

Can Network Theory-based Targeting Increase Technology

Adoption?

Lori Beaman
Northwestern Univ.

Ariel BenYishay
Coll. of William and Mary

Jeremy Magruder
UC-Berkeley

Ahmed Mushfiq Mobarak
Yale University

June 2015

Abstract

In order to induce farmers to adopt a productive new agricultural technology, we apply simple and complex contagion diffusion models on rich social network data from 200 villages in Malawi to identify seed farmers to target and train on the new technology. A randomized controlled trial compares these theory-driven network targeting approaches to simpler strategies that either rely on a government extension worker or an easily measurable proxy for the social network (geographic distance between households) to identify seed farmers. Both reduced form and structural estimates suggest a learning environment in which most farmers need to learn about the technology from multiple people before they adopt themselves.

JEL Codes: O16, O13

Keywords: Social Learning, Agricultural Technology Adoption, Complex Contagion, Malawi

* Contact: Beaman: l-beaman@northwestern.edu, BenYishay: abenyishay@wm.edu, Magruder: jmagruder@berkeley.edu, Mobarak: ahmed.mobarak@yale.edu. We thank the CEGA/JPAL Agricultural Technology Adoption Initiative (ATAI) and 3ie for financial support. We gratefully acknowledge the support and cooperation of Paul Fatch, Readwell Musopole and many other staff members of the Malawi Ministry of Agriculture. Thomas Coen, Niall Kelleher, Maria Jones, Ofer Cohen, Allen Baumgardner-Zuzik, and the IPA-Malawi country office provided invaluable support for data collection. Hossein Alidaee and Tetyana Zelenska provided excellent research assistance. We thank, without implicating, Arun Chandrasekhar, Matt Jackson, Kaivan Munshi, Chris Udry and numerous seminar audiences for very helpful comments.

1. Introduction

Technology diffusion is critical for growth and development (Alvarez et al. 2013, Perla and Tonetti 2014). Lack of credible information is one potential constraint to technology adoption, and social relationships can serve as important vectors through which individuals learn about, and are then convinced to adopt, new technologies.¹ Firms and extension services that market new products and technologies often rely on information cascades: they seed a few key entry points, and then allow the technology to diffuse via ambient social learning. This aspect of social learning could be manipulated to improve policy: if some entry points are better than others at inducing cascades, it would be valuable to identify ones that would maximize diffusion. We conduct a large-scale field experiment in which we randomly vary entry points in line with specific theories of network-based diffusion, in order to analyze the diffusion process for a productive new agricultural technology for arid regions of Africa.

The theoretical literature on diffusion processes is very rich², and for tractability, we refine our focus to an important class of diffusion models: threshold models, where individuals adopt if they are connected to at least a threshold number of adopters (e.g. Granovetter 1978; Centola and Macy 2007; Acemoglu *et al* 2011). Different formulations of the model yield different implications on whether entry points matter for diffusion, and which entry points would be most effective. Suppose an extension agency can train two farmers in a new technology. If a connection to just one adopter is sufficient to motivate adoption, then technology diffusion looks like viral infection, and the specific identity of the entry points in the network is not very critical: the diffusion rate will be similar regardless of who in the network is first seeded with the technology. Centola and Macy

¹ Large literatures in economics (Munshi 2008, Duflo and Saez 2003, Magruder 2010, Beaman 2012), finance (Beshears et al. 2013, Bursztyn et al. 2013), sociology (Rogers 1962), and medicine and public health (Coleman et al 1957; Doumit et al 2007) show that information and behaviors spread through inter-personal ties. The literature on diffusion of agricultural technologies (Griliches 1957, Foster and Rosenzweig 1995, Munshi 2004, Bandiera and Rasul 2006, Conley and Udry 2010) is most closely related to our work.

² See Jackson 2008 Chapter 7 for a review.

(2007) label this case “simple contagion.” In this case, if the extension agency wants to maximize adoption over the next 3-5 years, it would spread the entry points apart in the network to minimize repetition and redundancy in the same part of the network. Instead, if multiple connections are needed to encourage adoption, then the choice of the seed farmer becomes critical. Many potential seed pairings would yield no adoption at all, as no one would surpass the threshold for adoption. Centola and Macy (2007) label this case “complex contagion”. In this case, it is critical that multiple farmers are trained, and that these trained farmers are clustered together and share connections.

We first collect full social network census data on agricultural learning relationships in 200 villages in Malawi, and conduct simulations on those data to identify the theoretically optimal entry points (“seeds”) that would maximize adoption of a new technology, assuming the diffusion process is characterized by either simple contagion or complex contagion. Villages are then randomly assigned one of those targeting strategies, i.e. a version of the threshold model, and the Malawi extension services trained the seeds we selected via the theoretical simulations. Seeds were asked after training to disseminate the information. We then trace adoption patterns in these villages over the next 2-3 years. This exercise provides insights into the diffusion process, and also helps us examine whether network theory can be used to improve the effectiveness of public policy.

We compare the adoption via these “theoretically optimal” seeds against a benchmark treatment where agricultural extension agents use local knowledge to select seeds in another randomly selected set of villages. Typically, this involves asking village leaders to nominate a pair of extension partners and is similar to what many extension workers normally do outside of our study context. Interventions that rely on local institutions may use a great deal of valuable information in selecting influential people that is not available to contagion theory, such as their eagerness to try the new technology, their persuasiveness as communicators, or the trust other villagers place in their opinions. As such, this benchmark provides a demanding test for network-based diffusion theory:

our theoretically optimal partners were selected only by their position in the network, without the advantage of this additional local information.

As another comparison, we implement a fourth treatment in which we select optimal seeds who would maximize diffusion assuming geographic proximity proxies for social network connections. Unlike social network relationships, geographic location is easy for extension agents to observe, so we view this as a first step towards a policy-relevant alternative to the data intensive network-theory based approaches.

Banerjee et al (2013) highlight the role of entry points and document that in villages where leaders (shopkeepers, self-help group presidents, etc) occupy central positions in the village network, adoption of a microfinance product is higher. They focus on eigenvector and diffusion centrality, which are predicted to increase adoption under a wide range of diffusion models, including the model they estimate using non-experimental variation. Kim et al (2015) conduct an experiment to distribute public health coupons through high-degree or random individuals or nominated friends, but in a small sample of 32 villages. Our approach builds on this literature and tests whether diffusion models (in our case the threshold model) have *ex ante* predictive power: our target entry points were chosen in accordance with a particular diffusion model, and the theory is tested using a large-scale field experiment.

Our experiment focuses on the decision to adopt ‘pit planting’, a traditional West African technology which is largely unknown in Malawi. Pit planting has the potential to significantly improve maize yields in arid areas of rural Africa.³ Our test of the threshold diffusion model is therefore conducted in an important real-world setting with large implications for development.

³ It has been shown to increase productivity by 50-100% in tests conducted under controlled conditions (Haggblade and Tembo 2003); in large-sample field tests conducted under realistic “as implemented by government” conditions (BenYishay and Mobarak 2015), and using experimental variation among villagers in this study.

