1098	
# of clusters (Farmer group)	15	61	
Omnibus F-stat [p]			0.15
			[0.994]
			{0.998}

Table A29: Balance with respect to decentralizing demonstration is robust to randomization inference

	Mean (SD)			Coef. (SE) [p]	
	Regular demo in village	Shared demo in village	Decentralized demo in village	Difference	
	(1)	(2)	(3)	(2) - (1)	(3) - (1)
Age of HHH (Years)	48.71	48.32	48.53	-0.39	-0.18
	(11.91)	(11.85)	(12.56)	(0.97)	(1.14)
				[0.688]	[0.872]
				{0.706}	{0.852}
Any IAPP cultivation	0.240	0.246	0.160	0.006	-0.080
	(0.428)	(0.431)	(0.367)	(0.068)	(0.056)
				[0.934]	[0.153]
				{0.930}	{0.203}
Veteran cultivator	0.100	0.126	0.117	0.026	0.018
(Any promoted crop cultivation)	(0.300)	(0.332)	(0.322)	(0.046)	(0.041)
				[0.573]	[0.667]
				{0.620}	{0.740}
Any IAPP improved seed	0.076	0.052	0.040	-0.024	-0.036
	(0.266)	(0.223)	(0.196)	(0.028)	(0.027)
				[0.392]	[0.181]
				{0.369}	{0.179}
Any irrigation	0.367	0.406	0.419	0.039	0.052
	(0.483)	(0.492)	(0.494)	(0.118)	(0.119)
				[0.740]	[0.661]
				{0.736}	{0.662}
Total plot area (ha)	0.163	0.198	0.171	0.035	0.008
	(0.222)	(0.246)	(0.354)	(0.039)	(0.050)
				[0.369]	[0.873]
				{0.453}	{0.866}
# of connections	10.75	7.19	10.50	-3.56	-0.25
	(6.43)	(5.04)	(5.76)	(1.57)	(1.75)
				[0.024]	[0.888]
				{0.031}	{0.886}
# of observations (Farmer)	341	382	375		
# of clusters (Farmer group)	19	21	21		
Omnibus F-stat [p]				1.33	0.96
				[0.263]	[0.470]
				{0.416}	{0.630}

Table A30: Statistical precision of impacts of demonstration are robust to randomization inference, with exception of impacts on fertilizer use

	Promoted crops, One and two years post-demo					
	Improved seed	Culti-vation	Yield ('000 BDT/ha)	Fertilizer ('000 BDT/ha)	Profit ('000 BDT/ha)	Paddy cultivation
	(1)	(2)	(3)	(4)	(5)	(6)
Any demo in village	0.045	0.047	-3.401	2.019	-0.572	-0.111
	(0.013)	(0.034)	(3.344)	(0.884)	(0.681)	(0.095)
	[0.000]	[0.163]	[0.310]	[0.023]	[0.401]	[0.245]
	{0.006}	{0.135}	{0.462}	{0.127}	{0.501}	{0.231}
Control mean	0.031	0.130	42.746	5.149	1.874	0.438
Wave FE	X	X	X	X	X	X
# of observations (Farmer-crop-wave)	5,533	5,533	912	912	5,533	2,818
# of clusters (Farmer group)	96	96	81	81	96	96

Table A31: Statistical precision of impacts of decentralizing demonstration are robust to randomization inference

	Promoted crops					
	One year post-demo		Two years post-demo		One and two years post-demo	
	Improved seed					
	(1)	(2)	(3)	(4)	(5)	(6)
Any demo in village	0.056	0.028	0.026	0.022	0.045	0.027
	(0.017)	(0.020)	(0.012)	(0.014)	(0.013)	(0.015)
	[0.001]	[0.167]	[0.026]	[0.106]	[0.000]	[0.068]
	{0.014}	{0.337}	{0.064}	{0.190}	{0.007}	{0.198}
Shared demo in village		0.018		0.019		0.019
		(0.022)		(0.016)		(0.017)
		[0.422]		[0.236]		[0.278]
		{0.522}		{0.246}		{0.346}
Decentralized demo in village		0.065		-0.006		0.035
		(0.027)		(0.014)		(0.018)
		[0.018]		[0.658]		[0.057]
		{0.016}		{0.702}		{0.074}
Control mean	0.036	0.036	0.022	0.022	0.031	0.031
Wave FE	X	X	X	X	X	X
# of observations (Farmer-crop-wave)	3,289	3,289	2,244	2,244	5,533	5,533
# of clusters (Farmer group)	96	96	70	70	96	96

