The cost of favoritism in network-based markets

Kyle Emerick*

March 3, 2017

Abstract

This paper presents evidence showing that favoritism of close peers severely limits the ability of social networks to allocate a new agricultural technology. I use a three-year field experiment where a new and profitable seed variety was first delivered to a random subset of farmers in 82 Indian villages. The diffusion between farmers is then compared with demand that was revealed experimentally one year later. I find a large gap between adoption from neighbors and demand and this gap is largest for farmers that are the most socially distant from the original seed recipients. Applying an estimate of the causal impact of the technology to this adoption gap shows the costliness of caste-based favoritism. Specifically, the cost to the average farmer over the three-year period where diffusion was tracked is equivalent to about 20 days of casual work.

Keywords: Social Networks, Technology Adoption
JEL Codes: O17, L14

*Department of Economics Tufts University, 8 Upper Campus Rd, Medford, MA 02155. kyle.emerick@tufts.edu.
I am grateful to Elisabeth Sadoulet and Alain de Janvry for numerous discussions and feedback. I also thank Ted Miguel, Craig Mcintosh, Jeremy Magruder, Emily Oster, Kelsey Jack, Seema Jayachandran, Sofia Villas-Boas, Marshall Burke, Judson Boomhower, Ethan Ligon, and Michael Anderson for helpful comments and suggestions. I gratefully acknowledge financial support from the World Bank and the Agricultural Technology Adoption Initiative of the Bill and Melinda Gates Foundation. I thank Manzoor Dar for his help in setting up the experiment and sharing his abundant knowledge about rice production. I also thank Girija Prasad Swain, Anirudha Patra, Silu Saha, and the staff of the Balasore Social Service Society for their help with field work.

1 Introduction

Social relationships often guide the allocation of goods and services in developing countries. Common examples include bank credit (Fisman, Paravisini, and Vig, 2017), informal insurance (Mazzocco and Saini, 2012; Munshi and Rosenzweig, 2016), irrigation water (Anderson, 2011), and even information about new products or technologies (Conley and Udry, 2010; Banerjee et al., 2013a; Beaman et al., 2015; BenYishay and Mobarak, 2015; Cai, de Janvry, and Sadoulet, 2015). This practice may arise as a mechanism to alleviate information asymmetries, particularly in credit or insurance markets. Alternatively, it may reflect preferences for interaction with socially proximate individuals, such as within-caste marriage in India (Banerjee et al., 2013b) or favoritism of connected workers by firm managers (Bandiera, Barankay, and Rasul, 2009). This favoritism-based explanation potentially implies large costs if the end allocation is based on network connections rather than actual benefits.

The contribution of this paper is to estimate the magnitude of the economic losses that arise from favoritism based on social proximity. I take advantage of the fact that farmers often share seeds in remote parts of developing countries. These informal transactions are guided by network relationships — largely caste in my setting. I first provide a small amount of a profitable flood-tolerant rice seed called "Swarna-Sub1" to five farmers in each of 82 villages.[1] These five "original recipients" plant the seeds and produce enough output to diffuse to other farmers as seeds for the following year. Do the village seed markets show favoritism based on caste? If so, how costly is this favoritism? To answer these, I revealed demand in half the sample by selling the same seed using door-to-door sales. This arm of the experiment provides a revealed-preference benchmark of demand and allows me to compare demand with realized adoption, both for farmers that are more or less socially proximate to the original recipients.

I find a substantial gap between revealed demand and adoption in farmer networks. The magnitude of this effect is striking. Compared to revealed demand of around 40 percent, only 7 percent of farmers adopt after one year in villages where the seed diffuses directly between farmers. This gap closes over time, but remains significant after two additional seasons. I then apply recent experimental estimates of the technology's impact (Emerick et al., 2016) to this estimated adoption gap in order to estimate the losses from trading in farmer networks. This estimate suggests that over the 3-year period of the study, the average farmer loses an amount equivalent to about 20 days of casual work, or about 8 percent of the annual rice harvest.

The paper then turns to understanding why diffusion in networks performs poorly relative to actual demand. My primary results are consistent with caste-based favoritism. Specifically, original recipients tended to share seeds only with farmers that belonged to their same sub-caste or shared their same surname. The random selection of original recipients allows for estimation of these effects. In my preferred specification, having the same surname as an additional original recipient

[1] Recent experimental work has shown that Swarna-Sub1 is profitable for two reasons. First, it increases crop yield during flooding without affecting yield in normal years (Dar et al., 2013). Second, it also generates welfare gains by inducing farmers to invest more in inputs such as fertilizer and labor at planting (Emerick et al., 2016).

results in an approximate doubling in the one-year probability of adoption. Similarly, being part of the same sub-caste as an additional recipient leads to an approximate 53 percent increase in adoption. I observe no such peer effects in villages where demand was revealed with door-to-door sales. The result provides micro-level evidence that is consistent with the cross-country result that the diffusion of technology is slower in countries where networks are organized into distinct sub-networks or collectives (Fogli and Veldkamp, 2016). Additionally, the result empirically demonstrates the importance of network structure for trading outcomes — something that is consistent with results from laboratory experiments (Charness, Corominas-Bosch, and Frechette, 2007; Gale and Kariv, 2009).

Consistent with these findings, the welfare losses from network-based diffusion are smaller for farmers with better connections to the original recipients. As one example, I consider a farmer that shares the same surname with five other members of the sample. I first consider the unconnected farmer that shares the same surname with none of the five original recipients. The regression estimates indicate that this farmer loses 82.6 dollars in revenue over the three year period under the network model of diffusion. These predicted losses decrease by 60 percent to 32.7 dollars if the farmer instead shared the same surname with two original recipients. Thus, social connections were valuable in this context because they provided access to a profitable new technology.

In contrast to recent studies where network-based transactions resolve information asymmetries or provide enforcement of bilateral contracts (Anderson, 2011; Fisman, Paravisini, and Vig, 2017), the seed markets I study involve one-off transactions that minimize these issues of moral hazard and contract enforcement. As noted in Fisman, Paravisini, and Vig (2017), the competing explanation is favoritism based on preferences for within-network interaction. My findings therefore document and quantify the importance of caste-based favoritism as a friction that hinders trading in the village economy.

I then consider a number of possible alternative explanations. I start by documenting that original recipients had nothing to gain — and actually lost — by not diffusing seeds to other farmers. About one third of the harvest of original recipients — an amount larger than the total amount sold in the door-to-door sales — was sold to traders at an average price of 10.3 Rupees per kilogram. The door-to-door sales experiment showed that over 30 percent of farmers in the same village were willing to pay 14 Rupees per kilogram for the same seed. These basic numbers suggest a misallocation and rule out the explanation that alternative uses were more beneficial to original recipients.

I also show that price differences are unlikely to explain the findings. The prices for door-to-door sales were randomized at three levels: 10, 12, or 14 Rupees per kilogram. As mentioned above, the highest price is larger than the output price of rice. It is also larger than the implied price for farmer-to-farmer trades. Observing demand at the three price levels allows me to rule out that the seeds were simply cheaper in the door-to-door sales.

Another possibility is that there simply was not enough output to meet the demand of other farmers. A straightforward accounting exercise rules this out. The amount of seed planted in the second year of the study across the door-to-door villages averaged around 150 kilograms per village. This amounts to less than 10 percent of the amount harvested by original recipients. These numbers indicate that the slow diffusion between farmers cannot be merely an issue of capacity. In addition to these concerns about quantity, I also show that the results are less likely to be driven by salience effects or differences in seed quality.

I close the analysis by using satellite observations to show that adopters benefitted directly from the new seed. The 2013 season — immediately following the door-to-door sales — had variation in flood exposure across the sample. Using this variation, along with high-resolution satellite images of crop productivity, the data show that fields cultivated by adopters appear 8 percent "greener" to satellites.[2] This effect exists only during the growing season and only amongst fields that were affected by flooding — suggesting that the observed productivity effects of Swarna-Sub1 are not due to fixed differences between land types or differences in other fixed characteristics of adopters and non-adopters. Rather, these triple difference estimates confirm the value of the technology and provide further direct evidence on the costliness of the slow diffusion induced by caste-based favoritism.

The paper contributes mostly to the literature on how ethnicity and caste effect market outcomes in developing countries. There are several examples within India. Nagavarapu and Sekhri (2016) show that poor households are more likely to obtain subsidized food grains when shopkeepers are part of their same caste. Similarly, Fisman, Paravisini, and Vig (2017) find that matches in caste between borrowers and lending agents increases credit uptake in Indian banks. Anderson (2011) finds that rural groundwater markets are segmented by caste. Caste relationships also guide the provision of informal insurance (Mazzocco and Saini, 2012; Munshi and Rosenzweig, 2016). Outside of India, ethnicity in the Democratic Republic of the Congo provides a substitute enforcement mechanism for more formal government contracts (de la Sierra, 2017). In contrast to contract enforcement and monitoring, I show the importance of favoritism as an explanation for the concentration of transactions within castes or ethnicities. By revealing actual demand in a random set of villages, the main contribution of my paper is to quantify the magnitude of this effect.

This paper also contributes to the literature on the barriers to the adoption of agricultural technologies in developing countries (Jack, 2011). The important barriers emphasized in the literature include limitations to demand such as self control (Duflo, Kremer, and Robinson, 2011) and uninsured risk (Karlan et al., 2014; Emerick et al., 2016). Suri (2011) shows that constraints on the *supply* of technologies are also important. Specifically, farmers with the highest returns to hybrid maize in Kenya are those that have the highest cost of adoption due to low quality infrastructure. My findings also suggest the importance of supply barriers.

Turning to policy, the results demonstrate that network-based favoritism limits the effectiveness of the informal seed system in developing countries. Besides this paper on rice in India, direct transfers of seeds between farmers have been studied for cassava in Gabon (Delêtre, McKey, and

[2] Greenness is measured using the normalized difference vegetation index (NDVI) — a measure of vegetation density that is strongly correlated with crop yield in rice (Huang et al., 2013).

Hodkinson, 2011), maize in Mexico (Bellon, Hodson, and Hellin, 2011), and sorghum in Kenya (Labeyrie et al., 2016). McGuire and Sperling (2016) use primary data from 5 African countries to show that seed transfers from neighbors account for 6 percent of total seed supply in Kenya and this ranges to 22 percent in Zimbabwe. Given the push to make development interventions sustainable (Kremer and Miguel, 2007), and the absence of the formal seed market in many remote areas of developing countries, relying on seed transfers between farmers seems ideal because of its low cost. I demonstrate that this approach of identifying entry points and relying on social networks for seed diffusion fails to meet demand, at least in the short to medium run.

The rest of the paper is organized as follows. In section 2, I provide a more detailed description of the experimental design. Section 3 provides a simple model of diffusion that guides some of the empirical analysis in section 4. After establishing these main results, section 5 provides further analysis that considers the most likely explanations of the findings. Section 6 concludes.

2 Experimental Design and Data Collection

This section describes the details of the experimental design. I start by discussing the key properties of the new seed variety and some details on its availability. The section then outlines the steps of the experiment in chronological order. I focus on how the different steps of the experiment allow me to quantity the costs of slow diffusion in social networks.

The technology

The rice variety Swarna-Sub1 is suitable for the experiment for three reasons. First, results from a randomized experiment in the same part of India demonstrate its' profitability (Emerick et al., 2016). Those findings show that around 40 percent of the seed's returns are due to management effects: the seed reduces downside risk to farmers inducing them to increase investment on their fields. Specifically, the reduction in risk induces farmers to use more fertilizer, invest more labor at planting, and to avail of agricultural credit. These proven benefits rule out the simple explanation that the technology diffuses slowly because it is unprofitable.

Second, the flood tolerance property allows me to characterize agronomic returns based solely on past flooding. I use this observable heterogeneity in returns to test whether targeting is any more or less effective with network-based diffusion.

Third, Swarna-Sub1 was released in India only in 2009 and was not available when the experiment started. Farmers in the sample often obtain new seeds from local government offices. Seed supply in this channel is controlled by the state-run seed corporation and it takes them several years to multiply enough seed for all of the state's flood-prone areas. As a result, the seed was largely unavailable at the time of the long-term followup in 2015. The lack of availability from other sources ensures that diffusion between farmers is not affected by the presence of an additional option of obtaining seeds from other sources.

Overview of the experiment

The experiment was carried out in 82 villages in three blocks of the Bhadrak district of Odisha.[3] This region is a coastal low-lying area where flooding is frequent. Specifically, the villages selected into the study were flooded in either 2008 or 2011, based on satellite imagery. A majority of farmers in these villages grow a single rice crop during the Kharif (wet season) from June to December.[4] In addition, Swarna (the variety that is dominated by Swarna-Sub1) is widely grown. All of these characteristics are important because they combine to indicate that Swarna-Sub1 is suitable for the sample villages.

Each village was visited in May 2012 and farmers were invited to a meeting to discuss Swarna-Sub1. The meetings were open to any farmers cultivating rice and were attended by anywhere from 15 to 41 farmers, with average attendance being 22 — or approximately 22 percent of households.[5] During each meeting, enumerators provided a brief overview of the characteristics of Swarna-Sub1, described its similarity to the popular variety Swarna, and pointed to flood tolerance as its only known benefit. After the information was provided, each farmer was administered a short one-page baseline questionnaire. This questionnaire included a small set of household characteristics and a question where respondents were asked to list their closest social contacts.

The original recipients were selected at the end of the meeting via lottery. Attendees were informed that five farmers would be chosen to receive a five kilogram minikit. Minikits are a common approach to introducing a new seed variety in India (Bardhan and Mookherjee, 2011). Each minikit contained only five kg of Swarna-Sub1 seeds, which is enough to cultivate approximately 0.1-0.2 hectares. Despite being a small amount of seeds, 5 kilograms produces an average of around 300 kilograms after harvest. Each attendee then placed their name in a bucket and five names were selected. The seeds were then given to the winners immediately in front of all the meeting attendees. This approach was purposefully selected so that the identities of the five original recipients would be known to all participants. This eliminates the possibility that slow diffusion can be caused by lack of information on who was cultivating the new variety. Enumerators also informed farmers that Swarna-Sub1 would unlikely be available at local block offices where they most often buy new seeds. Therefore, any farmers that want to cultivate Swarna-Sub1 next year should purchase or exchange seeds with the original recipients.

The random selection of original recipients allows me to estimate peer effects and identify whether lack of connections to original recipients prevents adoption by farmers with demand. The tradeoff is that real-world methods of identifying injection points for new agricultural technology do not use randomization. The most common approach is to use local agricultural extension agents

[3]The total number of villages is 84. Two villages were used for piloting of surveys and interventions and are therefore not used in the analysis.

[4]26 percent of farmers have at least one plot that has access to lift irrigation, which is needed to cultivate dry-season crops.

[5]The households in the sample are fairly representative of the village. The average share of the population that is scheduled caste is 20 percent in both the sample and the matched 2001 census of villages. Average household size and male literacy are also similar between the sample and the census.

to identify "lead farmers".[6] Nonetheless, the randomization allows me to estimate the extent of diffusion when different types of farmers are selected as original recipients. I consider this in section 5.

The next step in the experiment was a followup visit with original recipients during the harvesting period in November of 2012. This visit had three purposes. First, enumerators verified that the crop had been planted. Most original recipients complied with the experiment by planting the minikit. Of the 396 farmers surveyed, 87 percent indicated that the minikit had been planted.[7] Second, original recipients were reminded of the potential opportunity to sell or trade Swarna-Sub1 seeds after the harvest. Third, enumerators collected a number of characteristics of original recipients to be used in heterogeneity analysis.

I then identified 15 non-recipients in each village by taking a random sample amongst the farmers that attended the meetings, but were not selected as original recipients.[8] This group of farmers serves as my main estimation sample for most of the analysis. The sample was drawn from the meeting attendees for two reasons. First, the meeting attendees witnessed the selection of original recipients and therefore focusing on attendees eases the concern that information about identities of original recipients can drive adoption patterns. Second, the benefits of Swarna-Sub1 were explained to meeting attendees. Focusing the sample of non-recipients on attendees therefore makes it less likely that non-adoption could be caused by lack of information on the variety or its benefits.

Enumerators then surveyed each of these 1,151 non-recipients during February-April 2013. There were three purposes of this survey. First, a plot-level record of flooding during the previous five years was collected in order to estimate the expected returns of the new technology. I return to the estimation of expected returns using these data below. Second, farmers were also reminded about the new variety and the potential to obtain it from other farmers in the village. These reminders limit the possibility that farmers chose not to adopt simply because they had forgotten or did not know about the technology. Third, all respondents were again informed about the flood tolerance benefits of Swarna-Sub1.

41 of the villages were randomly assigned to receive a door-to-door treatment where enumerators went to each of the 15 non-recipients and asked them if they were interested in purchasing Swarna-Sub1 seeds.[9] The door-to-door sales took place in late May and early June of 2013. Except for

[6]Beaman et al. (2015) show that using the structure of the social network to identify optimal injection points can increase adoption relative to this approach of using extension agents.