We find that the theory-driven targeting of optimal seed farmers leads to greater technology diffusion than the status-quo benchmark approach. Threshold theory based targeting increases adoption by 3 percentage points more than relying on extension workers to choose seeds, during the 3-year period of the experiment when pit planting adoption grew from 0% to about 10% in these villages overall.⁴

The complex contagion model suggests that one of the potential consequences of poor targeting is complete failure to adopt within the village. We observe no diffusion of pit planting in 45% of the ‘benchmark’ villages after 3 years. In villages where seeds were selected using the complex contagion model, there was a 56% greater likelihood that at least one person other than the seeds adopts in the village, relative to the benchmark. The results suggest that simply changing *who* is trained in a village on a technology on the basis of social network theory can increase the adoption of new technologies compared to the Ministry’s existing extension strategy. Even the low-cost geography-based targeting strategy generates some gains in adoption relative to the benchmark. However, physical proximity does not appear to be a good proxy for social connections. Developing other low-cost proxies for social network structure would be a useful avenue for future research.⁵

Finally, we structurally estimate the diffusion model using moments from the field experiment data in order to identify the distribution of the threshold parameter in our sample. Consistent with the reduced form results, our structural estimates suggest that a majority (65%) of people require multiple connections to adopt a new technology. The model that can best describe our data apparently lies in between the parameterizations for the threshold (1 for simple contagion and 2 for complex contagion) chosen in the design of our field experiment.

⁴ This rate of increase in adoption is not unusual for new agricultural technologies, including very profitable ones (e.g. Munshi 2007). Ryan and Gross (1943) show that it took 10 years for hybrid seed corn to be adopted in Iowa in the 1930s, and there was often 5 years between when a farmer heard of the technology and adopted it.

⁵ For example, promising results in Banerjee et al (2014) imply that households know who is central in their village, and this type of information may be easily elicited from a random sample of people. Kim et al (2015) use a related elicitation mechanism based on friends-of-friends.

Designing experimental treatments tightly linked to specific pre-formulated theoretical structures conveys several advantages. First, the use of the model *ex ante* in designing treatments commits us to testing a particular model. This eliminates the possibility of searching over potential theoretical models to *ex post* rationalize surprising (and possibly spurious) patterns in the data.

Second, as our treatment arms themselves incorporate the structure of the theoretical model, we can use these “structural experiments” to directly demonstrate the implications of simple or complex contagion models. Before showing empirical results, we simulate the model using network relationships data collected before the experiments are conducted, to precisely specify the pattern of estimates we should expect to observe *across all four treatment arms* under either simple or complex contagion theory. We present these simulated counterfactuals alongside our actual empirical results to allow readers to view what the experiment revealed about simple or complex contagion.

Third, while our approach is stylized, it may allow for a greater degree of external validity, because the theories we test were formulated independent of context. The results of other attempts in the literature to diffuse new technologies through social networks⁶ may be context-dependent because they rely on local institutions – such as local leaders or focus groups – to identify network partners. Carrel et al (2013) offer a cautionary tale for manipulating social diffusion in this way: they document that even among other classes at the same university, empirical patterns of peer effects do not create predictive power to design optimal classrooms.

The rest of the paper is organized as follows. We present the theoretical model on which the experimental design is based in Section 2. Section 3 discusses all field activities, including data collection and intervention implementation. Section 4 describes the characteristics and activities of the seed farmers and the performance of the technology in the field and documents that social

⁶ Kremer et al (2011) identify and recruit ‘ambassadors’ to promote water chlorination in rural Kenya, Miller and Mobarak (2014) first markets improved cookstoves to ‘opinion leaders’ in Bangladeshi villages before marketing to others, and BenYishay and Mobarak (2015) incentivize ‘lead farmers’ and ‘peer farmers’ to partner with agricultural extension officers in Malawi.

diffusion took place. Section 5 presents more detailed theoretical predictions and associated empirical results on the nature of contagion and network diffusion, while Section 6 presents structural estimates of the diffusion model. Section 7 concludes.

2. Theoretical model and experimental design

Our experiment takes place in 200 villages randomly sampled from three Malawian districts with largely semi-arid climates (Machinga, Mwanza, and Nkhotakota). Approximately 80% of Malawi's population lives in rural areas (World Bank 2011), and agricultural production in these areas is dominated by maize: more than 60% of the population's calorie consumption derives from maize, 97% of farmers grow maize, and over half of households grow no other crop (Lea and Hanmer 2009). Technology adoption and productivity in maize is thus closely tied to welfare.

The existing agricultural extension system in Malawi relies on Agricultural Extension Development Officers (AEDOs) who are employed by the Ministry of Agriculture and Food Security (MoAFS). Many AEDOs are responsible for upwards of 30-50 villages, which implies that direct contact with villagers is rare. According to the 2006/2007 Malawi National Agricultural and Livestock Census, only 18% of farmers participate in any type of extension activity. Against this backdrop of staff shortages, incorporating social learning in the diffusion process may be a cost-effective way to improve the effectiveness of extension.

2.1 Diffusion Models and Experimental Design

We use the linear threshold model (Granovetter 1978; Acemoglu *et al* 2011) as the theoretical basis for the dissemination strategies we develop in partnership with the Malawi Ministry of Agriculture. This model posits that an agent will adopt a new behavior once adoption among his connections crosses a threshold. The model was originally designed to study a wide array of collective behaviors including riots, voting, migration, and new technology adoption. The underlying

rationale for this formulation is either that the net benefit of adoption is a function of the number of neighbors who adopt (e.g. because neighbors' experiences with the technology can allow the farmer to refine his technique), or that farmers need to hear about the new technology from multiple sources before they are persuaded to adopt (when the threshold is above one).

The threshold model has served as a basic building block for diffusion theory. We designed field experiments around this model due to its influence in the literature (there are over 3000 citations to Granovetter 1978), and because the threshold formulation is more naturally micro-founded with a model of learning, relative to other canonical diffusion models.⁷ Further, the S-shaped diffusion pattern for new technologies noted by economists (e.g. Griliches 1957 for hybrid corn) is consistent with the number of contacts acting as a key driver of adoption decisions.⁸ However, our analysis will merely focus on understanding the process of diffusion using the experiments as a guide; we do not attempt to test the threshold model against other diffusion models in our empirical work. In fact, we show in Appendix A2 how specific parameterizations of the Susceptible-Infected model (another class of diffusion models widely discussed in the literature) are similar to the two parameterizations of the threshold model we selected for our experiments.

We employ two different versions of the threshold model in different arms of our experiment. The first version, called "simple contagion," postulates that the average individual needs to know only one other household who has adopted the technology in order to be convinced to adopt herself. Centola and Macy (2007) shows that some types of information – such as knowledge of job opportunities - spread through simple contagion. However, other behaviors may require

⁷ Both simple and complex contagion formulations are related to Bayesian learning models. In simple contagion, a single contact motivates adoption, suggesting that a person's prior (to not adopt) is not very strong. In contrast, complex contagion suggests that additional observations of adoption are necessary to move most people's priors. We opted for this simpler version over a formal Bayesian learning model as those models quickly become intractable in real world networks (Chandrasekhar et al 2012).