Table A32: Conditional convergence across demonstration modalities

	Promoted crops			
	One year post-demo		Two years post-demo	
	Improved seed			
	(1)	(2)	(3)	(4)
Veteran cultivator	0.067	0.067	0.040	0.040
	(0.029)	(0.029)	(0.014)	(0.014)
	[0.022]	[0.022]	[0.005]	[0.005]
Any demo in village	0.077	0.041	0.020	0.015
	(0.018)	(0.024)	(0.009)	(0.011)
	[0.000]	[0.081]	[0.030]	[0.180]
Shared demo in village		0.026		0.022
		(0.028)		(0.012)
		[0.354]		[0.072]
Decentralized demo in village		0.073		-0.004
		(0.034)		(0.011)
		[0.029]		[0.721]
Veteran cultivator × Any demo in village	-0.045	-0.028	0.029	0.043
	(0.035)	(0.037)	(0.022)	(0.032)
	[0.195]	[0.449]	[0.185]	[0.189]
Veteran cultivator × Shared demo in village		-0.035		-0.032
		(0.038)		(0.044)
		[0.358]		[0.465]
Veteran cultivator × Decentralized demo in village		-0.012		-0.009
		(0.043)		(0.038)
		[0.777]		[0.815]
[Any demo in village] + [Veteran cultivator × Any demo in village]	0.032	0.013	0.049	0.057
	(0.038)	(0.044)	(0.024)	(0.035)
	[0.403]	[0.768]	[0.047]	[0.107]
[Shared demo in village] + [Veteran cultivator × Shared demo in village]		-0.009		-0.010
		(0.046)		(0.047)
		[0.842]		[0.828]
[Decentralized demo in village] + [Veteran cultivator × Decentralized demo in village]		0.061		-0.013
		(0.048)		(0.041)
		[0.208]		[0.755]
Control mean	0.023	0.023	0.018	0.018
Wave FE	X	X	X	X
# of observations (Farmer-crop-wave)	2,223	2,223	2,199	2,199
# of clusters (Farmer group)	70	70	70	70

F Appendix Figures

Figure A1: Village map

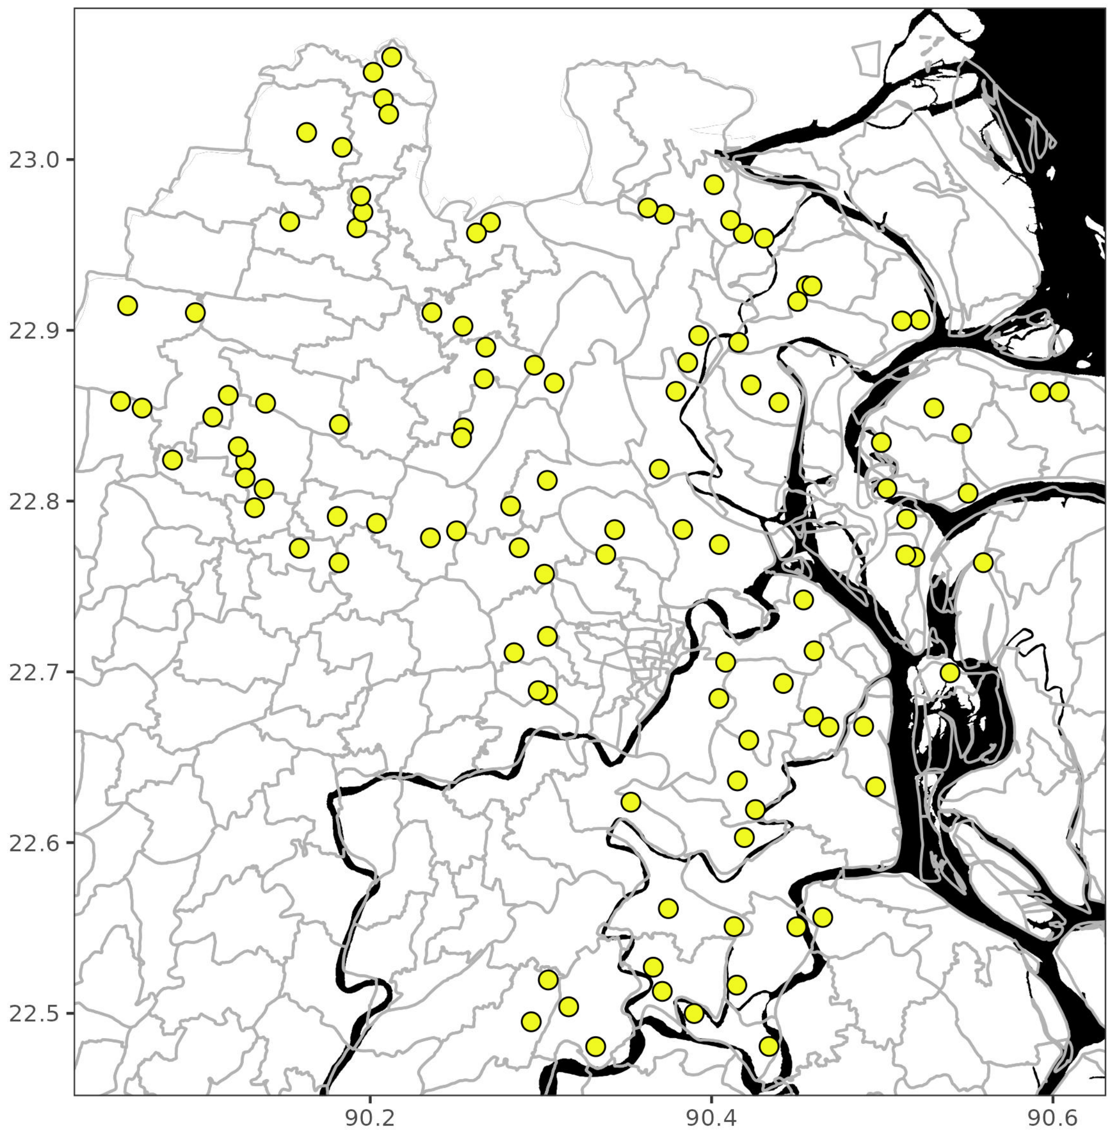

Notes: Our experimental sample of villages is plotted in yellow, with Union boundaries of Barisal in gray.

Figure A2: Floods