[7]14 of the 410 original recipients could not be reached because either the household had moved from the village or household members were away for work during survey visits. The most common reason reported for not cultivating the minikit was that the seedbed was damaged by drought or cows. The common method of planting rice in the area is transplanting, which involves preparing a small seedbed and uprooting the small seedlings approximately 3-4 weeks after emergence. The uprooted seedlings are then bundled and planted in the main field.

[8]All non-recipients were selected into this sample in villages where fewer than 20 farmers attended the meeting.

[9]The randomization of door-to-door sales was stratified by block — an administrative unit two levels above villages — and the relative importance of original recipients to non-recipients. Original recipients were defined as being relatively more important when the ratio of average degree (number of social contacts) of original recipients to the average degree of non-recipients was larger than the sample median. The degree is simply the number of links of a farmer, where two farmers are defined to be linked if either farmer stated that they would go to the other farmer

reminding farmers of the previous survey, enumerators gave farmers no additional details about Swarna-Sub1's benefits.

The purpose of these door-to-door sales was to compare revealed demand with adoption in the other 41 villages where diffusion between farmers was the only mechanism of exchange. One limitation of this approach is that door-to-door sales do not simulate a common or easily scalable method for selling seeds. At the same time, the door-to-door sales allow me to estimate demand in an environment without any frictions or costs to adoption. I return to this concern when calculating cost effectiveness in Section 4.5.

Prices are another important concern because the experiment requires the door-to-door offers to be made at prices that are equivalent to the prices paid in transactions between farmers. During piloting I found that farmers often exchange seeds of one variety for output of another.[10] The price in such a transaction corresponds to the farm-gate price of rice, which is the opportunity cost of the output that was exchanged for seed. The minimum support price set by the Indian government for the 2012-2013 season was 12.5 Rs per kilogram (1 USD $\approx$ 58 Rs).

Based on these output prices, the price for door-to-door sales was randomized at the village level to one of three values: 10, 12, or 14 Rs per kilogram. Importantly, I included the higher price of 14 Rs in order to estimate demand at a price that is *higher* than the output price of rice. The villages where demand was revealed at this price therefore allow me to establish whether there was demand at a price higher than prices for which Swarna-Sub1 was sold by original recipients.

Enumerators then revisited all 82 villages in July 2013 to track adoption — the main outcome variable of interest. Enumerators visited non-recipient households in all villages, including the door-to-door villages. Non-recipients in door-to-door villages could have also obtained seeds from original recipients and therefore an additional survey was needed to fully track their adoption status. In addition, carrying out the survey in all 82 villages ensures that adoption is measured in a uniform way across all villages, regardless of treatment status. I use adoption from this survey as the main outcome variable in most of the analysis.

In addition to this survey with non-recipients, enumerators were able to locate 394 of the original recipients. The main purpose of the survey with original recipients was to fully document the final uses of the Swarna-Sub1 output. Each original recipient was asked their harvest amount, amount sold as rice, amount sold or stored for consumption, amount saved for their own seed, and amount of seed transferred to others. Section 4.1 describes in more detail how the harvest of original recipients was allocated. Each original recipient was also asked for the identities of the farmers to which they transferred seeds. These data were used to verify the transactions that were reported by non-recipients.

A final survey was carried out two years later during July 2015. This second followup allows me to estimate longer term effects on take up during both the 2014 and 2015 wet seasons. These effects add to the analysis because it is expected that the adoption gap would dissipate over time as

for seeds, fertilizers, or other inputs.

[10]This pattern was true in the data as well. As I show in Section 4.1, 70 percent of the Swarna-Sub1 transactions reported by original recipients were 1:1 exchanges for other rice varieties.

more farmers start cultivating and sharing seeds with each other. The survey allows me to measure the difference between revealed demand and network-based diffusion over a three-year period. One notable difference is that the long-term followup included all rice-farming households, including those that did not participate in the village meetings and were thus not part of the original sample.

Summary Statistics

Table 1 verifies that the non-recipient farmers are similar in the network and door-to-door villages. In addition to this test of balance, Table 1 points out three notable features of the sample of non-recipients. First, Swarna is widely grown by around 70 percent of the sample. Swarna-Sub1 is superior to Swarna and therefore a maximum potential adoption rate of 70 percent is plausible. Second, the average non-recipient has 2.1 and 1.4 sub-caste and surname connections with original recipients, respectively. I return to these averages when interpreting the estimated peer effects in Section 4.2. Third, the average non-recipient did not have to travel far to obtain seeds, due to the small size of villages in the sample. Each respondent had just over one original recipient that lived within 50 meters from their household.

Seed sharing between farmers exists at baseline, although it is not the dominant method for obtaining seeds. 14 percent of farmers at baseline planted at least one plot with seeds that were obtained from a neighboring farmer. This figure stands at 66 percent for seeds saved from the previous harvest and 43 percent for seeds sold by local government offices. In other words, seeds from neighboring farmers account for one quarter of seed replacement, i.e. when farmers decide to plant new seeds. Therefore, sharing seeds is not a new practice for farmers.

I also use the short survey carried out during the village meeting to verify the similarity between non-recipients and original recipients (Table A1). These data also show that the characteristics of non-recipients and original recipients are similar in both network and door-to-door villages.[11]

3 A Simple Diffusion Model

In this section I formulate a model of technology adoption which guides the empirical analysis. The model has two features. First, network-based diffusion imposes some nonzero cost of adoption associated with trade between socially distant villagers. The literature on network-based trading has typically modeled these as costs of network formation, i.e. the costs of making links between buyers and sellers (Kranton and Minehart, 2001; Elliott, 2015). Second, the model also generates a prediction on differential targeting between network-based trade and door-to-door sales. I then discuss how both of these predictions are tested with the experimental data.

[11] I regress each of 8 household characteristics in Panel A of Table A1 on village-level treatment, an indicator for original recipients, and the interaction of these two variables. The F-statistics of these regressions range from 0.29 to 1.12 and thus the three variables do not jointly explain variation in any of the household characteristics.

3.1 Model Setup

The main benefit of adopting the new technology is improved flood tolerance. To formalize this, denote α_i as the probability that farmer i is affected by flooding. The agronomic return of the technology when flooding occurs is $r_i > 0$. Conversely, the return under non-flood conditions is zero — an assumption consistent with experimental results from farmer's fields (Dar et al., 2013). Therefore, the expected return of the technology is $R_i = \alpha_i r_i$.

In addition to the returns due to flood exposure, there is an idiosyncratic term, u_i, which measures benefits that are observed to the farmer, but not to the econometrician. For instance, some farmers are more risk averse or have stronger preferences for trying new technology. I assume that u_i is mean zero and independent of R_i. The difference in prices between the old and new technologies is v. I consider prices to be fixed. While there is a large literature on bargaining between buyers and sellers in networks (Corominas-Bosch, 2004; Manea, 2011; Abreu and Manea, 2011), I focus on network structure as the key potential barriers to adoption. I also show in the empirical analysis that differences in prices — the v term — are unlikely to explain the findings.

Under this simple setup the adoption probability in a perfect market where all demand is met is simply $P\{R_i + u_i > v\}$. While R_i can be approximated with data on past exposure to flooding, the u_i term is obviously unobservable. The purpose of the door-to-door sales is therefore to generate a revealed preference measure of demand that I can compare with network-based diffusion.

3.2 Diffusion in Networks

I assume that network-based diffusion imposes two sources of costs. The first term, $\underline{c}$, is the inconvenience of having to leave the house to obtain seeds. While likely important in many contexts, $\underline{c}$ seems less important in my sample due to the geographic setup of rural Indian villages. Half of the households in the sample have an original recipient that lives within 42 meters of their household. Also, over 90 percent of non-recipient households are located within 300 meters of an original recipient.

The second term, c_i, denotes trading costs due to the structure of the network. The value of c_i varies across the population because of varying degrees of connectedness between original recipients and non-recipients. As an example, a farmer may have a lower cost of adoption if a member of their own clan or kinship lineage was selected as an original recipient.

In contrast to the more real-world scenario of selling seeds at a shop, door-to-door sales eliminate both $\underline{c}$ and c_i by making transactions anonymous and bringing seeds directly to farmers. A standard market with some transportation costs would eliminate c_i, but not $\underline{c}$. I use various measures of connectedness to original recipients in the empirical analysis to show empirically that c_i is quantitatively important.

I assume that c_i, R_i, and u_i are distributed multivariate normally where the means of c_i and R_i are μ_c and μ_R. The parameter ρ denotes the correlation between R_i and c_i. The idiosyncratic term u_i has a mean of zero and is uncorrelated with both R_i and c_i.

The difference between revealed demand and adoption in network-based trading is simply $P\{R_i + u_i > v\} - P\{R_i + u_i > v + c_i + \underline{c}\}$. Comparing adoption probabilities between the two arms of the experiment allows me to estimate whether the frictions associated with network-based trade — the c_i and $\underline{c}$ terms — present a barrier to adoption. The first prediction that I therefore take to the data in Section 4.2 is whether these frictions exist and whether they drive any gap in adoption between revealed demand and diffusion in networks. Section 5 then considers the variation in c_i across the sample by estimating peer effects separately in network and door-to-door villages.

3.3 Differential Targeting

The expected return of adopters is a natural measure of targeting effectiveness that I can measure directly in the data. Conditional on the overall rate of adoption, the average return of adopters is a direct measure of how efficiently the technology is allocated. As shown in the appendix, the expected return of adopters in farmer-to-farmer diffusion is

$$E(R|R+u-c > v+\underline{c}) = \mu_R + \frac{\sigma_R(\sigma_R - \rho\sigma_c)}{\sqrt{\sigma_R^2 + \sigma_u^2 + \sigma_c^2 - 2\rho\sigma_R\sigma_c}} * M\left(\frac{v + \underline{c} - \mu_R + \mu_c}{\sqrt{\sigma_R^2 + \sigma_u^2 + \sigma_c^2 - 2\rho\sigma_R\sigma_c}}\right), \quad (1)$$

where $M(z) = \frac{\phi(z)}{1-\Phi(z)}$ is the inverse Mill's ratio. The expected return of adopters is above the average in the population if $-1 \leq \rho < \frac{\sigma_R}{\sigma_c}$. Intuitively, if returns and costs of adopting are negatively correlated, then the farmers facing the fewest barriers to adopting are those with the highest returns. Therefore, the adopters have higher average returns than the overall population. On the other hand, if the correlation between costs and returns is sufficiently large, then targeting in networks is less effective because the farmers with high costs are those with high returns.

The expected return of adopters when demand is revealed with door-to-door sales is

$$E(R|R+u > v) = \mu_R + \frac{\sigma_R^2}{\sqrt{\sigma_R^2 + \sigma_u^2}} * M\left(\frac{v - \mu_R}{\sqrt{\sigma_R^2 + \sigma_u^2}}\right). \quad (2)$$

Comparing equations (1) and (2), the average return of adopters will be larger in door-to-door sales if $\rho > \frac{\sigma_R}{\sigma_c}$. In this case, door-to-door sales crowd in farmers with high returns that did not adopt in networks. Conversely, if ρ is held constant, and μ_c is large relative to σ_c, then adoption in networks sends a stronger signal that returns are large because the farmer is willing to make the costly investment of adopting from other farmers.[12] Expected returns of adopters in networks will be larger in this case.

Overall, the model predicts that the difference in the average return of adopters between the two modes of exchange will depend upon the magnitudes of μ_c, σ_c^2, and ρ. However, the experimental design causes c and R to be uncorrelated when conditioning on network size. This results because

[12]To see this, note that $M(z)$ increases monotonically with z. Therefore, as μ_c increases and σ_c decreases, both M and $\frac{\sigma_R(\sigma_R - \rho\sigma_c)}{\sqrt{\sigma_R^2+\sigma_u^2+\sigma_c^2-2\rho\sigma_R\sigma_c}}$ increase.

random selection of original recipients generates random variation in connectedness conditional on network size. Applying this to the model, the targeting effectiveness of exchange in networks will depend only on the distribution of costs.

If frictions make network-based trade costly, then farmer-to-farmer diffusion could improve targeting. Any improvements in targeting would also offset any losses due to reduced adoption. The intuition is the same as in Alatas et al. (2016) who show that adding a small cost to claiming benefits in an Indonesian cash transfer program improves targeting towards the poorest households. I examine differential targeting in Section 4.3. In particular, I exploit the fact that agronomic returns can be estimated from past data on flooding, allowing me to estimate the difference between equations 1 and 2. This allows me to test whether network-based diffusion is any more (or less) effective at targeting the technology to the most flood-prone farmers.

4 Results

This section starts by quantifying the harvest of original recipients and documenting the exact uses of their output. I then argue that the two alternative uses of output — consuming it as rice or selling it for consumption — would make original recipients no better off in the absence of trading frictions. Sections 4.2 then shows the main results on adoption and Section 4.3 tests the model's auxiliary prediction on targeting. After having established the gap between adoption in networks and revealed demand at various prices, Section 4.4 considers diffusion over time and shows that adoption failed to catch up to revealed demand after two additional growing seasons. I then quantify the costs of this slow diffusion in Section 4.5. Finally, Section 4.6 shows that the productivity gains from adoption during the next season were measurable using high-resolution satellite imagery.

4.1 What did original recipients do with their harvest?

The average original recipient harvested just less than 350 kilograms of Swarna-Sub1. Figure 1 shows that this amount was allocated across four uses: transfers of seeds to other farmers, savings of seed for their own use, grain sold to either traders or government procurement centers, or grain consumed or stored for future consumption. The average farmer transferred only 6 kilograms — an amount roughly equivalent to what they were given — to other farmers. There are two notable features of these transfers. First, 80 percent of the transfers reported by original recipients were with farmers from the same village. Second, trading seeds for output of another variety is the dominant contract type. 69 percent of the reported transfers were exchanges of Swarna-Sub1 for another rice variety, while 26 percent were gifts and only 4 percent were sales for cash. Interestingly, the average original recipient saved over two times as much output for their own seeds. The 16 kilograms allocated for their own cultivation amounts to enough seed for approximately 30 percent of their land, suggesting that original recipients recognized the benefits of Swarna-Sub1.

The implicit price of these transactions (to the recipient farmers) is the opportunity cost of the unmilled rice that was traded for Swarna-Sub1. The government's minimum support price during this season was 12.5 Rupees per kilogram. As I note below, most farmers sell their output for lower prices to private traders — putting downward pressure on average prices. The implicit prices of the informal trades between farmers were therefore right in line with the prices in the door-to-door sales. Being able to observe the demand benchmark at 14 Rupees — a price higher than the implicit price of farmer-to-farmer trades — offers confirmation that results cannot be explained by lower prices in door-to-door sales.

Figure 1 also shows that the vast majority of the rice harvested was either sold for consumption (32 percent) or consumed directly by original recipients (62 percent). The entire experiment hinges upon these two alternative uses being no more profitable for original recipients relative to trading seeds. In other words, original recipients could not have obtained higher prices by selling for consumption or the consumption value (taste) of Swarna-Sub1 is no higher than the other varieties that farmers are willing to trade for Swarna-Sub1.

The first of these is directly testable in the data. I asked each original recipient that sold any Swarna-Sub1 as grain to report the transaction price. The average price across these transactions was 10.34 rupees per kilogram. This price is closest to the *lowest* price where seeds were offered during the door-to-door sales. In other words, original recipients were willing to sell their output at prices near or below the prices where demand was revealed in the door-to-door sales treatment. As a second piece of evidence, the seed varieties that farmers commonly traded for Swarna-Sub1 were no less valuable. I determined this by asking the name of the particular variety received for each trade reported by original recipients. I then calculated the average price across these varieties (weighted by share of trades) using data on sales prices collected in the same blocks as part of a related experiment (Emerick et al., 2016). The average market price of these varieties being received by original recipients was 10.15 rupees, a value quite similar to the price original recipients received when selling Swarna-Sub1. Thus, it seems unlikely that it was more profitable for original recipients to sell their output as grain for consumption.

The observed similarity in output prices also suggests that the consumption value of Swarna-Sub1 is similar to other varieties. We would expect clear price differences if Swarna-Sub1 has superior eating quality. Direct statements from farmers in Emerick et al. (2016) also support this conclusion. The percentage of adopters in that study reporting better taste as a reason for their adoption is nearly identical for Swarna, Swarna-Sub1, and other varieties.[13] An independent survey in two Indian states also found that amongst several statements about properties of Swarna-Sub1, adopters were the least likely to agree with the statement that it tastes better than other varieties (Yamano, Malabayabas, and Panda, 2013).

Original recipients did consume a large share of their harvest. However, this was not out of necessity. The median original recipient harvested 2,300 kilograms of rice during the year before the study — an amount more than enough to feed the household for a year. This result suggests that the output of Swarna-Sub1 represents a modest share of overall production. It therefore seems unlikely

[13]See Table A1 in the online appendix of Emerick et al. (2016).

that a subsistence constraint forced original recipients to consume their Swarna-Sub1 output.