⁸ The slow rate of diffusion in early stages can be explained by not many people in a network having multiple contacts who have adopted when a technology is new, but the probability of having multiple adopter contacts increases more rapidly as the technology spreads through the network.

multiple sources of information before they are adopted, and we explore this using a complex contagion model in a second arm of our experiment.⁹

To select two “seed farmers” in each village for the experiments, we first collect social network relationships data (described in detail in section 3) on the census of households in each village before launching any field intervention activities. The social network structures observed in these data allow us to construct network adjacency matrices for each of the 200 villages in our sample. Next we conduct technology diffusion simulations for all villages using these matrices, where each individual in the village draws an adoption threshold τ from the data, which is normally distributed $N(\lambda, 0.5)$ ¹⁰ but truncated to be strictly positive. We conduct simulations with $\lambda=1$ and $\lambda=2$ in all villages to observe optimal seeds under simple and complex contagion respectively. In these simulations, when an individual is connected to at least τ individuals who adopted, he adopts in the next period. Once an individual adopts, we assume that all other household members also adopt, since agricultural plots are held at the household level in Malawi. We also assume that these adoption decisions are permanent with no dis-adoption. We run the model for four periods.¹¹

The final step to prepare for the experimental interventions is to choose the “optimal” partner farmers for each village as prescribed by the theory randomly assigned to that village. To accomplish this, we conduct simulations to predict the distribution of the village-level adoption rate after four seasons under the specified contagion theory, for every possible pair of individuals acting

⁹ Several theory papers have explored the implications of this model. In contrast to the “strength of weak ties” in labor markets proposed by Granovetter (1978), strong ties may be important for the diffusion of behaviors that require reinforcement from multiple peers. Centola (2010) provides experimental evidence that health behaviors diffuse more quickly through networks where links are clustered, consistent with complex contagion. Acemoglu et al (2011) highlights that when contagion is complex, highly clustered communities will need a seed placed in the community in order to induce adoption

¹⁰ We allow for heterogeneity in the threshold across individuals within the village. One motivation for this heterogeneity is that individuals are likely to differ in their baseline beliefs of the returns to pit planting, and their beliefs will influence how many signals they will need to observe before adopting themselves.

¹¹ We collected data for up to three agricultural seasons after the interventions were implemented. Note that once we have identified the value of λ , a policymaker could use the model to maximize adoption over any timeframe they cared about, either more short-term or more long-term.

as the two seed farmers. Given the randomness built in to the model, we simulate the model 2000 times for each pair, and create a measure of the average adoption rate induced by these two seeds. We ultimately select the pair that yields the highest average adoption rate in our simulations. This procedure allows us to specify the optimal pair of seeds for each village under each potential treatment arm.

There are multiple diffusion models in the theoretical literature that can generate close-to-equivalent empirical patterns of diffusion in real-world networks. This experimental design does not allow us to make quantitative statements about whether the underlying diffusion model is a threshold model or an alternate model that behaves similarly. To compare to other parts of the literature, we conduct a simulation exercise in appendix A.2 that suggests specific parameterizations of the Susceptible-Infected (SI) model, where a single connection to an adopter non-deterministically induces adoption with some probability, $\rho < 1$, yields predictions that are similar to simple or complex contagion. The simple contagion model is approximately equivalent to the SI model when $\rho \approx 1$. The complex contagion model, while theoretically distinct from any SI model¹², performs very similarly to an SI model with $\rho \approx 0.1$ in the simulations we conduct using our village network data from Malawi. Intuitively, these models behave similarly because if ρ is small, the probability of transmission from any one contact is very low, and multiple contacts (as in complex contagion) are in practice required for adoption.

Of course, separating theoretically distinct but empirically quite similar models has limited policy relevance: if farmers have either a high threshold or a very low transmission rate, it will be difficult to get the diffusion process started, which is consistent with the challenge policymakers have faced in getting farmers to adopt many agricultural technologies, and policymakers will need to be thoughtful in selecting who they work with in introducing new information. Both models create a

¹² We discuss in the appendix the scenarios under which the SI model could be distinguished from the threshold model.

diffusion process in which farmers need to be close to multiple seed farmers or early adopters before they adopt themselves.

2.2 Interventions

The two seed farmers in each village are trained in the targeted technologies by the Malawi Ministry of Agriculture extension staff. Our experimental variation only changes the theoretical basis on which the seed farmer is chosen in each village, and holds all other aspects of the training constant. Within each district, we randomly assign villages to one of the following four treatment arms (or seed farmer selection process)¹³:

1. Simple Contagion: Simple diffusion ($\lambda=1$) model applied to the network relationship data
2. Complex Contagion: Complex diffusion ($\lambda=2$) model applied to network relationship data
3. Geo Treatment: Complex diffusion ($\lambda=2$) model applied to an adjacency matrix where geographic proximity proxies for a network connection
4. Status Quo Benchmark: Extension worker selects the seed farmers

Treatment arms 1 and 2 were described above. In treatment arm 3, the simulation steps are the same as in the Complex Contagion case, except that we apply the procedure to a different adjacency matrix that is generated by making the assumption that two individuals are connected if their plots are located within 0.05 miles of each other in our geo-coded location data. We chose a radius of 0.05 miles because this characterization produces similar values for network degree measures in our villages as using the actual network connections measures.

The fourth group is the status-quo benchmark, where AEDOs were asked to select two seed farmers as they normally would in settings outside the experiment. Comparing the adoption performance of network theory-based targeting against this benchmark constitutes a meaningful and

¹³ Randomization was implemented using a re-randomization procedure which checked balance on the following covariates: percent of village using compost at baseline; percent village using fertilizer at baseline, and percent of village using pit planting at baseline. Randomization was implemented in each district separately.

challenging test for the simple and complex contagion treatments. In principle, the AEDOs could use valuable information not available to researchers, such as the individual’s motivation to take on the role, to select highly effective seed farmers. It is not clear *ex ante* that the theory-driven diffusion strategies would out-perform this benchmark.¹⁴ The benchmark treatment is what the Malawi Ministry of Agriculture and other policymakers would normally do, so this is the most relevant counterfactual.¹⁵

Note that the Simple, Complex, and Geo seed farmer selection strategies were simulated in all 200 villages, so we know – for example – who the optimal simple contagion seed farmers would have been in a village randomly assigned to the complex contagion or the geo treatment. We label the counterfactual optimal farmers as “shadow seeds.” This is useful for analysis: in any regression where we examine decisions made by the actual seed farmers, the shadow seeds form the relevant comparison group since this utilizes the random variation created by the experiment, holding constant the network position of the seeds and shadows. When we report effects on the broader village population, we exclude both the actual and the shadow seeds from the analysis.

2.3 Technologies

In this section we describe the two technologies introduced to seed farmers and in section 4.1 we analyze data on crop yields to give further insights into the benefits of the technologies.

Pit Planting

Maize farmers in Malawi traditionally plant in either flat land or after preparing ridges. Ridging has been shown to deplete soil fertility and decrease agricultural productivity over time

¹⁴ Another option would have been to randomly select seed farmers from the population, but that would have constituted a weaker test, and one with little real-world relevance as extension programs rarely randomly choose their partners.

¹⁵ Normally the Ministry only trains one “Lead farmer” per village, not two. In most villages, the Lead Farmer will already be established, except for villages in which there hasn’t been an extension officer assigned to the village for a long time. The AEDO will have had to select a second seed farmer in benchmark villages due to the experiment.