Combining these descriptive data, there seems to be no evidence suggesting that original recipients were made worse off when transferring seeds to other villagers. The profitability of alternative uses — either selling as rice or consuming directly — can't explain limited diffusion of seeds to other farmers. The rest of the paper seeks to understand whether the observed diffusion in networks is a first-best outcome and if not, attempts to quantify the economic losses from slow diffusion between farmers.

4.2 What was the actual demand for seeds by non-recipients?

The door-to-door sales experiment serves as a benchmark so that I can compare revealed demand with diffusion directly between farmers. Again, I can make this comparison at three different price levels: 10, 12, and 14 rupees per kilogram. The experiment therefore reveals demand at a price equivalent to the one received when original recipients sold output as rice for consumption. In addition, I am able to estimate whether farmers were willing to pay prices higher than those received when original recipients chose to sell output as grain. I use the data from the follow-up survey with all non-recipients, which was carried out after planting for the second year of the study.

The revealed demand at all three price levels is over three times the adoption that took place when farmers shared informally amongst themselves. Figure 2 demonstrates this as simply as possible by showing the raw adoption rates in the network and door-to-door villages. Only 6.5 percent of non-recipient farmers adopted after the first year in network villages. At the same time, the above data showed that original recipients in the same villages were selling their output for an average of 10.34 rupees per kilogram. The striking finding is that 43.8 percent of farmers were willing to buy seeds at a price of 10 rupees and this falls to 41.7 percent and 35 percent at prices of 12 and 14 rupees, respectively. Therefore, this gap in adoption can not be explained by price differences. Demand for seeds existed at a price of 14 Rs, which is 35 percent above the price that original recipients chose to sell the very same variety as rice for consumption. These findings show that the costs of trading in networks do drive a wedge between actual demand for a new technology and the amount of trade that takes place between farmers. Returning to the model in Section 3.2, the costs of exchange c_i appear to be large enough to prevent adoption by around a third of farmers that expressed demand for the technology in the door-to-door sales experiment.

I next show the regression results that correspond to Figure 2. Formally, the regression specification is

$$adoption_{ij} = \beta_0 + \beta_1 Price\ 10_j + \beta_2 Price\ 12_j + \beta_3 Price\ 14_j + \varepsilon_{ij}, \tag{3}$$

where $adoption_{ij}$ is an indicator for adoption by farmer i in village j, and β_1, β_2, and β_3 measure the gaps in adoption at the three different price levels.[14] Recall that exchanging seeds was the most

[14] I focus on a binary adoption rate in the analysis because the amount used is only relevant for a single year. After one year, the harvest produced from only 1-2 kg of seed is enough to cultivate the average farmer's entire landholdings. In door-to-door villages, the adoption indicator is set to 1 if *either* the farmer purchased from an NGO representative, or adopted from a peer.

frequent arrangement in direct transactions between farmers. The opportunity cost to the seller of such a transaction is the 10.34 rupees that could have been obtained by instead selling that same output to traders for consumption. The coefficient β_1 therefore represents the difference between adoption in farmer-to-farmer networks and revealed demand at comparable prices.

The estimates in column 1 of Table 2 show that the revealed demand at all three price levels was significantly higher than the network adoption rate. Column 2 shows my main estimate which is the average effect across all three price levels. In particular, adoption is higher in door-to-door villages by 33 percentage points. The rate of adoption of 40 percent in door-to-door villages is larger than the adoption in the network by over five times. Not surprisingly, column 3 shows that the estimated adoption gap changes little when introducing household controls.

Are there some groups that are better off when trading occurs between farmers, or is the adoption gap similar across all farmers? I show in Table A2 that the treatment effect of door-to-door sales is smaller for lower caste farmers and farmers that are below the poverty line. It is therefore poorer households that lose the least from network-based trading. Liquidity constraints offer one possible explanation if poor households can't afford to purchase new seeds right before planting.

4.3 Are networks better at targeting to farmers with higher expected returns?

The theory model pointed to improved targeting as a possible outcome of trading in networks. The intuition is straightforward. If diffusion between farmers is costly then only farmers with higher returns will bear these costs and the average return of adopters increases relative to the door-to-door sales environment where adoption is costless. As a first step in testing this prediction, I use data on flooding during the past five years to generate a measure of expected returns for each farmer in the sample,

$$return_{ij} = \frac{\frac{1}{5} * \sum_{p=1}^{P_{ij}} \sum_{t=2008}^{2012} R(d_{ijpt}) * area_{ijp}}{\sum_{p=1}^{P_{ij}} area_{ijp}}. \tag{4}$$

The term d_{ijpt} represents the duration of flooding for farmer i in village j on plot p during year t, P_{ij} is the total number of plots cultivated, and the function $R(\cdot)$ is the expected agronomic return of Swarna-Sub1, relative to Swarna. The units of measurement of R are kilograms per hectare cultivated. I use estimates of R that were generated using data from a randomized experiment carried out in nearby villages during 2011. Specifically, I use nonparametric estimates of the treatment effect of Swarna-Sub1 as a function of flood duration.[15] Figure A1 shows that this metric of estimated returns varies both across and within villages.[16]

[15] See Dar et al. (2013) for the actual nonparametric estimates.

[16] One caveat is that this approach measures *agronomic* returns rather than *economic* returns. Emerick et al. (2016) show that access to Swarna-Sub1 causes farmers to change several production practices, leading to increases in yield even during years when flooding does not occur. Increases in investment are generally larger for farmers that have

I consider the average return of adopters as the most direct measure of targeting effectiveness. Figure 3 displays the densities of estimated returns for adopters across the different treatment groups. Visually, the distribution of estimated returns in the network villages shifts to the right when compared to door-to-door villages.

OLS regression estimates also suggest that farmer networks were slightly more effective at targeting. The regression results in column 1 of Table 3 show that the average return of adopters in door-to-door sales is smaller by 38 kg per hectare, an approximate 22 percent decrease.[17] The effect is reasonably large, but not quite statistically significant (p=0.12). Column 2 shows a similar, but slightly more precise, result when using log returns as the outcome variable. Turning to column 3, similar results are obtained when using a self-reported measure of flood risk for the plot where the new variety would be planted. Farmers were asked to assess on a scale from 1-10 how prone their Swarna-Sub1 plot is to flooding. The average value amongst adopters in the network villages is 5.25. The predicted decrease with door-to-door sales is 0.81, or 15 percent. The estimated effect with this separate measure is qualitatively similar, but also not statistically significant (p=0.12). I show in Table A3 that the estimated returns in networks are slightly affected by two farmers that were provided free Swarna-Sub1 seeds from a local disaster management office. Dropping these farmers from the analysis makes the targeting differential slightly more precise.

The results are in line with the model's prediction that costs of adopting in networks affects targeting effectiveness. However, the targeting effect is modest. It will therefore offset only a small amount of the losses from the large gap in adoption. This arises because the relative magnitudes of the adoption gap and targeting effect are strikingly different. The door-to-door sales increased adoption by over five times, but only reduced targeting effectiveness by around 20 percent. Therefore, the targeting differential is nowhere near large enough to offset the adoption gap.

4.4 Did the adoption gap persist over time?

There are two immediate reasons why these one-year impacts are complemented by a longer run analysis. First, it may simply take additional time for farmers to learn about the new technology, to learn who is cultivating it, and to approach these farmers to obtain the seed. Second, the randomly selected original recipients are representative of the average farmer and not necessarily the most entrepreneurial farmer. Allowing additional growing seasons to pass would allow entrepreneurial farmers to select into adoption, multiply the seeds, and sell them widely after one year. If either of these phenomena explain the short-term adoption gap, then we should expect this gap to disappear over time.

The longer run follow-up survey allows for this idea to be tested. The survey was carried out approximately two years after the first follow-up survey. It involved a door-to-door varietal adoption census with all households in each of the 82 villages. I can therefore measure adoption

more farmers in their peer group also cultivating the variety. Since networks favor adoption by peers, one advantage of farmer-to-farmer exchange is that it could facilitate these behavioral changes.

[17] Strata fixed effects are dropped in this regression in order to avoid absorbing selection effects.

for two additional years after the first survey. Attrition is not a concern. Around 99 percent of the households that were surveyed as part of the original survey were reached again during the second follow-up survey. I focus the analysis on the original sample of 15 farmers per village.

Figure 4 shows the differences in raw adoption rates over time.[18] There are two important observations from the figure. First, there has been some convergence. Despite this, there still remains a gap between revealed demand and adoption in farmer seed networks. While adoption increased to 32 percent of farmers in network villages, approximately 50 percent of farmers were cultivating the technology in door-to-door villages. This translates to a 56 percent increase. This persistence of an effect over time shows that the impediments to trading between farmers persist over at least a three-year period.

Both this long-run survey and the one-year followup measured adoption in the same way across all villages. Each household was directly asked about their adoption of Swarna-Sub1. While adoption was directly observable from the door-to-door sales, the survey was still necessary for two reasons. First, farmers in door-to-door villages could still obtain seeds from original recipients. The survey was needed to detect adoption of this type. Second, using the survey in all villages ensures that adoption is being measured in the same way across the two treatments.

4.5 How costly is slow diffusion in networks?

The results thus far establish a gap between revealed demand and take up when diffusion occurs informally between farmers. I next take this one step further to quantify the losses that result from farmer-to-farmer diffusion in networks. This additional exercise requires only an estimate of the causal impact of Swarna-Sub1. The experiment in Emerick et al. (2016) generated such an estimate. Specifically, Swarna-Sub1 reduces risk farmers face and this induces several changes in agricultural management. These changes include applying more fertilizer and investing more at planting. Overall, the estimated impact of Swarna-Sub1 due to these crowd-in effects is around 283 kilograms per hectare cultivated. Put differently, the average Swarna-Sub1 adopter will gain around 283 kilograms of output for each hectare of rice they cultivate.[19]

I use this value along with the estimated returns from flood tolerance (calculated above) to calculate a per-farmer measure of the revenue gains from Swarna-Sub1. I convert this measure from kilograms to rupees by multiplying by the Indian government's minimum support price for rice during each year. More specifically, the measure of estimated revenue gains is

$$\Delta revenue_i = \sum_{t=2013}^{2015} \delta^{t-2013} * price_t * adopt_{it} * (return_i * ss1_{it} + farmsize_i * 283.45), \qquad (5)$$

where $adopt_{it}$ is an indicator for adoption of Swarna-Sub1, $return_i$ is the return (in kg) from flood tolerance estimated in Equation 4, $ss1_{it}$ is the area planted with Swarna-Sub1, and $farmsize_i$ is

[18]The complete regression results including strata fixed effects are given in Table A4.

[19]This value is output per hectare cultivated with rice, not output per hectare cultivated with Swarna-Sub1.

the rice area of farmer i.[20] I use a discount factor of $\delta = 0.95$. I report both the difference in annual revenue gains for each of the three years and the cumulative effect as written in Equation 5.

The estimates demonstrate the costliness of slow diffusion in networks. This is first shown graphically in Figure 5 while the complete regression results are shown in Table A5. My best estimate indicates that relative to the demand benchmark, network-based diffusion led to a cost of around 21 dollars per farmer in 2013, 18 dollars in 2014, and 12 dollars in 2015. This amounts to a non-trivial magnitude of 51 dollars cumulatively over the entire three-year period. The amount is equivalent to around 20 days of casual work or around 8 percent of the annual rice harvest for an average farmer.[21]

The door-to-door sales were not meant to test a scalable intervention. Nonetheless, it is useful to benchmark these estimated revenue losses against the acutal costs of door-to-door sales. The cost of each door-to-door visit from my NGO partner was approximately 9 dollars. The foregone revenue gains from missed trading opportunities are larger than the cost of door-to-door sales by a factor of five. This finding is striking and suggests that the network-based diffusion is slow diffusion enough to make even a door-to-door sales intervention cost effective.

4.6 Did adopters benefit during the next season's flood?

I next show direct measures of benefits of adoption during the next season. The 2013 season (immediately following the door-to-door sales) involved variation in flooding across the areas of my sample. I overcome the lack of survey-based data on crop yield by using high resolution satellite imagery that allows me to observe the greenness of fields on various days starting in April of 2013 (before planting) and continuing until March 2014 during the dry season when only irrigated fields are cultivated.[22] These data provide estimates of the Normalized Difference Vegetation Index (NDVI) — a standard measure of plant greenness that is strongly correlated with survey-based measures of rice yield (Huang et al., 2013). In addition, the spatial resolution of 30 meters allows me to generate a plot-specific measure of greenness by overlaying each image with the centroids of the farmers' plots.[23]

I combine three sources of variation to measure the impact of Swarna-Sub1 on the NDVI. First, there is variation in adoption: 59 percent of farmers for which I have plot locations did not adopt during the 2013 season. Second, there is differential exposure to flooding. Satellite images of flooded areas identify 35 percent of plots that were flooded during the season.[24] Third and finally,

[20] I observe 2015 rice area and use it to estimate rice area for the other two years.

[21] The casual labor benchmark is based on the 2015-2016 Odisha wage for NREGS, the government's employment guarantee scheme. The rice output benchmark is based on the measured average output in the sample of 3.3 tons.

[22] The data are derived from the Landsat 8 8-day NDVI composites made available on Google's Earth Engine public data catalog (code.earthengine.google.com, accessed September 2016).

[23] Survey teams collected the centroids of plots that were cultivated with Swarna-Sub1 by original recipients during the 2012 season. This differed for non-recipients where survey teams collected locations for the most flood-prone plot. This implies some measurement error in identifying the plots that were actually cultivated with Swarna-Sub1 during the 2013 season when the satellite images become available. This would only affect the analysis in the unlikely event that this source of measurement error is correlated with flooding.

[24] I define flooded plots using daily images of flood areas (see http://csdms.colorado.edu/pub/flood_observatory/

I observe images during both the wet and dry seasons and flooding is non-existent during the dry season. This third source of variation therefore allows for a weaker identification assumption. I basically compare the difference between adopters and non-adopters in flooded versus non-flooded areas in both the wet and dry seasons. This triple differences estimate eliminates concerns that the differential return to adoption in flooded areas is due to adopters in flooded areas putting the technology on different types of land — such as irrigated fields — since these fields would appear more productive during the dry season.

Figure 6 helps visualize these data by showing the NDVI measures at four points in time: once before the 2013 wet season, twice during the season, and once after the completion of the wet season. For each of these images I basically compare the difference in greenness between adopters and non-adopters on flooded plots (the blue dots) to that on plots that were not flooded (the black dots). The flood-tolerance benefit of the technology is apparent when making this comparison. Figure 7 shows this by plotting the difference in log NDVI between adopters and non-adopters over time and separately for plots that were affected by flooding. Visually, there is little difference between adopters and non-adopters in areas that were not flooded (the black line in the figure). In contrast, there is a noticeable increase in greenness for adopters on fields that were flooded and this increase exists only during the growing season. It can therefore not be attributed to fixed differences of the land such as irrigation access, since these characteristics would be visible during the dry season.

I estimate the following triple differences regression in order to quantify this effect:

$$\begin{aligned} log(NDVI_{it}) = \beta_0 &+ \beta_1 flood_i * growseason_t * adopt_i + \beta_2 adopt_i + \beta_3 growseason_t + \\ &\beta_4 flood_i + \beta_5 growseason_t * flood_i + \beta_6 growseason_t * adopt_i + \\ &\beta_7 adopt_i * flood_i + \varepsilon_{it}. \end{aligned} \tag{6}$$

In this specification $NDVI_{it}$ is the the observed NDVI for farmer i on date t, $flood_i$ is a time-invariant indicator for farmers that were affected by flooding during the wet growing season, $adopt_i$ is an indicator for adoption, and $growseason_t$ is an indicator for observations during the wet growing season, i.e. the 7 images from August 13th to November 17th. The coefficient of interest β_1 measures the differential effect of adoption on plant greenness in areas that were flooded and for observations during the growing season. This triple differences estimate can be interpreted causally as long as there are no unobserved differences between adopters and non-adopters that both influence crop productivity and exist *only* in areas that were affected by flooding and during the rainy growing season.

Table 4 shows these results for two different measures of exposure to flooding. In column 1 flood exposure is defined as any field within 250 meters of any flooded area since this corresponds with the resolution of the flooding data. Using this definition I find a 10.1 percent increase in greenness

MODISlance/080e030n/, accessed September 2016). The resolution of these images is 250 meters. I therefore consider a plot to be flooded if it is within 250 meters of the nearest area that was flooded any time during the growing season.

associated with adoption of Swarna-Sub1. The point estimate of β_5 shows that the average flooded field is 31 percent less green during the growing season and therefore the flooding imagery capture variation in flooding that is strongly associated with crop health. Access to Swarna-Sub1 eliminates around one third of this negative effect of flooding on greenness. Column 2 demonstrates that the results are similar when alternatively defining flooded fields as those that are within 500 meters of any flooded area. . I show in the appendix that similar results are obtained when controlling for time-invariant unobservables with either village or farmer fixed effects (Table A6).

This additional analysis verifies that the adoption induced by the experiment was immediately beneficial to farmers. The above cost-benefit calculations relied on impact estimates from a randomized experiment in nearby villages during the previous two seasons. While the impact estimates are not directly comparable because the outcomes are different, i.e. survey-based measures of crop yield vs. satellite-based measures of NDVI, these additional findings strengthen the argument that more rapid diffusion of seeds between farmers in networks would have led to immediate gains in crop yield.

5 Mechanisms

I have shown that relative to demand revealed in a door-to-door sales experiment, diffusion of seeds between farmers is slow. This slow diffusion occurs with a profitable technology. Combining these two, the costs of slow diffusion over a 3-year period are approximately equivalent to offering 20 days of casual work to the average farmer in the sample. These results suggest there is some friction that prevents trading between farmers. In this section I look at several possible frictions that would be compatible with the results. I start by showing several pieces of evidence consistent with network structure and the propensity to trade with close peers as being serious impediments leading to slow diffusion. I then consider four plausible alternative explanations of my findings: capacity constraints, the random choice of recipients, differences in seed quality, and increased salience caused by door-to-door sales. Overall, I don't find strong evidence in support of these four alternative explanations. Nonetheless, the evidence is merely suggestive and I can not strongly eliminate any of the other possible explanations.

5.1 Favoritism

The random selection of original recipients allows me to test whether peer effects partially explain the slow diffusion between farmers. If they do then being part of the same social group as an original recipient should have a positive effect on adoption and this effect should be reduced when actual demand is revealed with door-to-door sales. This is tested by estimating

$$\begin{aligned} adoption_{ij} = &\beta_0 + \beta_1 door\ to\ door_j + \beta_2 linksOR_{ij} + \beta_3 links_{ij} + \beta_4 linksOR_{ij} * door\ to\ door_j \\ &+ \beta_5 links_{ij} * door\ to\ door_j + \varepsilon_{ij}, \end{aligned} \tag{7}$$

where $linksOR_{ij}$ is the number of peers of farmer i that were selected as original recipients and $links_{ij}$ is the total number of links of farmer i. I focus on two closely related measures of social links: sharing a common surname and belonging to the same sub-caste, or jati.[25] Surnames offer a slightly finer measure than sub-caste. While members of the same sub-caste may have several surnames, sharing a common surname measures clan association or kinship. Importantly for identification of β_2 and β_4, the random selection of original recipients guarantees that the number of links with original recipients is random when conditioning on the total number of links, thus avoiding the classic reflection problem discussed in Manski (1993). A finding of $\beta_2 > 0$ and $\beta_4 < 0$ would indicate that network structure contributed to the slow diffusion between farmers.

Network structure does influence diffusion. In column 1 of Table 5, sharing the same surname with an additional original recipient causes a 3.5 percentage point, or 50 percent, increase in the probability of adoption in network villages. The peer effect decrease significantly by 7.5 percentage points when door-to-door sales were carried out. Turning to column 2, the effects increase when including village fixed effects.[26] Having the same surname as one additional original recipient results in a 106 percent increase in adoption with network-based diffusion. Again, this peer effect is eliminated with the door-to-door sales.

Belong to the same sub-caste as an original recipient also impacts adoption in the network villages. In column 3, having one additional original recipient from the same sub-caste leads to a 4 percentage point increase in the probability of adoption, representing a 57 percent effect. The estimated coefficient on the interaction between the door-to-door indicator and the number of original recipients in the same sub-caste is negative and of similar magnitude as the effect in the network villages. Thus, the sub-caste peer effect becomes effectively zero when door-to-door sales are made. Column 4 shows qualitatively similar results when estimating the regression with village fixed effects.

These estimated network effects are robust to two alternative estimation strategies. First, accounting for the dichotomous nature of the dependent variable by using a probit specification has little impact on the estimates (columns 3 and 4 in Table A7). Second, an alternative approach would be to use the *share* of original recipients with the same surname or subcaste that were selected as recipients. As shown in Table A8, using this approach actually improves precision of the estimates.

Table 6 uses the additional followup to show that surname connections guide adoption even over a three-year period. The results in column 1 are reassuringly similar to the estimates from the initial follow-up survey.[27] That is, each additional surname connection with an original recipient causes

[25] There is substantial variation in surnames within villages. The average number of unique surnames per village is 5.6. Therefore, each farmer in the sample shares a surname with approximately 3.3 other farmers in the sample.

[26] This likely occurs because the villages with little variation in adoption and where most farmers share the same surname receive less weight in the identification. In Table A7 I show that the estimated peer effects are much larger in the sample of villages where there was at least one adopter (columns 1 and 2). The results are also more similar to fixed effects results when discarding the 5 percent of observations where over 15 of the farmers in the village have the same surname (not shown).

[27] The 2013 adoption rate measured in the 3 year follow-up is larger in both types of villages by around 10 percentage points. One reason for the difference between the two surveys is that the original follow-up survey in 2013 was carried

a 5.4 percentage point increase in adoption in network villages. This peer effect is erased entirely in door-to-door villages. In fact, the peer effects in door-to-door villages are if anything negative in 2013 and 2014 and statistically insignificant in 2015.[28] These additional findings point to the underlying structure of the social network as an actual friction that limits diffusion. The door-to-door sales experiment revealed demand by farmers that were unable to access the technology through social networks. The farmers with this inability were those lacking connections to recipients and facing seemingly high costs of making those connections.

Interestingly, sub-caste association is a much weaker predictor of long-run adoption in the village economy. Table A9 shows that the number of recipients in the same sub-caste is positively, but weakly related to adoption. One explanation — which is consistent with results in the online appendix — is that surnames are a stronger measure of social connections.

How much of the revenue losses from slow diffusion can be explained by rigidity in network structure? Table 7 shows regressions where the outcome variable is predicted gains from the technology, as in Figure 5. I compare estimated gains from network-based diffusion for two farmers, one that is less connected and one that is more connected to original recipients. I measure connectedness using residuals from the regression of number of original recipients with the same surname on the total number in the sample with the same surname. More precisely, the farmer at the 5th percentile of the residual distribution shares the same surname with five other people in the sample, none of which were selected as original recipients. The farmer at the 95th percentile shares the same surname with six other villagers and three of these were recipients. Focusing on the marginal effects from column 4, the relatively unconnected farmer is expected to lose 83 dollars in revenue when diffusion occurs via farmer networks. The farmer at the 95th percentile is on the other hand expected to lose only 16 dollars from network-based diffusion, relative to actual demand. This compares to the median farmer that loses about 49 dollars from network-based diffusion. This pattern indicates that variation across the sample in connectivity explains some — but not all — of the losses that arise from slow diffusion between farmers.

As a final test, I consider whether farmers adjusted their stated networks by making links with original recipients. I find no evidence of such behavior. The follow-up visits included a module where each respondent was asked whether they would go to each other respondent for sharing seeds or other inputs. I use these dyadic data to test whether links are more likely in dyads with a single original recipient. Table A10 shows that dyads with a single original recipient are about 6 percent more likely to involve a link, however this difference is not statistically significant. In contrast, original recipients are more likely to report links amongst themselves. This effect could be generated by the desire of recipients to seek information and other inputs from other farmers that are known to be using a new technology. The dyadic regressions also show significant homophily

out after farmers had planted seedbeds but before the main rice fields had been planted. Thus, there was an additional opportunity for farmers to obtain seedlings (as opposed to seeds) from other farmers. Nonetheless, Table A4 shows that the estimated adoption gap between treatment and control villages is similar between the two surveys.

[28]The p-values for the overall effect of surname connections in door-to-door villages are 0.054, 0.029, and 0.48 for 2013, 2014, and 2015 respectively.

in these village networks. Farmers sharing the same surnames, being part of the same sub-caste, or living close to each other are more likely to report being linked (Table A11). Of these variables, sharing a common surname is the most robust predictor of link formation.

5.2 Alternative Explanations

Capacity constraints

Capacity offers an explanation for the results if the output of original recipients was insufficient to meet seed demands of other villagers. I consider this by asking whether the output of original recipients would have been enough to meet the village's seed demand if it had all been allocated to its most efficient use as seeds. The data suggest that this is the case. The average harvest across all original recipients during year one was 1,700 kilograms, indicating that each original recipient harvested an average of 340 kilograms. The average farmer in the door-to-door sales experiment procured 2.83 kilograms.[29] Taken together, the amount produced was sufficient to meet the seed demands of 600 farmers — a number larger than village size by a factor of 6.

Figure A2 shows this across door-to-door villages by showing the distribution of the differences between the total harvest and the total amount of Swarna-Sub1 planted in each village *after* door-to-door sales were made. The average amount harvested exceeded the amount planted during the next year by an average 14 times. There was only one village where the harvest was particularly poor and the amount purchased exceeded the amount harvested. These straightforward calculations are inconsistent with capacity-based explanations of the findings.

Selection of original recipients

The original recipients were selected randomly to allow for causal identification of network effects. But, such random selection it is not a policy-relevant method since it would never be used in practice. An alternative approach would be to purposefully select the original recipients that are theoretically desirable for diffusion. Recent studies have considered the impacts of using network theory to target injection points (Beaman et al., 2015) as well as directly surveying villagers about optimal injection points (Banerjee et al., 2014).

I exploit the random selection of original recipients to test whether adoption was higher when recipients were more important in a network sense. I partition villages into two groups according to the ratio of the average degree of recipients to non-recipients. Villages where recipients are more central are defined as those where this ratio is greater than the sample median.[30] One important caveat is that the number of links is not the optimal measure of how effective a particular node in a network is as an injection point (Beaman et al., 2015). Nonetheless, the analysis helps to answer

[29] This figure is unconditional, i.e. it includes farmers choosing not to purchase.

[30] Randomization of village-level treatment was stratified by this degree ratio for purposes of investigating heterogeneity with respect to importance of original recipients. Using the ratio of average degrees carries one additional advantage since the social network in each village was only partially sampled. Chandrasekhar and Lewis (2011) show that the bias in average degree due to partial sampling of network data is proportional to the sampling rate. Using the ratio of average degrees should therefore minimize concerns regarding biases.

the question of whether a more targeted approach of selecting the most central original recipients explain the low adoption in the network-based model of diffusion.

I find no evidence that the gap in adoption is any smaller when original recipients were randomly more connected to other villagers (Table A12). If anything, the effect goes in the other direction. The effect of the door-to-door treatment is 25.6 percentage points in villages where original recipients were less important and 41 percentage points in villages where they were more important. This difference is also statistically significant at the 10 percent level. A simple explanation of this finding is that more connected original recipients are indeed better at generating additional demand. However, they are no better at actually sharing seeds with other villagers.

The random selection of original recipients also allows me to test whether adoption was faster when the recipients differed on observable characteristics from the non-recipients. Table A13 shows evidence that the relative wealth of original recipients influenced adoption. The adoption rate increased by 7.7 percentage points in villages where recipients were wealthier than non-recipients. This effect increases by 7 percentage points in door-to-door villages, yet the difference is not statistically significant. The finding suggests that targeting entry points based on wealth can increase diffusion. There are two candidate explanations. First, wealthy farmers could be less likely to store or consume their harvest and thus serve as better recipients for transmitting seeds to other farmers. Second, wealthy farmers could be better at demonstrating the technology and thus creating demand. The data seem more consistent with the latter explanation since the relative wealth of original recipients influences both adoption in networks and demand revealed in the door-to-door sales arm.

Information on returns

Was diffusion slow because farmers lacked information on the key benefits of the technology? All farmers were instructed on the benefits of Swarna-Sub1 during the initial village meeting. In addition, farmers were reminded of this information when surveyed in between the harvest of year one and the door-to-door sales of year two. The demand revealed in the door-to-door sales suggests that the information was at least partially internalized and used when making purchasing decisions.

I also use a non-experimental test to suggest that information on returns was not a key constraint. Section 4.6 demonstrated variation across the sample in flooding exposure during the 2013 season. In addition, the returns to adopters from Swarna-Sub1 were higher in villages exposed to flooding. Adoption during the 2014 and 2015 seasons should therefore increase in flooded villages if this event provided new information to farmers. I estimate separate regressions for each season where adoption is regressed on an indicator for being in a flooded village and strata fixed effects.[31] Figure A3 shows that being in a flood-affected village is unrelated to adoption in 2014 and 2015. This suggests that lack of information about the flood tolerance property did not constrain adoption decisions.

[31] Flooded villages are those for which at least one plot was flooded, where the definition of flooding corresponds to that in Section 4.6.

Quality

Diffusion between farmers would be less effective if the seeds traded between farmers are lower quality relative to the seeds sold in the door-to-door sales. There are two aspects of quality to consider. First, information asymmetries could drive the results if Swarna-Sub1 is indistinguishable from other seeds. If this were true then diffusion occurring only between socially proximate farmers would be a solution to the information asymmetries that arise from not being able to tell the two types of seeds apart. Second, the adoption gap could result from failure of original recipients to take basic measures to ensure quality seed production, such as removing seeds of "off types" or varieties from the neighboring field.

A unique property of Swarna-Sub1 makes the first candidate explanation unlikely. Swarna-Sub1 has a white husk, making it easily distinguishable from the reddish husk of Swarna. Figure A4 shows the visual comparison between the two varieties. The white husk color of Swarna-Sub1 makes it easily distinguishable and thus makes it unlikely that lack of trust and counterfeiting explain the results.

Unobserved seed quality still remains as a possible explanation. The seeds that were exchanged between farmers were second generation, i.e. output from the first year's harvest, while the seeds sold in door-to-door sales were procured directly from a private seed company. If farmers fail to produce quality seeds, this could potentially explain low adoption between themselves.[32]

I test whether the effect of door-to-door sales varies as a function of stated and revealed preference measures for new and certified seeds. I use two measures of preferences. First, approximately 42 percent of farmers purchased certified seeds from local government offices for the 2012 season.[33] Given the higher quality standards for certified seeds, this serves as a revealed preference measure of demand for seed quality. As a second measure, I use responses to a survey question where demand was elicited hypothetically for two scenarios: one where certified Swarna-Sub1 seeds were available and another where seeds were produced by another villager. I define those who indicated that a larger quantity of certified seeds would be procured as having a preference for new seeds. This group represents approximately half of the sample.[34]

I find no evidence that the door-to-door sales effect varies with these two measures. Table A14 shows that the correlation between the two measures of preference for new seeds and adoption in the village economy is small and statistically insignificant. Further, the effect of door-to-door sales is no larger for this group of farmers. While inconsistent with the unobserved quality explanation, these findings are only suggestive since the two variables are imperfect measures of actual preferences.

[32]As an example, if seed is stored without proper drying, then germination ability and vigor of seedlings are negatively affected. Other practices that farmers can do to improve seed quality and purity are hand sorting to remove weeds and seeds of other varieties, winnowing to remove empty grains and chaff, and careful storage to avoid moisture absorption and damage by pests.

[33]Seeds that are certified are produced following certain guidelines that ensure purity and higher quality.

[34]These two measures are not strongly correlated. A regression of one characteristic on the other produces a point estimate that is small and statistically insignificant.

Independent effect of door-to-door sales on demand

Simply going door-to-door to sell seeds could have increased awareness about the technology or sent a signal to farmers about its potential value. If this is true then the door-to-door sales would have created demand, rather than measured it as a benchmark to compare with network-based diffusion. Two steps were taken to minimize this effect. First, the original recipients were selected at a village meeting to make all farmers aware of their identities. Second, farmers were reminded about the technology and its flood-tolerance property during the midline survey that occurred three months before the door-to-door sales.

To test salience effects, I take advantage of the fact that while door-to-door visits were only made to a randomly selected group of 15 farmers per village, it was well known that NGO staff were moving between houses to offer seeds. Houses in the sample villages are small and located in close proximity. For instance, there is an average of over two other houses in the sample within a 25 meter radius of each sample household. If the door-to-door visit itself increased perceived benefits of the technology, then adoption rates of farmers outside the sample should be larger in door-to-door villages. In addition, original recipients would be more likely to continually adopt over time in door-to-door villages if farmers perceived benefits of the technology to be higher due to door-to-door sales.

There is no evidence of salience effects in the data. I use data from the final survey with original recipients to test whether those in door-to-door villages transacted with a larger number of farmers from outside the sample. Table A15 shows that the effect of door-to-door sales on the number of trading partners from outside the sample is negative and statistically insignificant. Also, Figure A5 shows that the door-to-door sales had no effect on the long-run adoption of original seed recipients. villagers that did not receive door-to-door sales visits, or farmers that had adopted at the time of the original followup survey. In combination, these results suggest that the results are not due to demand effects created by the door-to-door sales themselves.

6 Conclusions

I have shown that there is a gap between demand and actual adoption in the rural Indian setting where a new seed variety diffused between farmers. The revealed demand for an improved seed variety was over five times larger than the adoption rate in a network setting where farmers traded amongst themselves. Moreover, this limited diffusion is costly. My best estimate is that the average farmer paid a price equivalent to around 20 days of casual labor by having to rely on fellow villagers as a source of seeds. The data suggest that the seeds were misallocated. Much of the output was sold to traders at prices below the prices that other villagers were willing to pay.