(Derpsch 2001, 2004). In contrast, pit planting involves planting seeds in a shallow pit in the ground, in order to retain greater moisture for the plant in an arid environment, while minimizing soil disturbance. The technique is practiced elsewhere in Africa, and has been shown to greatly enhance maize yields both in controlled trials and in field settings, with estimated gains of 50-113% in yields (Haggblade and Tembo 2003, BenYishay and Mobarak 2014). In section 4.1 we offer further evidence on yield impacts in our sample of villages. The enhanced productivity is thought to derive from three mechanisms: (1) reduced tillage of topsoil, which allows nutrients to remain fixed in the soil rather than eroding, (2) concentration of water around the plants, which aids in plant growth during poor rainfall conditions, and (3) improved fertilizer retention. The gains from the first mechanism over a counterfactual of continued ridging are thought to accumulate over time, while the gains from the second and third are expected to accrue even in the very short run.

Practicing pit planting may involve some additional costs. First, only a small portion of the surface is tilled with pit planting, and hand weeding or herbicide requirements may increase in principle, though focus groups suggested that weeding demands were reduced substantially relative to ridging. Second, digging pits is a labor-intensive task with potentially large up-front costs. However, land preparation becomes easier over time, since pits should be excavated in the same places each year, and estimates suggest that land preparation time falls by 50% within 5 years (Haggblade and Tembo 2003). BenYishay and Mobarak (2014) find that in Malawi, labor time decreases while the change in other input costs are negligible in comparison.

Crop Residue Management

Seed farmers were also trained in crop residue management (CRM), a set of farming practices which largely focus on retention of crop residues in fields for use as mulch. Alternative practices commonly used by farmers include burning the crop residues in the fields and removing

them for use as livestock feed and compost. The trainings emphasized the value of retaining crop residues as mulch to protect topsoil, reduce erosion, limit weed growth, and improve soil nutrient content and water retention. The trainings also addressed potential concerns about modifications in semi-arid areas (where there are fewer residues available), pest infestation, fire prevention, and alternative sources of livestock feed. There is little experimental evidence on the impacts of CRM on soil fertility, water retention, and yields in similar settings.

3. Field Activities: Implementation of Interventions and Data Collection

3.1. Training of Seed Farmers

After we produced the lists of seed farmers for each village using the procedures described above, the AEDO assigned to the village trained the two seed farmers.¹⁶ We provided AEDOs with two seed farmer names for each village in experimental arms 1-3, and then replacement names if either of the first two refused to participate. Refusal was uncommon: we trained 93% of the selected seeds or their spouses. We conduct intent-to-treat analysis using the original seed assignment. The trainings took place in April-May of 2011 for Machinga and Mwanza districts, and March-April of 2012 for Nkhotakota. Following the training of seed farmers by AEDOs, all seed farmers were also informed that they would receive a small in-kind gift (valued at US\$8) if they themselves adopted pit planting in the first year (and that the gift would be given only in the first year). The gift was given at the time of follow up data collection and verified on the farm by the enumerator. There was no gift or incentive offered or provided on the basis of others' adoption in the village.

3.2 Data

The interventions were designed on the basis of social network census data collected from all sample villages at baseline. After training the seed farmers, we collected up to three rounds of longer

¹⁶ As the technologies themselves were new, the AEDOs were themselves trained by staff from the Ministry's Department of Land Conservation.

household survey data for sub-samples of the village populations. Appendix Figure A1 shows the timeline of these data collection activities. We describe each major data source in turn.

Social Network Census Data

Targeting based on different network characteristics—including relational statistics of these networks—requires relatively complete information on network relationships within the village (Chandrasekhar and Lewis 2011). To collect this data, our field teams completed a social network census within each village, attempting to interview one man and one woman in each household. In practice, we reached more than 80% of households participating in the census in every sample village.¹⁷

The main focus of the social network census was to elicit the names of people each respondent consults when making agricultural decisions. General information on household composition, socioeconomic characteristics of the household, general agriculture information, and work group membership was also collected. The individual questionnaires asked about agricultural contacts several ways: first by asking in general terms about farmers with whom they discuss agriculture. To probe more deeply, we also asked them to recall over the last five years if they had: (i) changed planting practices; (ii) tried a new variety of seed, for any crop; (iii) tried a new way of composting; (iv) changed the amount of fertilizer being used for any crop; (v) tried a new crop, such as paprika, tobacco, soya, cotton, or sugar cane; or (vi) started using some other new agricultural technology. If they responded affirmatively, we asked respondents to name individuals they knew had previously used the technique in the past and whether they had consulted these individuals. Finally we asked them if they discussed farming with any relatives, fellow church or mosque members, or farmers whose fields they pass by on a regular basis. We also elicited contacts with

¹⁷ We interviewed at least one household member from 89.1% of households in Nkhotakota, 81.4% in Mwanza and 88.6% in Machinga. We interviewed both a man and a woman in about 30% of households.

whom they share food and close friends. These responses were matched to the village listing to identify links. Individuals are considered linked if either party named each other (undirected graph), and all individuals within a household are considered linked.

Sample Household Survey Data

We collected survey data on farming techniques, input use, yields, assets, and other characteristics for a sample of approximately 5,600 households in the 200 sample villages. We attempted to survey all seed and shadow farmers in each village, as well as a random sample of 24 other individuals, for a total of 30 households in each village.¹⁸ In villages with fewer than 30 households, all households were surveyed. Three survey rounds were conducted in Machinga and Mwanza in October-December of 2011, 2012 and 2013. In Nkhotakota, two survey rounds were conducted in October-December of 2012 and 2013.¹⁹ The initial rounds referenced agricultural production in the preceding year—thus capturing some baseline characteristics—as well as current knowledge of the technologies, which could reflect the effects of training. Since the data was collected at the start of a given agricultural season, we observe 3 adoption decisions for pit planting for farmers in Mwanza and Machinga, and 2 decisions for farmers in Nkhotakota. Since crop residue management (CRM) decisions are made the end of an agricultural season after harvest, we observe CRM decisions for two agricultural seasons in Mwanza and Machinga, and one in Nkhotakota.

Randomization and Balance

Appendix Table A1 shows how observable characteristics from the social network census vary with the treatment status of the village. The table shows the results of a regression of the dependent variable listed in the column heading on indicators for the respondent residing in a benchmark,

¹⁸ In Simple, Complex and Geo villages there were 6 (2x3) seed and shadow farmers to interview, while in Benchmark villages there were 8 (2x4) seeds and shadows.

¹⁹ Unanticipated delays in project funding required us to start training of AEDOs and seed farmers in Nkhotakota in 2012 instead of 2011 as in Mwanza and Machinga.

simple, complex, or geo treatment villages. District fixed effects are included in the regression, and standard errors clustered at the village level. P-values from statistical tests comparing across the different treatment groups as well as a joint test of all treatment groups are displayed. Given the large number of comparisons made in Table A1, few differences across treatment groups are statistically significant. Farm size, in column (9), is the most concerning: farmers in the benchmark villages have larger farm sizes on average than farmers in Simple and Complex villages, and the joint test across the treatment variable is significant at the 10% level.

4. Seed Farmers: Characteristics, Adoption, Communication, and Diffusion

In this section, we provide some background information from our data that shows the field experiment “worked” as intended, which helps set the stage for the diffusion analysis in Section 5. First, we show that the technologies we promoted improved agricultural yields, so that the diffusion results presented later can be interpreted as learning about a positive attribute. We then describe empirically the baseline characteristics of the seeds associated with each treatment arm, and the rates at which the seeds adopt the new technologies themselves. Next, we show that the seed farmers disseminate information on pit planting within the village. Finally, we show that individuals close to the trained seeds are more likely to adopt, so the individual-level adoption patterns appear to show that diffusion takes place.

4.1. Effect of Technology Adoption on Crop Yields

We collected data on maize yields in our follow-up surveys, and we use this to show in Appendix Table A2 that the technologies we promoted led to an increase in output. This allows us to establish that the information about pit planting that diffused through the networks was likely positive on average.

We compare seed farmers to shadow farmers to study yield effects, exploiting the randomization in the experimental design.²⁰ In an intent-to-treat specification, maize yields among seed farmers (who were both trained on the technologies and promised a small reward to adopt) are 13% greater than the yields experienced by the comparable shadows. We report the local average treatment effect using an IV regression in the second column in which we instrument pit planting adoption with an indicator for being randomly assigned the role of actual seed farmer who was trained and incentivized to adopt (rather than a shadow). In this specification, pit planting adoption is associated with a 44% increase in maize yield. However, we cannot rule out that CRM adoption also increased yields, potentially violating the exclusion restriction in the IV estimation.

4.2 Characteristics of the Seed Farmers under each Treatment

The simulations of the simple and complex contagion models generated different optimal seeds in most but not all cases. In 50% of villages, there was at least 1 seed who was judged as optimal in more than one (simple, complex or geo) models. Appendix table A3 describes the frequency of overlap in seeds across treatments. The most common scenario is that a simple seed is also a complex seed, which happens about 25% of the time. Even though the extension workers could have chosen central individuals, benchmark seeds are also simple seeds only 9% of the time and also complex seeds only 11%. Even more striking (not shown in table A3) is that in cases where an individual is both a simple and a complex seed, i.e. they are likely to be a very obvious and central person in the village, (s)he is chosen as a benchmark seed only 11% of the time. The least overlap is between the Geo seeds with all others. In cases with overlap in the seeds, the treatment arms are naturally less distinguishable. Nevertheless, the treatment arms generated different types of seed

²⁰ Whenever we report comparisons between seeds and shadows, we restrict analysis to Simple, Complex and Geo treatment villages, because both seeds and shadows were identified via simulations in those villages and we don't observe Benchmark shadow farmers.

farmers in general, as discussed below. They also generated different clustering patterns. For example, 35% of our random household sample has a connection to a simple seed, and 6% are connected to both simple seeds. However, 18% of households are connected to two complex seeds and 28% are connected to one complex.²¹

Table 1 compares the seed farmers chosen under the four different targeting strategies in terms of observable characteristics such as wealth and land size from our survey data, and in terms of centrality measures computed from our social network census data. The table includes all seeds and shadow farmers, irrespective of treatment status. Note that there are only 100 benchmark farmers since we never observe shadow benchmark farmers in non-benchmark villages. The most striking pattern in Table 1 is that the seeds selected under the geographic treatment are much poorer than other seeds. This is because many households live on one of their plots in Malawi. Households who are geographically close to lots of people will mechanically have less land, and these households tend to be poorer overall. Therefore while the idea of using geography as a proxy for one's network may be intuitive, the implications of geographic centrality may be highly context-specific.

Seed farmers selected through the complex contagion simulations are the most “central” across all measures of network centrality we compute (columns 3-5). Seed farmers in the complex contagion villages have three to four (30%) more direct connections to others in the village than the seed farmers chosen by the extension workers. Seeds in complex contagion villages also possess the highest between-ness and eigenvector centrality measures, which imply that they are important nodes in these villages.²² Simple seeds have similar between-ness centrality as complex seeds, but lower eigenvector centrality.

²¹ For the geo-based seeds, 20% of households are connected to one, 10% connected to two and a larger fraction than in either network theory-based treatments – 70% – are connected to no seed.

²² Eigenvector Centrality is weighted sum of connections, where each connection's weight is determined by its own eigenvector centrality (like Google pagerank). Betweenness centrality captures that a person is important if one has to go through him to connect to other people. Therefore it is calculated as the fraction of shortest paths between individuals in the network that passes through that individual. See Jackson (2008) for more details.

Figure 1 shows five example villages from our data with network links mapped and the locations of the simple, complex and geo seeds within the village social networks. One feature common across these villages is that the simple seeds tend to be more distant from one another than do the complex or geo seeds. In village 45, for example, one central household was chosen as a seed in both Simple and Complex models, but the selection of the second farmer reveals the main difference between these models. In complex contagion, the second seed farmer is directly connected to the first seed and is also quite central in the network. This is a necessary condition for a cascade under complex contagion: without shared connections, no one would adopt at all. The second simple seed, however, is far more removed from the giant component in the network. Under simple contagion, training the first seed is sufficient to induce the diffusion process to occur within the main cluster in the village, and the second - more removed farmer - was otherwise unlikely to adopt without being directly targeted.²³

All the example villages in Figure 1 show that the geo seeds are generally close to one another. This is because the underlying diffusion model selecting these sides is complex contagion, which would ensure that they live near each other. It appears there is some correlation in the geo and social network maps since these neighbors also show up as socially connected in Figure 1. However, they are located in more peripheral locations within the network, as anticipated given the summary statistics in Table 1, since the geo-based measurement of the network misses key structural facts about the network. Figure 2 shows four example villages which also include Benchmark seeds. As Table 1 suggests, Benchmark farmers are more central in the network than Geo farmers, but less central than Complex farmers. Most importantly, they are rarely sufficiently clustered in the network to spark the diffusion process if decisions are governed by the complex contagion model.

²³ The second simple farmer would be more central in a village which had multiple distinct cliques. However, we rarely observe this network structure in our data. This is discussed further in the appendix.

4.3 Do Seed Farmers Adopt the Technology Themselves?

Table 2 shows that seed farmers assigned to treatment in the experiment were more likely to adopt the technologies. The sample is restricted to seed and shadow farmers only and excludes benchmark villages (as we cannot identify the counterfactual benchmark farmers), so this specification captures the causal effect of the intervention, and not differences in adoption across farmers at different positions within the network. Columns (1)-(3) focus on pit planting and columns (4)-(5) on crop residue management. Seed farmers who are trained on pit planting adopt at a similar rate (31%) in all three years. This represents an 18-26 percentage point improvement over comparable shadow farmers across all three years, as the adoption rate among shadows was 5% in year one, rising to 14% by year 3. We provided an in-kind incentive for the seed to adopt pit planting in the first year but not thereafter. We never provided the seeds any incentive to adopt CRM, but the trained farmers were also 14 percentage points more likely to try CRM in the first year after training. The rate for shadow farmers was high to begin with (32%), so managing crop residues was not as new or unfamiliar a concept as digging pits.²⁴ However, CRM adoption declined quickly in the second year among both actual seeds and the shadows (from 46% to 26% for seeds and from 32% to 21% for shadows), which is strong evidence that it was not deemed as useful a technology as pit planting.