Additional results showed that caste-based favoritism is an important friction driving this misallocation. Farmers that shared the same surname or belonged to the same sub-caste with the randomly-selected original recipients were significantly more likely to adopt — despite their demand being no higher when demand was revealed with door-to-door sales. Put differently, the gap between adoption and demand was largest for farmers that lacked social connections to the original recipients.

The findings suggest that favoritism based on social proximity interferes with the functioning of markets in settings where caste and ethnic identity remain important. Focusing on India, the prevalence of within-caste transactions is prevalent in both credit markets and village markets for informal insurance. The ability to transact with close peers can increase efficiency in these cases because monitoring and enforcement is better within social groups. While I observe that seed transfers are more likely to occur within social groups, this is arguably a transaction where moral hazard is less perverse. Therefore, the propensity to trade seeds with socially proximate farmers is more likely to reflect favoritism rather than a need to resolve information asymmetries. The paper has shown that the costs of this favoritism — in terms of slow technological diffusion — are large.

References

Abreu, D. and M. Manea. 2011. "Bargaining and efficiency in networks." *Journal of Economic Theory* .

Alatas, Vivi, Ririn Purnamasari, Matthew Wai-Poi, Abhijit Banerjee, Benjamin A Olken, and Rema Hanna. 2016. "Self-Targeting: Evidence from a Field Experiment in Indonesia." *Journal of Political Economy* 124 (2):371–427.

Anderson, S. 2011. "Caste as an Impediment to Trade." *American Economic Journal: Applied Economics* 3 (1):239–263.

Bandiera, Oriana, Iwan Barankay, and Imran Rasul. 2009. "Social connections and incentives in the workplace: Evidence from personnel data." *Econometrica* 77 (4):1047–1094.

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

———. 2014. "Gossip: Identifying central individuals in a social network." Tech. rep., National Bureau of Economic Research.

Banerjee, Abhijit, Esther Duflo, Maitreesh Ghatak, and Jeanne Lafortune. 2013b. "Marry for What? Caste and Mate Selection in Modern India." *American Economic Journal: Microeconomics* 5 (2):33–72.

Bardhan, P. and D. Mookherjee. 2011. "Subsidized Farm Input Programs and Agricultural Performance: A Farm-Level Analysis of West Bengal's Green Revolution, 1982-1995." *American Economic Journal: Applied Economics* 3 (4):186–214.

Beaman, Lori, Ariel BenYishay, A. Mushfiq Mobarak, and Jeremy Magruder. 2015. "Can Network Theory-based Targeting Increase Technology Adoption?" Unpublished.

Bellon, Mauricio R, David Hodson, and Jon Hellin. 2011. "Assessing the vulnerability of traditional maize seed systems in Mexico to climate change." *Proceedings of the National Academy of Sciences* 108 (33):13432–13437.

BenYishay, Ariel and A Mushfiq Mobarak. 2015. "Social Learning and Incentives for Experimentation and Communication." Tech. rep., National Bureau of Economic Research.

Cai, Jing, A. de Janvry, and E. Sadoulet. 2015. "Social Networks and the Decision to Insure: Evidence from Randomized Experiments in China." *American Economic Journal: Applied Economics* 7 (2):81–108.

Chandrasekhar, Arun G and Randall Lewis. 2011. "Econometrics of sampled networks." Working Paper.

Charness, G., M. Corominas-Bosch, and G.R. Frechette. 2007. "Bargaining and network structure: An experiment." *Journal of Economic Theory* 136 (1):28–65.

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

Corominas-Bosch, Margarida. 2004. "Bargaining in a network of buyers and sellers." *Journal of Economic Theory* 115 (1):35–77.

Dar, MH., A. de Janvry, K. Emerick, D. Raitzer, and E. Sadoulet. 2013. "Flood-tolerant rice reduces yield variability and raises expected yield, differentially benefitting socially disadvantaged groups." *Scientific Reports* 3:3315.

de la Sierra, Raul Sanchez. 2017. "Ethnic Contract Enforceability." *Working Paper* .

Delêtre, Marc, Doyle B McKey, and Trevor R Hodkinson. 2011. "Marriage exchanges, seed exchanges, and the dynamics of manioc diversity." *Proceedings of the National Academy of Sciences* 108 (45):18249–18254.

Duflo, Esther, Michael Kremer, and Jonathan Robinson. 2011. "Nudging Farmers to Use Fertilizer: Theory and Experimental Evidence from Kenya." *American Economic Review* 101 (6):2350–2390.

Elliott, Matthew. 2015. "Inefficiencies in networked markets." *American Economic Journal: Microeconomics* 7 (4):43–82.

Emerick, Kyle, Alain de Janvry, Elisabeth Sadoulet, and Manzoor Dar. 2016. "Technological innovations, downside risk, and the modernization of agriculture." *American Economic Review* 106 (6):1537–1561.

Fisman, Raymond, Daniel Paravisini, and Vikrant Vig. 2017. "Cultural proximity and loan outcomes." *American Economic Review* 107 (2):457–492.

Fogli, Alessandra and Laura Veldkamp. 2016. "Germs, social networks and growth." NBER Working Paper.

Gale, D.M. and S. Kariv. 2009. "Trading in networks: A normal form game experiment." *American Economic Journal: Microeconomics* :114–132.

Huang, Jingfeng, Xiuzhen Wang, Xinxing Li, Hanqin Tian, and Zhuokun Pan. 2013. "Remotely sensed rice yield prediction using multi-temporal NDVI data derived from NOAA's-AVHRR." *PloS one* 8 (8):e70816.

Jack, Kelsey. 2011. "Market inefficiencies and the adoption of agricultural technologies in developing countries." *White paper, Agricultural Technology Adoption Initiative (Abdul Latif Jameel Poverty Action Lab/MIT, Cambridge, MA* .

Karlan, Dean S, Robert Osei, Isaac Osei-Akoto, Christopher Udry et al. 2014. “Agricultural decisions after relaxing credit and risk constraints.” *Quarterly Journal of Economics* :597–652.

Kranton, R.E. and D.F. Minehart. 2001. “A theory of buyer-seller networks.” *The American Economic Review* :485–508.

Kremer, M. and E. Miguel. 2007. “The Illusion of Sustainability.” *The Quarterly Journal of Economics* 122 (3):1007–1065.

Labeyrie, Vanesse, Mathieu Thomas, Zachary K Muthamia, and Christian Leclerc. 2016. “Seed exchange networks, ethnicity, and sorghum diversity.” *Proceedings of the National Academy of Sciences* 113 (1):98–103.

Manea, M. 2011. “Bargaining in stationary networks.” *The American Economic Review* 101 (5):2042–2080.

Manski, C.F. 1993. “Identification of endogenous social effects: The reflection problem.” *The Review of Economic Studies* 60 (3):531–542.

Mazzocco, Maurizio and Shiv Saini. 2012. “Testing efficient risk sharing with heterogeneous risk preferences.” *The American Economic Review* 102 (1):428–468.

McGuire, Shawn and Louise Sperling. 2016. “Seed systems smallholder farmers use.” *Food Security* 8 (1):179–195.

Munshi, Kaivan and Mark Rosenzweig. 2016. “Networks and misallocation: Insurance, migration, and the rural-urban wage gap.” *American Economic Review* 106 (1):46–98.

Nagavarapu, Sriniketh and Sheetal Sekhri. 2016. “Informal monitoring and enforcement mechanisms in public service delivery: Evidence from the public distribution system in India.” *Journal of Development Economics* 121:63–78.

Suri, T. 2011. “Selection and comparative advantage in technology adoption.” *Econometrica* 79 (1):159–209.

Yamano, Takashi, Maria Luz Malabayabas, and Architesh Panda. 2013. “Perception, Adoption, and Realized Benefits of Swarna-Sub1 in Eastern India.” *Stress Tolerant Rice for Africa and South Asia (STRASA) Economic Brief* .

Figures

Figure 1: Allocation of Swarna-Sub1 output across different uses by original recipients

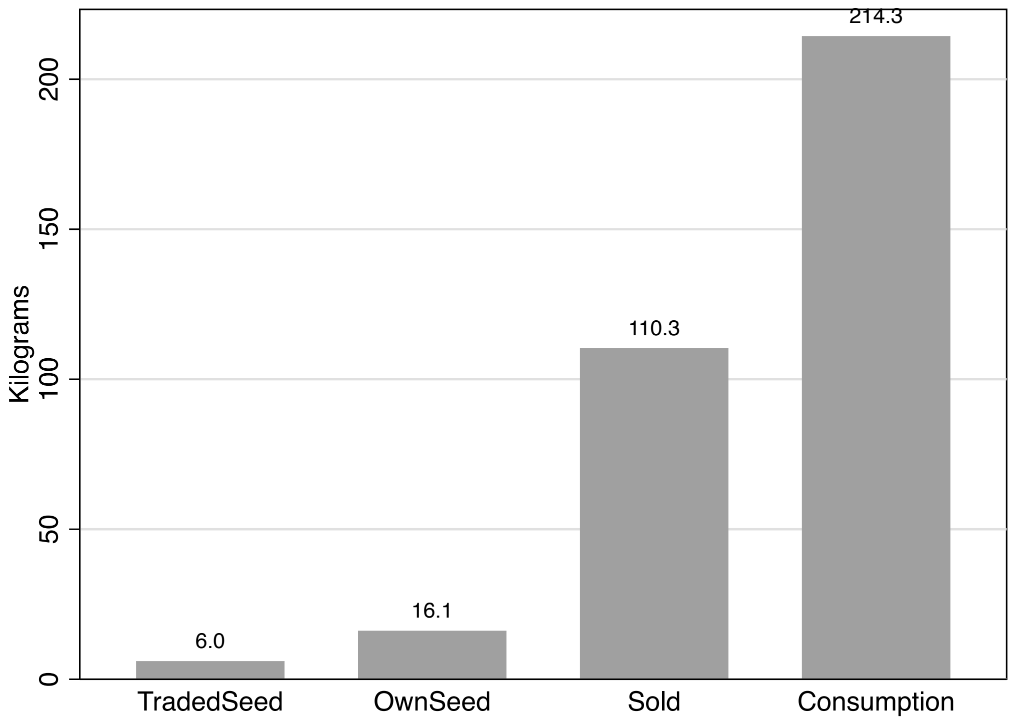

Notes: Data are from July 2013 follow-up survey where each original recipient was asked about their total Swarna-Sub1 harvest and its allocation across uses. Bar heights are average values across the 404 farmers surveyed. “Traded seed” is the output that was either sold, traded, or given to other farmers to be used as seed. “Own seed” is the output for use as seed in the upcoming season. “Sold” is output that had been sold as grain for consumption while “Consumption” is the output that the farmer’s household had already consumed or set aside for their own consumption.

Figure 2: Adoption in door-to-door sales and farmer-to-farmer networks

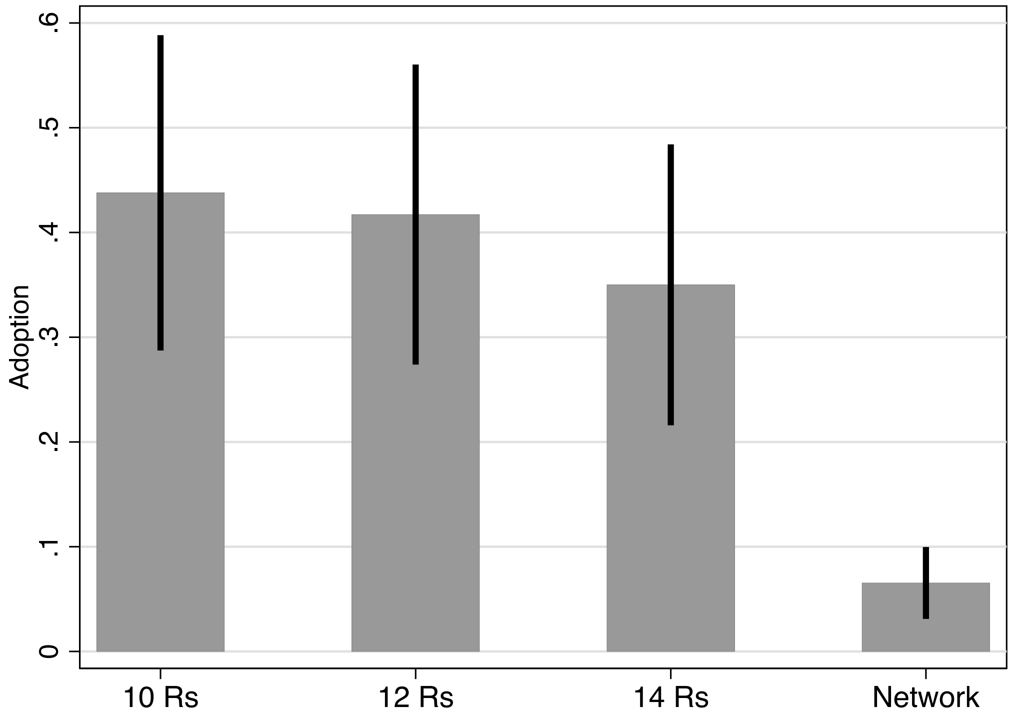

Notes: Figure displays the raw adoption rates for the sample of non-recipients during the 2013 agricultural season. The bands represent 95 percent confidence intervals where standard errors are clustered at the village level.

Figure 3: Densities of estimated returns of adopters, by treatment

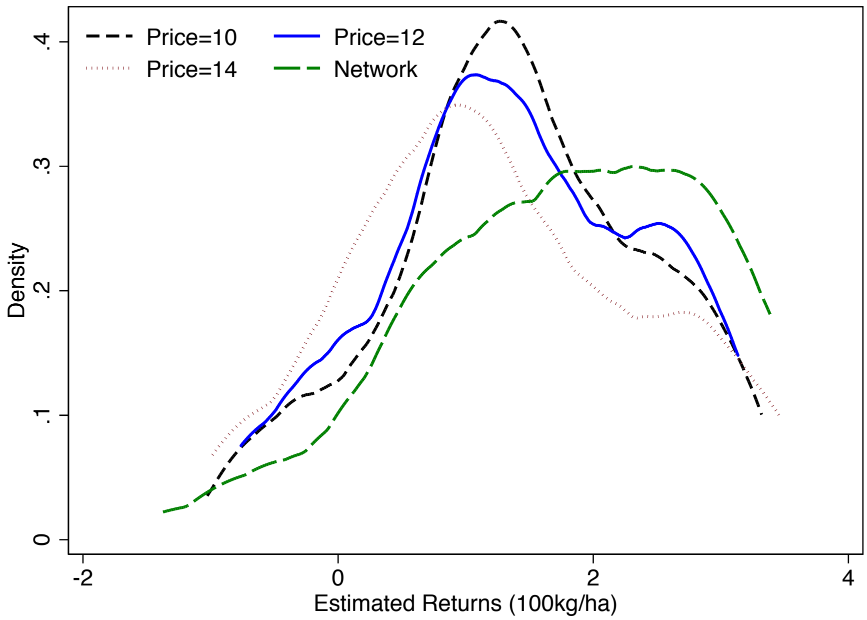

Notes: Figure displays kernel densities of estimated returns, by treatment group. Densities are estimated only for the group of farmers that adopted Swarna-Sub1 for the 2013 wet season.

Figure 4: Adoption rates over time as a function of treatment status

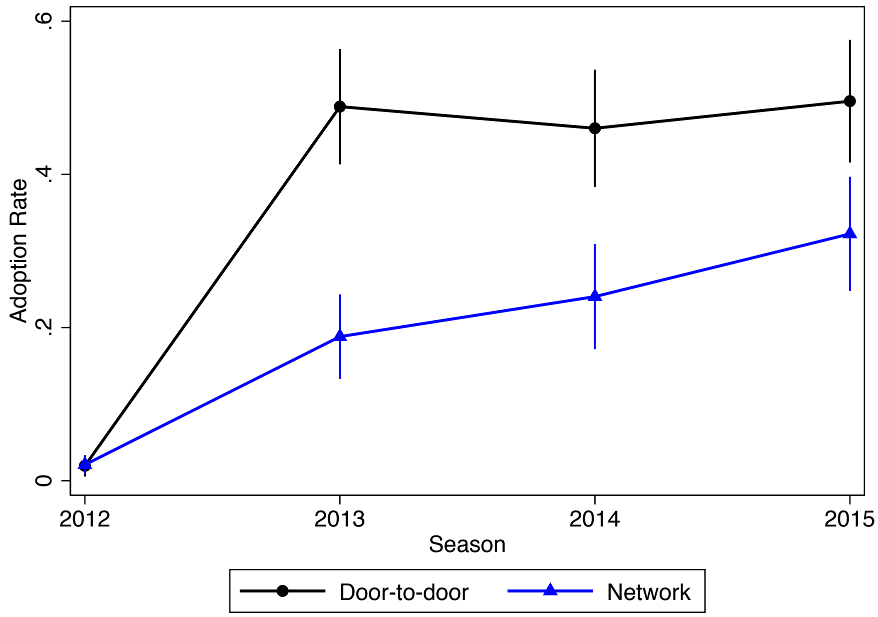

Notes: Figure displays the adoption rate as a function of time for door-to-door and network villages. Adoption was estimated using the long-term folllowup survey in July 2015. The sample consists of 1,139 farmers from the original sample of buyers that were reached again during this survey.