Given that CRM was already a familiar technology, and that there were high rates of dis-adoption, our threshold diffusion model does not appear to be the appropriate tool to analyze social

²⁴ While pit planting is a new, largely unknown technology in Malawi (0.5% of farmers are practicing at baseline), farmers were using a range of strategies to deal with crop residues, including burning fields, leaving residues in fields, using residues as mulch, feeding residues to livestock, using residues to make compost, and, most commonly, burying residues in fields as they prepare new ridges. Several of these strategies overlap with the recommendations provided in our CRM training, creating measurement problems. Unlike pit planting, which is readily observable and distinct from other practices, whether farmers follow our CRM guidelines is not as easy to decipher in our data. Further, the optimal crop residue technique depends on household-specific factors like livestock owned, and farmers and extension experts disagree about the best practice. Ministry officials do agree that burning residue is a bad idea, and we observe burning frequency decreasing from 20% to 9% over the 3-year study period.

learning about CRM. If the technology was not suitable, the message “do not adopt” was likely passed within the network, in which case, adoption propensity is not the right outcome variable to study. Our estimates indicate that, if anything, our treatments reduce adoption of CRM by year 2, and the threshold model is not very useful to interpret this dynamic. In contrast, pit planting is a new and unknown technology for which information constraints were probably more relevant. Pit planting adoption among those trained was also persistent, which suggests that the seed farmers found the method useful. We therefore focus on adoption results only for pit planting, and present CRM adoption results in Appendix tables A5 and A6.

Table 3 restricts the sample to only seed farmers who were trained (and drops all shadows) to examine whether adoption behavior varies across the seeds in the four experimental arms. In the first year, there are no differences in adoption propensities (or in the likelihood of recalling the existence of the technology) across the four types of seeds. This implies that it is unlikely that any observed differences in village-wide adoption patterns are driven by initial differences in seed behavior. Columns (2) and (3) show that seed farmers in simple contagion villages become relatively more likely over time to adopt the technology, which may be an outcome of the experiment. Their adoption propensity is significantly higher than the AEDO-selected seeds in years 2 and 3, which is striking because AEDOs could have screened partner farmers based on their interest in using the new technology. Note that differences in years 2 and 3 could also be an outcome of the experiment, which limits our ability to interpret these differences. Columns (4)-(5) show that there are no significant differences in adoption in seasons 1 or 2 for crop residue management.

4.4 Seeds Farmers’ interactions with other villagers

In this section we examine whether seed farmers disseminate information about pit planting to their neighbors in the village, a likely necessary step to jumpstart a diffusion process. Table 4 uses data

collected in the first follow-up survey in 2011-12 on conversations about pit planting that respondents had with others in the village. Each respondent was asked questions about seven other individuals in their village, whether they knew them, and what they had discussed. The seven individuals comprised of the two seed farmers, some randomly selected shadow farmers, and a random sample of other village residents. We exploit the random variation in the experiment, and compare conversations with the (say) complex farmers who were assigned the role of seed farmer by our intervention, to communication with the complex shadows in other villages who are observably similar, and who *would have been* the seed had those comparison villages been assigned to the complex contagion treatment.²⁷ In other words, we test whether a potential seed *being trained* on pit planting increases the likelihood that he talks to others about pit planting.

Table 4 shows that the experiment did induce the seed farmers to discuss pit planting with fellow villagers. Column (1) shows that there are more discussions with the “simple seed” in both Simple and Complex treatment villages compared to the benchmark villages. As expected, the effect is significantly larger in Simple treatment villages (4.7 percentage points) than in Complex treatment villages (2.0 percentage points), and these represent large increases over the mean value (1.9 percentage points) in the benchmark villages. We observe a treatment effect even in Complex villages because, as mentioned above, there is considerable overlap in the optimal seeds chosen through the complex contagion and the simple contagion simulations.²⁸ Columns (2) and (3) show, analogously, increases in conversations about pit planting with the complex farmer in Complex treatment villages (a 3.6 percentage point increase compared to benchmark) and with the geo farmer

²⁷ While all sample respondents in Simple treatment villages were asked about simple farmers, not all respondents in the remaining villages were, since we chose a random subset of shadow farmers. This is analogously true for complex and geo villages. We therefore flexibly control for the number of simple (complex, geo) farmers we asked about in the regression where the dependent variable is talking about pit planting with the simple (complex, geo) farmer.

²⁸ If we exclude from this regression villages where there is overlap in the optimal farmers, we observe an increase in conversations with Simple farmers only in the Simple contagion treatment villages.

in Geo villages (3.1 percentage points). In summary, the seed farmers trained in the pit planting method discussed the technology with others in their villages as a result of our experiment.

4.5 Technology diffusion within the village

If adoption is a social contagion, individuals close to the seeds should be the first who are persuaded to adopt before information percolates through the rest of the network. To explore this, we estimate the coefficients β_1 and β_2 in the following equation:

$$Y_{ij} = \alpha + \beta_1 1TSeeds + \beta_2 2TSeeds + \beta_3 1Simple + \beta_4 2Simple + \beta_5 1Complex + \beta_6 2Complex + \beta_7 1Geo + \beta_8 2Geo + \theta_j + \varepsilon_{ijt}$$

1TSeeds is an indicator for the respondent being directly connected to a *trained* seed farmer (assigned in the experiment) and *2TSeeds* indicates the respondent was directly connected to two *trained* seed farmers. Since network position is endogenous, we also control for whether an individual is connected to one or two Simple, Complex or Geo (actual or shadow) seeds.²⁹ We therefore use shadows to control for the respondent’s network position, and only using variation in closeness to the *trained* seed generated by the experiment for our identification. To illustrate, we compare, say, two farmers who are both connected to exactly two ‘Simple seeds’, but where one farmer is in a village randomly assigned to the “simple contagion” treatment (so that his connections were actually trained on the technology), while the other was not.

Table 5 shows these results for each of two outcomes: awareness of pit planting and adopted pit planting. In season 1, the training led to more information transmission to those directly connected to seeds, and in particular, those who have a direct connection to both the seed farmers who were trained on the technologies. Respondents with two connections are 7.1 percentage points more likely to have heard of pit planting than those with no connection to a seed. This represents a

²⁹ For example, *1Simple* indicates that the respondent is directly connected to one simple seed while *2Simple* says that the respondent has connections to two simple seeds. *1Complex*, *2Complex*, *1Geo* and *2Geo* are analogously defined for complex seeds and geo seeds. Coefficients on these control variables (β_3 through β_8) are not reported in the table.

33% increase in knowledge relative to the mean familiarity among unconnected individuals. This effect is also statistically significantly different from the effect of being connected to one seed ($p=.02$). The increased awareness of pit planting among households connected to two seeds persists into season 2 (column 4), and two connections is again significantly more advantageous than one connection ($p=.05$). However, we see no effect in column (2) of the information targeting on adoption among individuals directly connected to either one or two seeds, relative to those with no connections.³⁰

This comparison between connections with one versus 2 seeds is informative, because the simple and complex contagion models differ sharply in what sort of pattern they suggest across these two parameters: If everyone behaves as though adoption is a simple contagion, we should expect having two connections to seeds to be no more effective than having a single connection. In contrast, the complex contagion model suggests that in the first period, only people with multiple connections to seeds would be encouraged to adopt. Taking these sharp theoretical predictions to the data requires some care: measurement error in our exposure variable could bias coefficients in either direction. Nonetheless, we interpret these results as providing suggestive support that learning has complex contagion attributes.

The information effect in year 1 translates into an adoption effect in year 2. Column (5) shows that households with two connections to trained seeds are 4.2 percentage points more likely to adopt in the second season than those with no connections, which represents a 76% increase in adoption propensity. Though the point estimate of the effect of 2 connections is considerably larger than the effect of a connection to one seed (4.2 pp compared to 1.5 pp), we cannot statistically reject that households with a connection to only one treated seed adopt less frequently ($p=.17$).

³⁰ The control variables show that individuals with certain positions – such as those with one direct connection to a simple seed – are more like to adopt and hear of pit planting even when that seed is not trained on the new technologies. This highlights the importance of using the variation induced by the experiment since unobserved factors, related to one's position in the network or characteristics correlated with it, also affect adoption.

By season 3, however, we no longer see differences in either adoption or knowledge. This may be because the diffusion process has progressed to individuals further from the seeds by the third year. Looking at the means, we observe that both the adoption rate and awareness of pit planting has increased among individuals with no direct contacts (to 6.4% and 4.3% respectively), thus eroding the difference between direct and indirect contacts as information spreads further out from the seeds over time. To look at this more formally, columns 3, 6 and 9 also include an indicator for whether the household is within path length 2 of a seed (that is, a friend of a friend) as well as controls for being within path length 2 of each type of potential seed. We see that initially in year 1, there is no difference in adoption between those within path length 2 compared to those who are more distant to seeds. By season 3, column (9) shows that there is no longer a difference between households directly connected to seeds compared to those who are within path length 2, and households within path length 2 are more likely (3.7 pp) to have adopted over those who are socially more distant.

In summary, analysis using individual-level data demonstrates that individuals who are initially close to the trained seeds are more likely to adopt than individuals with no direct connections – as one would expect if the experiment is inducing social network-based diffusion. The data also suggest that having two direct connections – and not just one – is important for diffusion, as would be the case in the complex contagion version of the threshold model.

5. Village-Level Experimental Results: Does Theory-based Targeting increase Adoption?

In this section, we report aggregate village-level outcomes in the four treatments we implemented, and compare those against predictions from simple and complex contagion theory. This allows us to

both (a) test whether threshold diffusion theory can improve targeting and technology diffusion, and (b) characterize the underlying diffusion process through the lens of contagion theory.

5.1 The Advent of Diffusion under Simple and Complex Contagion

Our approach to analysis is to first conduct adoption simulations, to generate predictions for all 200 sample villages (regardless of treatment status) assuming that simple contagion is the right model. This is useful, because the simple contagion model has specific predictions how our seed choices in complex, geo or benchmark villages will behave, given that we know the network positions of the seeds chosen in those villages. Similarly, we can also simulate the adoption patterns in all villages, assuming that the learning environment is complex. We then compare our data to these simulated predictions.

One key feature of the threshold formulation that helps distinguish complex from simple contagion is that almost any pair of partners will guarantee a start to the diffusion process under simple contagion, but if the world is complex, many potential pairs of seeds would never generate any social diffusion. This is because agents need to cross a higher threshold ($\lambda > 1$) if the world is complex, and under many network structures, almost no individual may be connected to the required number of seeds. Productive technologies may die quickly and never diffuse if the learning environment is complex. We therefore focus on the advent of diffusion in our sample villages as a key outcome, using the variable “any adoption”, which is an indicator for villages which have at least one household (other than the seeds) that adopted pit planting.

The left part of Figure 3 shows the *predicted* fraction of villages with “Any Adoption” from simulating the model for all sample villages when $\lambda=1$ (Simple contagion) and $\lambda=2$ (Complex

contagion) for years 2 and 3.³¹ Since the goal is to compare these simulations to the actual data, we design the simulations to reflect the fact that we only observe a random sample of households in these villages.³² The right part of Figure 3 shows the empirical counterpart: the actual (observed) values for “any adoption” in years 2 and 3.

When the threshold is set to 1 on average in the simulations (i.e. assuming simple contagion), diffusion is predicted to be widespread: in year 2, 85% of villages where Geo and Benchmark partners were trained are predicted to have some measured adoption, and that rate goes up to 94% with Simple and Complex partners. The predicted rates of ‘any adoption’ are even higher in year 3.

The risk of no adoption increases in the complex contagion simulation where we increase the (median) threshold from one to two. If the learning environment is complex, then the model predicts that more than half the villages assigned to Simple, Geo or Benchmark partners will not see any adoption at all in year 2. In contrast, when complex seeds are trained, 70% of villages are predicted to experience some adoption. The fraction of villages with “no adoption” is predicted to be dramatically lower (16.5%) in the villages randomly assigned to the Complex treatment, relative to the Simple (53%) and Geo (45%) villages, respectively.

The right side of Figure 3 shows the actual fractions of villages with “any adoption” during years 2 and 3 in our data under each experimental arm. The data appears to match the shares of villages with any adoption simulated under complex contagion much more closely than those generated under simple contagion, in three distinct ways. First, the simple contagion simulations

³¹ These simulations exclude 12 villages where at least one of the extension worker chosen seeds (benchmark) was not observed in our social network census. This occurred because the spatial boundaries of villages are not always clearly delineated in Nkhotakota.

³² The simulations use the full social network (that we observe) to predict adoption. We then sample from the full network to better mimic our data. In the model, the rate of any adoption is identical in years 2 and years 3. If there was no adoption by year 2, there is no way there will be any additional adoption taking place in year 3. The sampling process, however, generates the increase over time observed in the figure. If the rate of adoption is low, as is empirically the case as well, then a random sample may miss all adopters. As the number of adopters increases over time (only in villages which are experiencing diffusion, so holding the extensive margin constant), the random sample is more likely to pick up an adopter and hence the rate of any adoption increases over time in the figure.

suggest that we should observe a much higher fraction of villages with some adoption than is true in the data across treatment groups. Second, simple contagion predicts that the “any adoption” outcome should not be very sensitive to the identity of the seed farmer who is initially trained. In contrast, the identity of the seed farmer dramatically alters this outcome in the data. Finally, the complex contagion simulation predicts that the complex partners will maximize the fraction of villages with some adoption, which is exactly what we observe in the data. In summary, the village-level adoption patterns look more consistent with simulations that assumed that the learning environment is complex.

5.2 Adoption Rates across Treatment Arms

We now document treatment effects for a broader set of adoption outcomes. Table 6 displays regression results based on empirical data from our experiments. Column (1) shows that both simple and complex contagion villages have higher adoption rates relative to the benchmark as of season 2. Compared to the benchmark rate of 3.8%, complex villages experience 3.6 percentage point higher adoption rates and simple villages experience 3.6 percentage points. We cannot reject that the adoption rates are the same in Simple and Complex villages. The adoption rate in Geo treatment is very similar to that in Complex and Simple villages. Column (2) looks at the adoption rate in season 3, and the adoption rate increases across all four types of villages (adoption increases in the benchmark villages, the omitted category, from 3.6% to 7.5% from season 2 to 3). With the smaller sample size, we cannot reject that the adoption rate is the same across all treatment types, though the point estimate on Complex remains the largest.