Figure 5: Estimated losses from slow diffusion in farmer-to-farmer networks

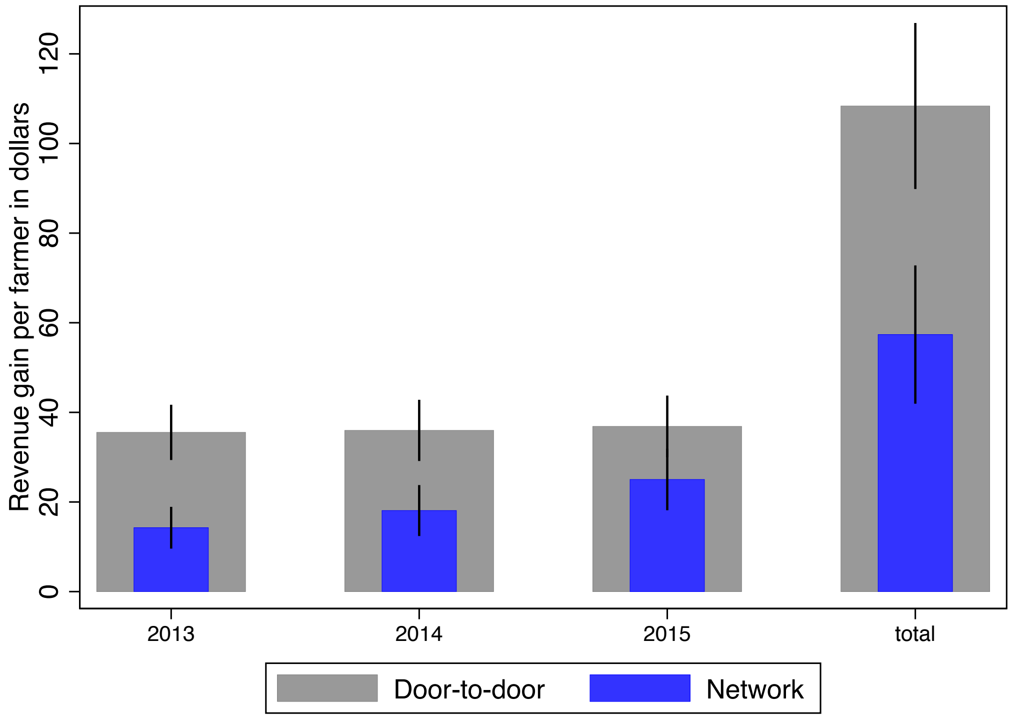

Notes: The height of each bar is the per-farmer estimated revenue gain from Swarna-Sub1. The grey bars are for door-to-door villages and the blue bars are for network villages. The differences between the two bars represent the per-farmer losses from slow diffusion in networks. The black bands represent 95 percent confidence intervals. The total revenue gain (fourth bar) is the sum of the revenue gains from 2013-2015.

Figure 6: Evolution of satellite-based greenness on fields cultivated by sample farmers

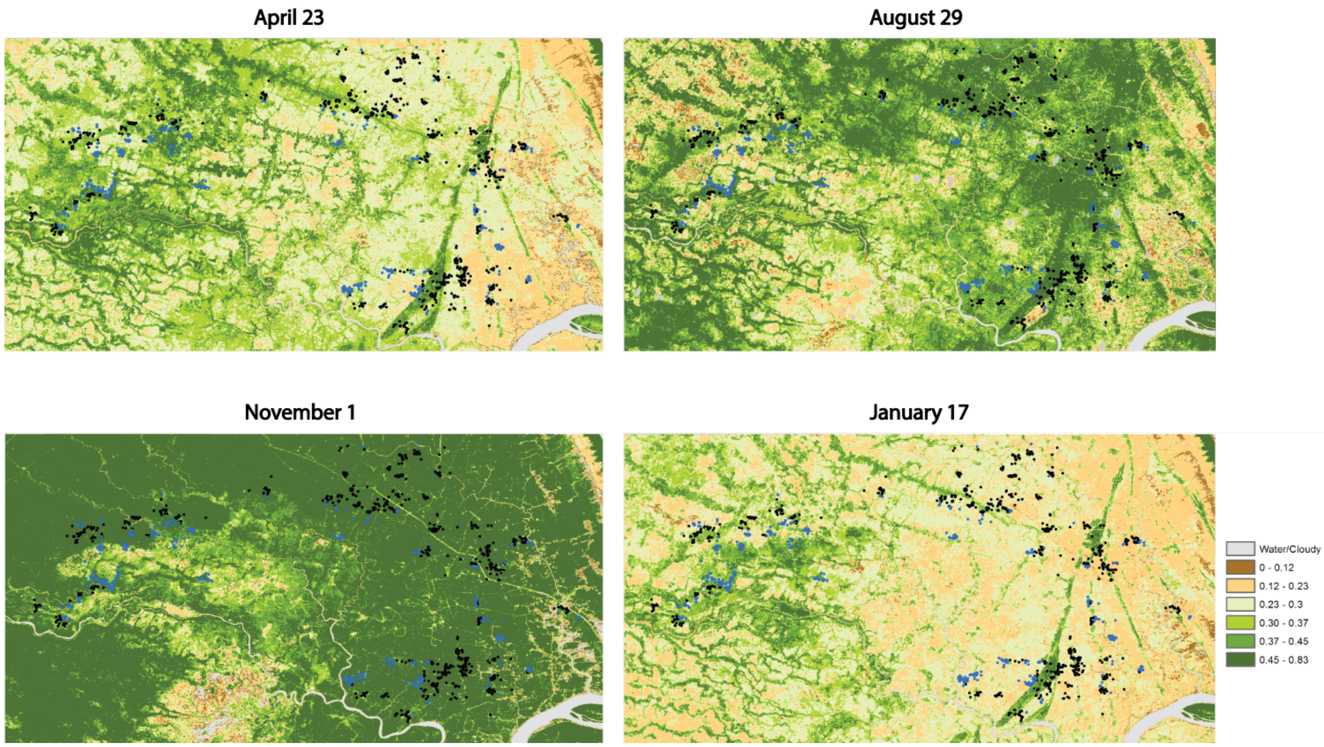

Notes: Figure shows satellite measures of NDVI before the growing season (April 23) during the growing season (August 29th and November 1) and after the growing season (January 17th). The images are from the 2013-2014 agricultural season, i.e. the first season after the door-to-door sales. The spatial resolution is 30 meters, implying that each pixel is around 0.09 hectares, or slightly less than the amount of area that can be cultivated with 5 kg of seed. The dots are fields cultivated by farmers in the sample, where black dots represent fields that were not affected by flooding and blue dots represent flooded fields.

Figure 7: Visualization of the effect of adoption on satellite-based measure of greenness

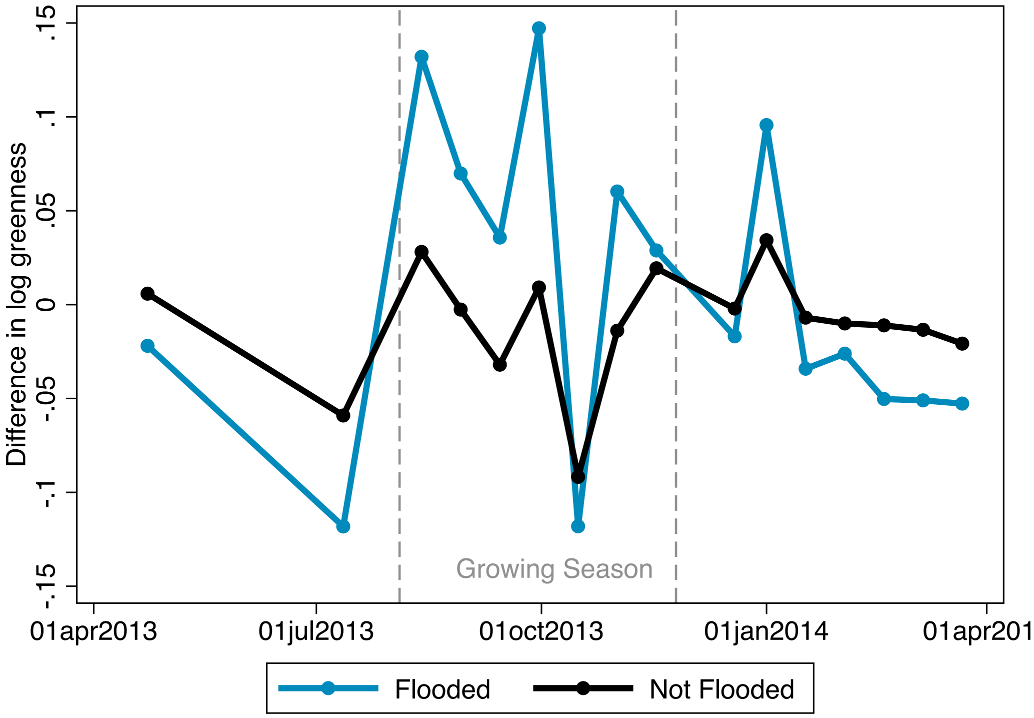

Notes: The figure shows the difference in vegetation greenness (log NDVI) between adopters and non-adopters. The light blue line are these gains from adoption in areas that were flooded and the black line is the gains from adoption in areas that were not flooded.

Tables

Table 1: Summary statistics of non-recipient sample

	Means		
	Network	Door-to-door	p-value
Area cultivated in 2012 (acres)	2.714	3.055	0.075
Land owned (acres)	1.674	1.835	0.326
Swarna user in 2012	0.723	0.695	0.624
Rice yield in 2012 (kg per acre)	1146.918	1188.065	0.448
Ag cooperative member	0.448	0.457	0.820
Monthly income of highest earning member (Rs)	3492.811	3685.515	0.450
Farmer is SC	0.240	0.168	0.266
Household head at least primary education	0.667	0.693	0.520
Thatched roof	0.689	0.775	0.016
Owns private tubewell	0.231	0.148	0.124
Access to electricity	0.893	0.894	0.962
Below the poverty line card	0.610	0.649	0.368
Villagers in sample same subcaste	8.176	8.890	0.492
Villagers in sample same surname	5.476	5.586	0.895
Original recipients same subcaste	2.121	2.133	0.970
Original recipients same surname	1.417	1.485	0.770
Original recipient houses w/in 50 meters	1.176	1.048	0.523
Bought seeds from government	0.441	0.419	0.624
Used seeds from previous harvest	0.652	0.675	0.526
Obtained seeds from neighboring farmer	0.126	0.159	0.305

The data are from the February 2013 survey with non-recipients. The survey took place 3 months before the door-to-door sales. Columns 1 and 2 show mean values of each characteristic in network and door-to-door villages, respectively. Column 3 gives the p-value for the joint test of equality where standard errors are adjusted for clustering at the village level.

Table 2: Estimated difference between adoption in networks and demand revealed in door-to-door sales

	(1)	(2)	(3)
Door-to-door and Price=10	0.380***		
	(0.077)		
Door-to-door and Price=12	0.357***		
	(0.066)		
Door-to-door and Price=14	0.275***		
	(0.061)		
Door-to-door treatment		0.336***	0.337***
		(0.043)	(0.043)
Farmer is SC			-0.060
			(0.040)
Farmer has BPL card			-0.054*
			(0.030)
Land cultivated in 2012			0.004
			(0.007)
Ag. cooperative member			-0.019
			(0.023)
Swarna user in 2012			0.090***
			(0.033)
Strata Fixed Effects	Yes	Yes	Yes
Mean of Dep Variable: Network	0.07	0.07	0.07
Number of Observations	1150	1150	1134
R squared	0.190	0.185	0.203

Dependent variable is 1 if farmer adopted Swarna-Sub1 for 2013 wet season. *Door-to-door treatment* is 1 for villages where farmers had the opportunity to obtain the technology from a door-to-door salesperson. Land cultivated in 2012 is measured in acres. Standard errors are clustered at the village level. Asterisks indicate statistical significance at the 1 percent ***, 5 percent **, and 10 percent * levels.

Table 3: Relative targeting effectiveness

	All adopters		
	(1)	(2)	(3)
	Return	Log Return	Flood severity (1-10)
Door-to-door treatment	-0.384	-0.281***	-0.813
	(0.241)	(0.098)	(0.518)
Constant	1.742***	0.581***	5.250***
	(0.219)	(0.063)	(0.463)
Mean of Dep Variable: Network	1.742	0.581	5.250
Number of Observations	266	233	267
R squared	0.016	0.022	0.026

The data are limited to the sample of farmers that cultivated Swarna-Sub1 for the 2013 wet season. The dependent variable in column 1 is the expected return of Swarna-Sub1, measured in quintiles (1 quintile = 100 kg) per hectare. The dependent variable in column 2 is log of the expected return. The dependent variable in column 3 is a subjective measure of the flood severity of the plot where Swarna-Sub1 was being planted. This variable ranges from 1-10 and was collected during the final follow-up survey with all non-recipients. *Door-to-door treatment* is 1 for villages where farmers had the opportunity to obtain the technology from a door-to-door salesperson. Standard errors are clustered at the village level. Asterisks indicate statistical significance at the 1 percent ***, 5 percent **, and 10 percent * levels.

Table 4: Effects of adoption on satellite-based measures of plant greenness

	Flooding threshold:	
	(1)	(2)
	250 m	500 m
Adopter * Flooded * Growing Season	0.101**	0.091**
	(0.049)	(0.043)
Adopter * Flooded	-0.025	-0.018
	(0.035)	(0.029)
Adopter * Growing Season	-0.009	-0.014
	(0.024)	(0.027)
Growing Season * Flooded	-0.314***	-0.231***
	(0.057)	(0.056)
Growing Season	0.427***	0.432***
	(0.029)	(0.029)
Adopter	-0.012	-0.013
	(0.019)	(0.017)
Flooded	0.087**	0.054*
	(0.035)	(0.032)
Mean of Dep. Variable	-1.18	-1.18
Number of Observations	18689	18689
R squared	0.125	0.117

The dependent variable in both columns is the log of the NDVI value of the field. The data consist of the 8 day NDVI composites from the Landsat 8 satellite (available via Google Earth Engine API). The coordinates of each plot were matched to Landsat images from 4/23/2013, 7/12/2013, 8/13/2013, 8/29/2013, 9/14/2013, 9/30/2013, 10/16/2013 11/1/2013, 11/17/2013, 12/19/2013, 1/1/2014, 1/17/2014, 2/2/2014, 2/18/2014, 3/6/2014, and 3/22/2014. The growing season extends from late July (transplanting) to mid November (harvesting). Flooded plots were identified using daily flood layers generated from NASA's Modis satellite. Standard errors are clustered at the village level. Asterisks indicate statistical significance at the 1 percent ***, 5 percent **, and 10 percent * levels.

Table 5: Estimated peer effects in network and door-to-door villages

	(1)	(2)	(3)	(4)
Door-to-door Treatment	0.332*** (0.056)		0.367*** (0.065)	
Door-to-door Treatment * Original recipients w/ same surname	-0.075* (0.043)	-0.111** (0.045)		
Original recipients w/ same surname	0.035 (0.026)	0.081** (0.031)		
Total number w/ same surname	-0.008 (0.008)	-0.025*** (0.009)		
Door-to-door Treatment * Total number w/ same surname	0.021 (0.014)	0.035** (0.014)		
Door-to-door Treatment * Original recipients same sub-caste			-0.056* (0.030)	-0.051 (0.035)
Original recipients same sub-caste			0.040* (0.021)	0.044** (0.020)
Total number same sub-caste			-0.010 (0.007)	-0.011 (0.007)
Door-to-door Treatment * Total number same sub-caste			0.011 (0.009)	0.015 (0.010)
Strata Fixed Effects	Yes	No	Yes	No
Village Fixed Effects	No	Yes	No	Yes
Mean of Dep Variable: Network	0.07	0.07	0.07	0.07
Mean Original recipients w/ same surname	1.45	1.45		
Mean Total number w/ same surname	5.53	5.53		
Mean Original recipients same sub-caste			2.12	2.12
Mean Total number same sub-caste			8.53	8.53
Number of Observations	1135	1135	1135	1135
R squared	0.191	0.410	0.192	0.404

Dependent variable is 1 if farmer adopted Swarna-Sub1 for 2013 wet season. *Door-to-door treatment* is 1 for villages where farmers had the opportunity to obtain the technology from a door-to-door salesperson. Standard errors are clustered at the village level. Asterisks indicate statistical significance at the 1 percent ***, 5 percent **, and 10 percent * levels.

Table 6: Long run peer effects

	(1) 2013	(2) 2014	(3) 2015
Door-to-door Treatment	0.259*** (0.064)	0.176** (0.071)	0.090 (0.071)
Door-to-door Treatment * Original recipients w/ same surname	-0.108*** (0.034)	-0.141*** (0.038)	-0.081** (0.040)
Original recipients w/ same surname	0.054** (0.021)	0.072*** (0.022)	0.058*** (0.022)
Total number w/ same surname	-0.020*** (0.007)	-0.027*** (0.007)	-0.027*** (0.007)
Door-to-door Treatment * Total number w/ same surname	0.036*** (0.011)	0.045*** (0.012)	0.037*** (0.013)
Strata Fixed Effects	Yes	Yes	Yes
Mean of Dep Variable: Network	0.19	0.24	0.32
Number of Observations	1137	1137	1137
R squared	0.112	0.081	0.081

Dependent variable in all columns is 1 if the farmer adopted Swarna-Sub1 in the given year. Standard errors are clustered at the village level. Asterisks indicate statistical significance at the 1 percent ***, 5 percent **, and 10 percent * levels.

Table 7: Effects of door-to-door treatment on estimated revenue gains as a function of surname connections

	(1) 2013	(2) 2014	(3) 2015	(4) Total
Door-to-door Treatment	18.483***	16.829**	5.559	40.871**
	(5.798)	(6.685)	(7.359)	(17.776)
Door-to-door Treatment * Original recipients w/ same surname	-8.316**	-10.707***	-5.932	-24.955**
	(3.196)	(3.440)	(4.218)	(9.784)
Door-to-door Treatment * Total number w/ same surname	2.681**	2.996***	2.681**	8.357***
	(1.090)	(1.050)	(1.317)	(3.033)
Total number w/ same surname	-1.395***	-1.917***	-2.019**	-5.331***
	(0.527)	(0.602)	(0.932)	(1.772)
Original recipients w/ same surname	3.393*	4.842**	3.684	11.919*
	(1.934)	(2.078)	(3.065)	(6.284)
Constant	17.129***	21.784***	30.920***	69.833***
	(3.374)	(4.065)	(5.217)	(10.496)
Mean of Dep Variable: Network	14.25	18.09	25.02	57.36
Number of Observations	1137	1137	1137	1137
R squared	0.071	0.050	0.022	0.054

Dependent variable in columns 1-3 is the estimated revenue gain from Swarna-Sub1 in the year corresponding to the column label. The dependent variable in column 4 is the sum of columns 1 to 3. The unit of the dependent variable is dollars in all regressions. Standard errors are clustered at the village level. Asterisks indicate statistical significance at the 1 percent ***, 5 percent **, and 10 percent * levels.

Appendix

Derivation of expected returns of adopters

The expected return of adopters in social networks is $E(R|R+u-c>v+\underline{c})$. Using properties of the multivariate normal distribution, this is written as

$$E(R|R+u-c>v+\underline{c}) = \mu_R + \sigma_R\tilde{\rho}E\left(\frac{R+u-c-\mu_r+\mu_c}{\sqrt{\sigma_R^2+\sigma_c^2+\sigma_u^2-2\rho\sigma_c\sigma_R}}\middle| R+u-c>v+\underline{c}\right), \quad \text{(A1)}$$

where $\tilde{\rho}$ is the correlation between R and $R+u-c$. This expression can be rewritten as

$$\begin{aligned}E(R|R+u-c>v+\underline{c}) =& \mu_R + \frac{\sigma_R}{\sqrt{\sigma_R^2+\sigma_c^2+\sigma_u^2-2\rho\sigma_c\sigma_R}}\tilde{\rho}E(R+u-c|R+u-c>v+\underline{c}) \quad \text{(A2)}\\ &- \frac{\tilde{\rho}\sigma_R(\mu_R-\mu_c)}{\sqrt{\sigma_R^2+\sigma_c^2+\sigma_u^2-2\rho\sigma_c\sigma_R}}.\end{aligned}$$

Given that $R+u-c$ is distributed normally, $E(R+u-c|R+u-c>v+\underline{c})$ can be written as

$$\mu_R-\mu_c+\sqrt{\sigma_R^2+\sigma_u^2+\sigma_c^2-2\rho\sigma_R\sigma_c} * M\left(\frac{v+\underline{c}-\mu_R+\mu_c}{\sqrt{\sigma_R^2+\sigma_u^2+\sigma_c^2-2\rho\sigma_R\sigma_c}}\right), \quad \text{(A3)}$$

where $M(z)=\frac{\phi(z)}{1-\Phi(z)}$ is the inverse Mill's ratio. Reinserting this into Equation A2 gives

$$E(R|R+u-c>v+\underline{c}) = \mu_R + \tilde{\rho}\sigma_R M\left(\frac{v+\underline{c}-\mu_R+\mu_c}{\sqrt{\sigma_R^2+\sigma_u^2+\sigma_c^2-2\rho\sigma_R\sigma_c}}\right). \quad \text{(A4)}$$

Since $\tilde{\rho}$ is the correlation between R and $R+u-c$, $\tilde{\rho}$ simplifies to

$$\tilde{\rho} = \frac{\sigma_R-\rho\sigma_c}{\sqrt{\sigma_R^2+\sigma_u^2+\sigma_c^2-2\rho\sigma_R\sigma_c}}. \quad \text{(A5)}$$

Combining Equations A4 and A5,

$$E(R|R+u-c>v) = \mu_R + \frac{\sigma_R(\sigma_R-\rho\sigma_c)}{\sqrt{\sigma_R^2+\sigma_u^2+\sigma_c^2-2\rho\sigma_R\sigma_c}} * M\left(\frac{v+\underline{c}-\mu_R+\mu_c}{\sqrt{\sigma_R^2+\sigma_u^2+\sigma_c^2-2\rho\sigma_R\sigma_c}}\right). \quad \text{(A6)}$$

This establishes the result. A similar derivation is used to verify the formula for $E(R|R+u>v)$ in the main text.

Appendix Figures and Tables

Table A1: Baseline characteristics for original recipients and non-recipients

	(1) Non-recipient	(2) Original recipient	(3) p-value: (1)-(2)
Rice acres in Kharif 2011	3.88	3.80	0.53
Acres flooded 4 days or less in Kharif 2011	1.25	1.25	0.94
Acres flooded 5 days or more in Kharif 2011	2.63	2.56	0.52
Acres grown with Swarna in Kharif 2011	1.95	1.88	0.34
Farmer is Scheduled Caste (SC)	0.20	0.18	0.46
Age of farmer	48.96	49.07	0.86
Farmer is lead farmer	0.09	0.11	0.29
Network degree	4.19	4.37	0.21

Data are from the short baseline survey that took place during the village meeting in May or June 2012. Column 1 gives mean values for farmers that were not selected as original recipients. Column 2 gives mean values for farmers that were randomly selected as original recipients. Column 3 gives the p-value for the test of equality of means. Network degree is defined as the number of links of a farmer from the baseline survey (undirected).

Table A2: Heterogeneity in adoption effects by household characteristics

	(1)
Door-to-door treatment	0.409***
	(0.093)
Farmer is SC	0.016
	(0.044)
Farmer has BPL card	-0.014
	(0.033)
Land cultivated in 2012	0.007
	(0.006)
Ag. cooperative member	-0.020
	(0.027)
Swarna user in 2012	0.032
	(0.026)
Education above primary	-0.006
	(0.021)
<u>*Door-to-door treatment interacted with:*</u>	
Farmer is SC	-0.197**
	(0.076)
Farmer has BPL card	-0.103
	(0.065)
Land cultivated in 2012	-0.001
	(0.014)
Ag. cooperative member	0.009
	(0.046)
Swarna user in 2012	0.115*
	(0.068)
Education above primary	-0.114**
	(0.048)
Strata Fixed Effects	Yes
Mean of Dep Variable: Network	0.07
Number of Observations	1131
R squared	0.224

Dependent variable is 1 if farmer adopted Swarna-Sub1 for 2013 wet season. *Door-to-door Treatment* is 1 for villages where farmers could either obtain the technology from door-to-door sales or from peers. Standard errors are clustered at the village level. Asterisks indicate statistical significance at the 1 percent ***, 5 percent **, and 10 percent * levels.

Table A3: Relative targeting effectiveness when dropping two farmers that obtained Swarna-Sub1 from local government office

	Adopters from peers or door-to-door		
	(1)	(2)	(3)
	Return	Log Return	Flood severity (1-10)
Door-to-door	-0.524***	-0.281***	-0.945*
treatment	(0.155)	(0.098)	(0.499)
Constant	1.882***	0.581***	5.382***
	(0.117)	(0.063)	(0.441)
Mean of Dep Variable: Network	1.882	0.581	5.382
Number of Observations	264	233	265
R squared	0.029	0.022	0.033

The data are limited to the sample of farmers that cultivated Swarna-Sub1 for the 2013 wet season and either obtained it from the door-to-door sales experiment or directly from a peer. The dependent variable in column 1 is the expected return of Swarna-Sub1, measured in quintiles (1 quintile = 100 kg) per hectare. The dependent variable in column 2 is log of the expected return. The dependent variable in column 3 is a subjective measure of the flood severity of the plot where Swarna-Sub1 was being planted. This variable ranges from 1-10 and was collected during the final follow-up survey with all non-recipients. *Door-to-door treatment* is 1 for villages where farmers had the opportunity to obtain the technology from a door-to-door salesperson. Standard errors are clustered at the village level. Asterisks indicate statistical significance at the 1 percent ***, 5 percent **, and 10 percent * levels.

Table A4: Long run effects of door-to-door treatment on adoption

	(1) 2013	(2) 2014	(3) 2015
Door-to-door Treatment	0.301*** (0.048)	0.222*** (0.051)	0.175*** (0.051)
Strata Fixed Effects	Yes	Yes	Yes
Mean of Dep Variable: Network	0.19	0.24	0.32
Number of Observations	1139	1139	1139
R squared	0.102	0.064	0.069

Dependent variable in all columns is 1 if the farmer adopted Swarna-Sub1 in the year corresponding to the column label. Standard errors are clustered at the village level. Asterisks indicate statistical significance at the 1 percent ***, 5 percent **, and 10 percent * levels.

Table A5: Effects of door-to-door treatment on estimated revenue gains from the new technology

	(1) 2013	(2) 2014	(3) 2015	(4) Total
Door-to-door	21.389^{***}	18.119^{***}	11.883^{***}	51.391^{***}
Treatment	(3.933)	(4.366)	(4.440)	(11.799)
Strata Fixed Effects	Yes	Yes	Yes	Yes
Mean of Dep Variable: Network	14.25	18.09	25.02	57.36
Number of Observations	1137	1137	1137	1137
R squared	0.065	0.050	0.053	0.058

Dependent variable in columns 1-3 is the estimated revenue gain from Swarna-Sub1 in the year corresponding to the column label. The dependent variable in column 4 is the sum of columns 1 to 3. The unit of the dependent variable is dollars in all regressions. Standard errors are clustered at the village level. Asterisks indicate statistical significance at the 1 percent ***, 5 percent **, and 10 percent * levels.

Table A6: Robustness of productivity effects to village and farmer fixed effects

	(1) Village FE	(2) Farmer FE
Adopter * Flooded * Growing Season	0.103**	0.098*
	(0.049)	(0.053)
Adopter * Flooded	-0.032	
	(0.030)	
Adopter * Growing Season	-0.010	-0.007
	(0.025)	(0.026)
Growing Season * Flooded	-0.321***	-0.324***
	(0.058)	(0.061)
Growing Season	0.421***	0.425***
	(0.029)	(0.031)
Adopter	-0.002	
	(0.020)	
Flooded	0.034	
	(0.029)	
Mean of Dep. Variable	-1.18	-1.18
Number of Observations	18689	18689
R squared	0.176	0.280

The dependent variable in both columns is the log of the NDVI value of the field. The data consist of the 8 day NDVI composites from the Landsat 8 satellite (available via Google Earth Engine API). The coordinates of each plot were matched to Landsat images from 4/23/2013, 7/12/2013, 8/13/2013, 8/29/2013, 9/14/2013, 9/30/2013, 10/16/2013 11/1/2013, 11/17/2013, 12/19/2013, 1/1/2014, 1/17/2014, 2/2/2014, 2/18/2014, 3/6/2014, and 3/22/2014. The growing season extends from late July (transplanting) to mid November (harvesting). Flooded plots were identified using daily flood layers generated from NASA's Modis satellite. Standard errors are clustered at the village level. Asterisks indicate statistical significance at the 1 percent ***, 5 percent **, and 10 percent * levels.

Table A7: Robustness of estimated peer effects to different subsamples and nonlinear model

	Variation in adoption		Full sample	
	(1) OLS	(2) OLS	(3) Probit	(4) Probit
Door-to-door Treatment	0.268*** (0.080)	0.206** (0.095)	0.349*** (0.055)	0.357*** (0.059)
Door-to-door Treatment * Original recipients w/ same surname	-0.159** (0.061)		-0.076** (0.036)	
Original recipients w/ same surname	0.110** (0.053)		0.027* (0.016)	
Total number w/ same surname	-0.015 (0.011)		-0.010 (0.013)	
Door-to-door Treatment * Total number w/ same surname	0.034** (0.016)		0.021 (0.014)	
Door-to-door Treatment * Original recipients same sub-caste		-0.120** (0.045)		-0.063** (0.027)
Original recipients same sub-caste		0.075* (0.041)		0.024* (0.013)
Total number same sub-caste		-0.024 (0.017)		-0.014 (0.012)
Door-to-door Treatment * Total number same sub-caste		0.031* (0.017)		0.017 (0.012)
Strata Fixed Effects	Yes	Yes	Yes	Yes
HH controls	Yes	Yes	Yes	Yes
Mean of Dep Variable: Network	0.18	0.18	0.07	0.07
Number of Observations	744	744	1134	1134
R squared	0.120	0.118		

Data in columns 1 and 2 are limited to villages where at least one farmer adopted Swarna-Sub1 for 2013 wet season. Dependent variable in all columns is 1 if farmer adopted Swarna-Sub1 for 2013 wet season. Columns 3 and 4 present marginal effects calculated from probit coefficients, along with standard errors calculated from the delta method. Control variables are indicator for SC, indicator for holding BPL card, land area cultivated in 2012 wet season, indicator for member of agricultural cooperative, and indicator for Swarna cultivator in 2012 wet season. Standard errors are clustered at the village level. Asterisks indicate statistical significance at the 1 percent ***, 5 percent **, and 10 percent * levels.

Table A8: Robustness of estimated peer effects to measurement of peer influence in shares rather than levels

	(1)	(2)	(3)	(4)
Door-to-door	0.439^{***}		0.439^{***}	
Treatment	(0.047)		(0.054)	
Door-to-door	-0.373^{***}	-0.346^{***}		
Treatment * Share of same surname that are recipients	(0.112)	(0.130)		
Share of same	0.206^{**}	0.202^{**}		
surname that are recipients	(0.080)	(0.095)		
Door-to-door			-0.398^{**}	-0.411^{**}
Treatment * Share of same sub-caste that are recipients			(0.167)	(0.184)
Share of same			0.125	0.174^{*}
sub-caste that are recipients			(0.090)	(0.099)
Strata Fixed Effects	Yes	No	Yes	No
Village Fixed Effects	No	Yes	No	Yes
Household controls	Yes	Yes	Yes	Yes
Mean of Dep Variable: Network	0.07	0.07	0.07	0.07
Number of Observations	1008	1008	1055	1055
R squared	0.220	0.435	0.218	0.434

Dependent variable in all columns is 1 if farmer adopted Swarna-Sub1 for 2013 wet season. Control variables are indicator for SC, indicator for holding BPL card, land area cultivated in 2012 wet season, indicator for member of agricultural cooperative, and indicator for Swarna cultivator in 2012 wet season. Standard errors are clustered at the village level. Asterisks indicate statistical significance at the 1 percent ***, 5 percent **, and 10 percent * levels.

Figure A1: Distribution of expected returns of Swarna-Sub1

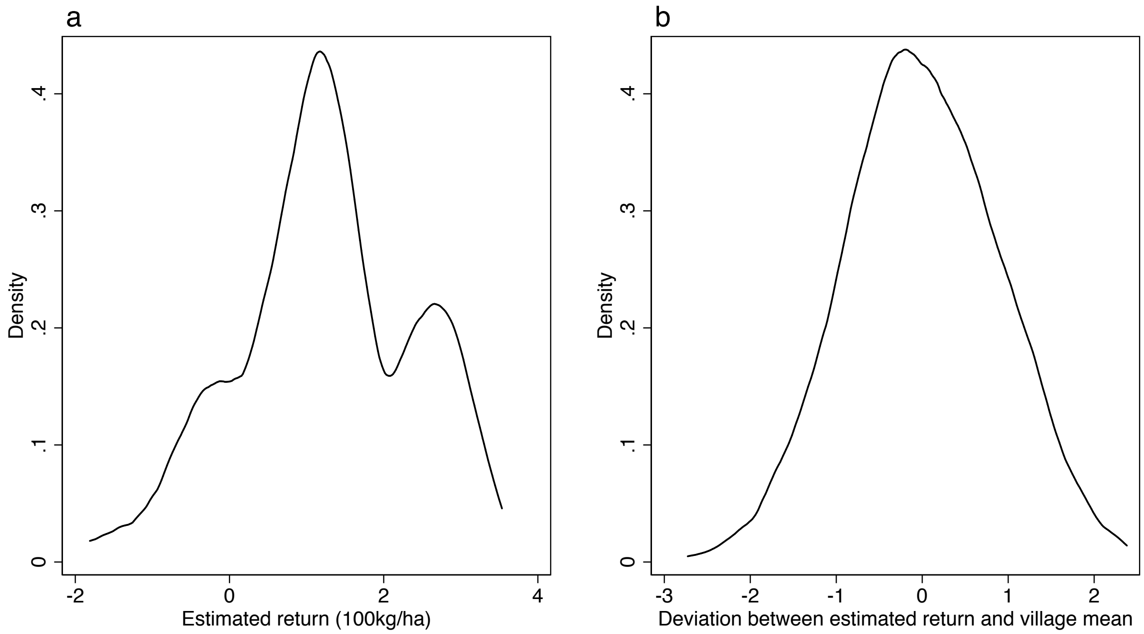

Notes: Figure shows densities of raw estimated returns (Panel A) and deviations between estimated returns and village averages (Panel B). Plot-level recall on flood duration and impact estimates in Dar et al. (2013) were used to calculate expected returns for each farmer in the sample. The only source of variation in expected returns using this methodology is exposure of the farmers' land to flooding.

Table A9: Long run peer effects when measuring connectivity using sub-caste association

	(1) 2013	(2) 2014	(3) 2015
Door-to-door Treatment	0.336***	0.275***	0.169**
	(0.072)	(0.075)	(0.081)
Door-to-door Treatment * Original recipients same sub-caste	-0.045	-0.049	-0.006
	(0.038)	(0.040)	(0.040)
Original recipients same sub-caste	0.012	0.024	0.010
	(0.030)	(0.032)	(0.032)
Total number same sub-caste	-0.005	-0.003	-0.000
	(0.010)	(0.011)	(0.012)
Door-to-door Treatment * Total number same sub-caste	0.007	0.006	0.002
	(0.012)	(0.014)	(0.014)
Strata Fixed Effects	Yes	Yes	Yes
Mean of Dep Variable: Network	0.19	0.24	0.32
Number of Observations	1137	1137	1137
R squared	0.109	0.070	0.070

Dependent variable in all columns is 1 if the farmer adopted Swarna-Sub1 in the year corresponding to the column label. Standard errors are clustered at the village level. Asterisks indicate statistical significance at the 1 percent ***, 5 percent **, and 10 percent * levels.

Table A10: Dyadic regressions of network formation

	(1)	(2)
One farmer is recipient	0.013	0.022
	(0.014)	(0.015)
Both farmers are recipients	0.182***	0.207***
	(0.030)	(0.035)
Same sub-caste		0.035*
		(0.018)
Same surname		0.124***
		(0.018)
Houses within 25 m		0.006
		(0.017)
Plots within 100 m		0.009
		(0.015)
Village Fixed Effects	Yes	Yes
Mean of Dep Variable	0.380	0.385
Number of Observations	27633	24837
R squared	0.073	0.088

Data are from follow-up social network survey of all farmers. Dependent variable is 1 if either farmer in the dyad indicated a sharing link (i.e. an undirected network). Standard errors are clustered at the village level. Asterisks indicate statistical significance at the 1 percent ***, 5 percent **, and 10 percent * levels.

Table A11: Effects of different household characteristics on the probability of link formation

	(1)	(2)	(3)	(4)	(5)
Same sub-caste	0.079***				0.036**
	(0.016)				(0.018)
Same surname		0.136***			0.127***
		(0.015)			(0.017)
Houses within 25 m			0.043***		-0.002
			(0.015)		(0.017)
Plots within 100 m				0.021	0.006
				(0.014)	(0.014)
Village Fixed Effects	Yes	Yes	Yes	Yes	Yes
Mean of Dep Variable	0.380	0.380	0.380	0.384	0.385
Number of Observations	27633	27633	27427	24979	24837
R squared	0.071	0.080	0.066	0.066	0.080

Data are from follow-up social network survey of all farmers. Dependent variable is 1 if either farmer in the dyad indicated a sharing link (i.e. an undirected network). Standard errors are clustered at the village level. Asterisks indicate statistical significance at the 1 percent ***, 5 percent **, and 10 percent * levels.

Figure A2: Distribution of difference between total harvest of Swarna-Sub1 in year 1 and amount planted in year 2 in door-to-door villages

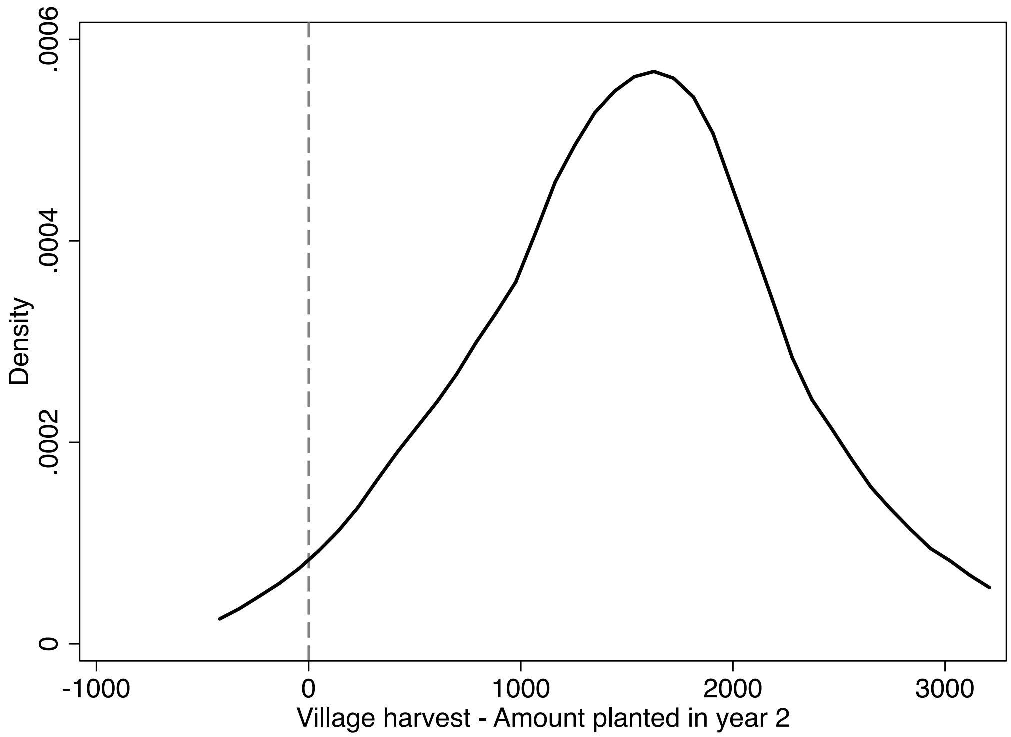

Notes: Data are for door-to-door villages. Figure shows the kernel density of difference between total year 1 harvest of Swarna-Sub1 by original recipients and aggregate amount of Swarna-Sub1 planted in village during year 2 (in kg). The amount planted during year 2 includes amount purchased from door-to-door sales, amount obtained directly from original recipients (by all farmers, not only farmers in the sample), and amount planted by original recipients.

Table A12: Heterogeneous effects according to network connectivity of original recipients

	(1)	(2)
Door-to-door treatment	0.256^{***}	0.256^{***}
	(0.063)	(0.063)
1 if recipient degree / non-recipient degree > median	-0.054	-0.047
	(0.035)	(0.038)
Door-to-door treatment*1 if recipient degree / non-recipient degree > median	0.159^{*}	0.157^{*}
	(0.088)	(0.088)
Farmer is SC		-0.071^{*}
		(0.041)
Farmer has BPL card		-0.061^{*}
		(0.032)
Land cultivated in 2012		0.004
		(0.007)
Ag. cooperative member		-0.025
		(0.024)
Swarna user in 2012		0.074^{**}
		(0.032)
Block Fixed Effects	Yes	Yes
Mean of Dep Variable: Network	0.07	0.07
Number of Observations	1135	1134
R squared	0.182	0.199

Dependent variable is 1 if farmer adopted Swarna-Sub1 for 2013 wet season. 1 if recipient / non-recipient degree > median is a village-level indicator for ratio of average degree of recipients to average degree of non-recipients being larger than the median. *Door-to-door treatment* is 1 for villages where farmers had the opportunity to obtain the technology from a door-to-door salesperson. Land cultivated in 2012 is measured in acres. Standard errors are clustered at the village level. Asterisks indicate statistical significance at the 1 percent ***, 5 percent **, and 10 percent * levels.

Table A13: Heterogeneous effects according to relative wealth of original recipients

	(1)	(2)
Door-to-door treatment	0.309***	0.317***
	(0.056)	(0.056)
Recipients realtively wealthy	0.077**	0.100**
	(0.036)	(0.038)
Door-to-door treatment * Recipients relatively wealthy	0.070	0.052
	(0.088)	(0.085)
Block Fixed Effects	Yes	Yes
Household controls	No	Yes
Mean of Dep Variable: Network	0.07	0.07
Number of Observations	1120	1119
R squared	0.194	0.213

Dependent variable is 1 if farmer adopted Swarna-Sub1 for 2013 wet season. *Recipients relatively wealthy* is an indicator for villages where the average wealth of original recipients divided by the average wealth of non-recipients is larger than the median. Wealth is defined as monthly average income for the highest earning household member. *Door-to-door treatment* is 1 for villages where farmers had the opportunity to obtain the technology from a door-to-door salesperson. Standard errors are clustered at the village level. Asterisks indicate statistical significance at the 1 percent ***, 5 percent **, and 10 percent * levels.

Figure A3: Relationship between village-level flooding in 2013 and adoption of Swarna-Sub1

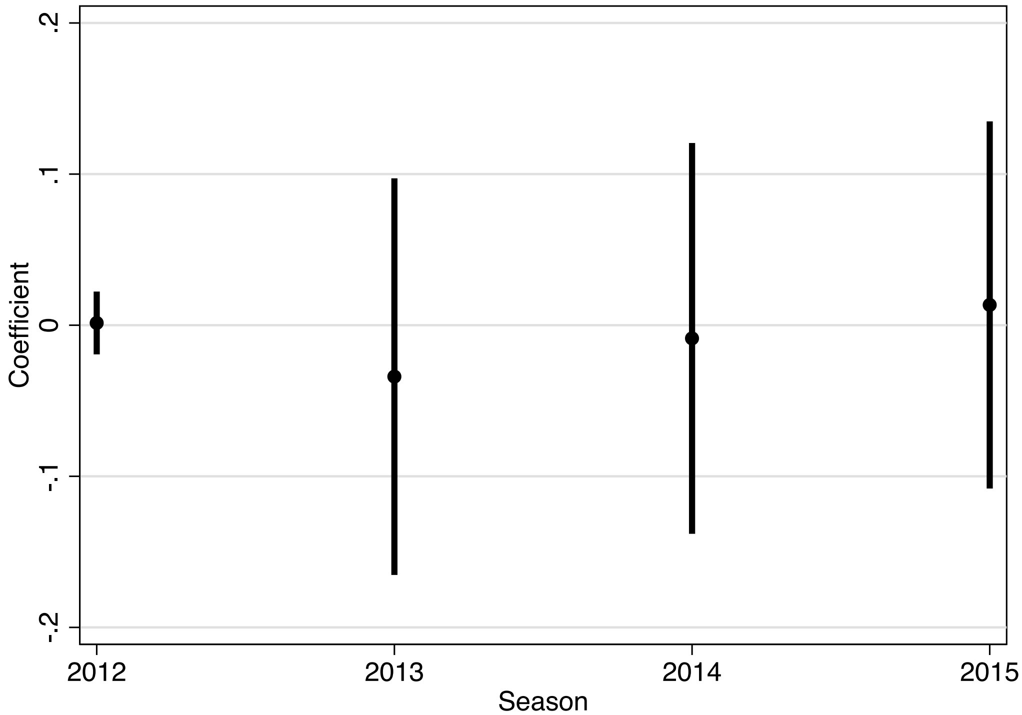

Notes: The figure shows estimates from separate regressions of adoption on an indicator for flood-affected villages and strata fixed effects. Flood-affected villages are defined as those where at least one field was flooded during the 2013 floods. Each point in the figure is the point estimate on the indicator for flood-affected villages and the bands denote 95 percent confidence intervals. Standard errors are clustered at the village level in each regression.

Figure A4: Color distinction between Swarna and Swarna-Sub1 seeds

Notes: The image shows unmilled rice seeds of Swarna (on the left) and Swarna-Sub1 (on the right). The color distinction between Swarna and Swarna-Sub1 eliminates the possibility that farmers fail to trade because of mistrust on the type of variety being exchanged.

Table A14: Heterogeneity of adoption effect according to preferences for quality seeds

	(1)	(2)
Door-to-door treatment	0.352***	0.376***
	(0.051)	(0.050)
Door-to-door treatment*Seed buyer in 2012	-0.036	
	(0.050)	
Seed buyer in 2012	-0.021	
	(0.024)	
Door-to-door treatment*Quality preference		-0.078
		(0.051)
Quality preference		-0.012
		(0.027)
Farmer is SC	-0.063	-0.054
	(0.041)	(0.039)
Farmer has BPL card	-0.055*	-0.057*
	(0.031)	(0.030)
Land cultivated in 2012	0.004	0.005
	(0.007)	(0.007)
Ag. cooperative member	-0.016	-0.007
	(0.024)	(0.023)
Swarna user in 2012	0.101***	0.091***
	(0.032)	(0.033)
Strata Fixed Effects	Yes	Yes
Mean of Dep Variable: Network	0.07	0.07
Number of Observations	1134	1134
R squared	0.206	0.209

Dependent variable is 1 if farmer adopted Swarna-Sub1 for 2013 wet season. *Door-to-door treatment* is 1 for villages where farmers had the opportunity to obtain the technology from a door-to-door salesperson. Standard errors are clustered at the village level. Asterisks indicate statistical significance at the 1 percent ***, 5 percent **, and 10 percent * levels.

Table A15: Effect of door-to-door sales on sales and exchanges to farmers outside the sample

	(1)	(2)
Door-to-door treatment	-0.057	-0.047
	(0.075)	(0.073)
Swarna-Sub1 harvest (100 kg)		0.056^{***}
		(0.018)
Farmer is SC		0.268^{**}
		(0.115)
Age of farmer		-0.002
		(0.002)
Farmer has BPL card		0.034
		(0.067)
Education above primary		-0.046
		(0.075)
Strata Fixed Effects	Yes	Yes
Mean of Dep Variable: Network	0.29	0.29
Number of Observations	394	393
R squared	0.024	0.101

Data are from the final survey with original recipients. Dependent variable is the number of farmers from outside the sample that a given original recipient sold or exchanged seeds with. *Door-to-door treatment* is 1 for villages where farmers had the opportunity to obtain the technology from a door-to-door salesperson. Standard errors are clustered at the village level. Asterisks indicate statistical significance at the 1 percent ***, 5 percent **, and 10 percent * levels.

Figure A5: Adoption rates over time as a function of treatment status

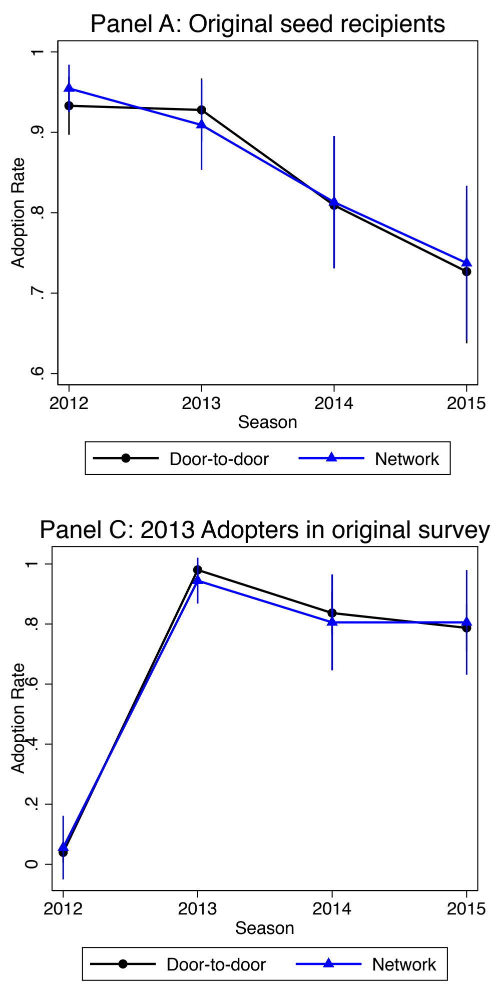

Notes: Figure displays the adoption rate as a function of time for door-to-door and network villages. Adoption was estimated using the long-term folllowup survey in July 2015. Panel A includes the 392 farmers that received seed minikits in 2012. Panel B includes the 4,738 farmers that did not receive minikits and were not part of the original sample that was surveyed in 2013. Panel C includes the 238 farmers that adopted Swarna-Sub1 at the time of the original 2013 follow-up survey.