Columns (3) and (4) examine differences in the number of adopters across the treatment arms. This specification uses sampling weights to account for the fact that we sampled the same number of respondents irrespective of village size. In both season 2 and 3, there are on average an

additional 2 to 3 adopters in Complex villages relative to the benchmark. This represents a 140% increase in the number of adopters in season 2, and this difference is statistically significant. Neither Simple nor Geo villages are statistically different from the Benchmark villages in either season, but qualitatively we observe the point estimates in both treatment groups becoming smaller (relative to Benchmark) from season 2 to season 3. In season 3, the number of adopters in Complex villages is statistically higher than in Geo villages.

Finally, columns (5)-(6) replicate figure 3 in a regression framework and examine the extensive margin – whether anyone in the village sample adopted pit planting. In season 2, this rate is significantly higher in Complex villages relative to Benchmark villages, but not relative to the other two treatment groups. The point estimate for Simple villages is 0.155, compared to 0.252 for Complex. These effects are large compared to a Benchmark any adoption rate of 42%. In season 3, Simple, Complex and Geo villages all attain a higher rate of “any adoption” than Benchmark villages. 85% of Complex villages had at least one non-seed adopter, compared to 73% of Simple and Geo villages and 54% of Benchmark villages.

To interpret these empirical results we report in Table 6 using the lens of our model, we run the analogous regressions in appendix Table A4 on simulated data that applies the theory to the network adjacency matrices. The simulations we conduct predict that the complex treatment should perform best in terms of all three outcomes (adoption rate, number of adopters and advent of diffusion) if the learning environment is complex. If the learning environment is instead Simple, then we should expect to see few statistical differences in adoption across targeting strategies by season 3, since the choice of seed partners is relatively less important under our simple contagion treatment. The empirical results appear more consistent with a complex learning environment, but our parameterization of $\lambda=2$ does not appear to provide a perfect fit for the data.

5.3 Adoption Rates across Treatments when the Technology is Novel

Pit planting is in general a new technology in Malawi, although there is some heterogeneity within our sample in how novel it is. It seems reasonable to expect that the threshold for adoption may be higher when the technology is more novel, and people have had less opportunity to learn about its benefits. Table 7 shows the results of re-estimating the program evaluation regression from Table 6, but for the subset of villages that were least familiar with the new technologies at baseline (where the proportion of villagers who had ever heard of pit planting was below the median value of 4.3%). These are the villages where lack of information is more likely to be a deterrent to adoption, and thus where our models are most applicable. The Complex treatment performs relatively better in this sub-sample of villages. Moreover, the adoption rate is statistically differentiable from both the Simple contagion and benchmark treatment in year 3. In the sub-sample, the empirical results closely align with the predictions from simulations that assumed a complex learning environment. This suggests that the complex contagion model may be a good description of diffusion behavior when technologies are truly novel.

5.4 Is geographic targeting a viable alternative to data-intensive methods?

The empirical results suggest that targeting seeds using network theory (the threshold model in particular) can improve on technology diffusion relative to the status quo design of agricultural extension services. However, achieving these gains in adoption is not necessarily cost effective: the theory-based procedure is data intensive, and eliciting social network connections in each village is expensive. The geography-based treatment arm serves as a first step towards assessing how much of the benefit derived from applying theory could be appropriated without having to resort to expensive data collection methods. Using households' physical location within a village, which is

more easily observable to policy-makers and marketing agents than network relationships, would be less expensive and scalable.

Table 6 suggests that though the geographic proxy is generally not as effective as the network theory based treatments, it does generate some adoption gains relative to the Benchmark. By season 3, the Geo treatment villages exhibit a statistically significant 19 percentage point increase in the probability that any non-seed farmer adopts the technology, relative to benchmark villages. However, the Geo treatment does not perform as well as the data-intensive Complex treatment in generating a larger *number* of non-seed adopters. The point estimate in season 3 is considerably smaller than in the Complex villages.

Table 1 provides some insights into potential reasons for these differences. The Geo seed farmers are on average much poorer, and are often more remote in terms of their network position (as indicated by lower eigenvector centrality values in Table 1 and illustrated in figures 1 and 2). The two Geo seeds in each village are generally clustered together, since their selection process employed a simulation based on the complex contagion model. Thus theoretically we would expect some diffusion to their geographic neighbors, increasing the extensive margin of adoption as observed empirically. However, since these seeds are less connected and in a less dense part of the network than are the simple, complex and benchmark seeds, there would be a slower pace of diffusion.

Overall, we conclude that there remains a need to develop other simple and inexpensive procedures that can identify individuals who our social network data (combined with theory) chose as seed farmers in order to make network-based targeting more policy relevant and scalable. Recent evidence indicates that less expensive approaches may be feasible. For example, Banerjee et al (2014) have shown that in India a simple question like “if we want to spread information about a new loan product to everyone in your village, to whom do you suggest we speak?” is successful in identifying individuals with high eigenvector centrality and diffusion centrality.

6. Structural Estimation of the Threshold Model

The results in tables 5-7 suggest that the data are more consistent with the predictions generated from complex rather than simple contagion theory, though neither model fits the data perfectly. The two specific parameterizations we chose in designing our treatments ($\lambda=1$ and $\lambda=2$) may not provide the best fit; a different distribution for the threshold parameter may be able to better match patterns in the data. We therefore structurally estimate the model, and search for the distribution of the adoption thresholds (τ) that best matches certain moments in the data. We only use variation generated by the experiment to conduct this Method of Simulated Moments estimation, and focus on adoption rates and the “any adoption” measure under the four randomly assigned treatments as the moments to match.

We begin this exercise by extending the model to allow for the possibility that a person is informed about the technology, and thereby passing his threshold, but chooses not to adopt (following Banerjee et al, 2013).³³ We observe many seed farmers, for example, who do not adopt pit planting. We treat the probability of adoption conditional on being informed as an auxiliary second stage. The diffusion process is therefore about information and not about adoption directly. If farmers become informed, by crossing their threshold, we assume they adopt with some fixed probability, determined by the average probability with which seeds adopt (32%).³⁴

Using this framework, we estimate the distribution of τ . We focus on a vector of 8 moments: whether anyone adopted and the mean adoption rate in villages which were randomly selected to receive the complex treatment, the simple treatment, the geographic treatment and the benchmark treatment. For simplicity, we focus on the year 2 adoption data and assume that

³³ Not accounting for this may lead to an overestimate of the mean λ as many people may be informed about the new technology but choose to not adopt, suggesting that transmission thresholds (for being informed) may be lower.

³⁴ In our data, few observable characteristics significantly predict adoption among the seeds so we cannot follow Banerjee et al (2013) and use a logit with observable characteristics to predict adoption.

information has passed through the network once by year 2. Since each village only has two seeds, our identifying variation comes off of differences in expected effects of targeting depending on the fraction of people with $\tau = 1$ and the fraction of people with $\tau = 2$, but we have little variation to identify the fraction of people with values of $\tau > 2$. We therefore estimate the distribution of τ under two assumptions: (i) fraction q_1 has $\tau = 1$ and q_2 has $\tau = 2$, or (ii) $\tau \sim N(\lambda, \sigma)$ truncated to be positive (following modeling assumptions). This process has several steps. For each treatment t of complex, simple, control, and geo:

1. We discretize (λ, σ) and (q_1, q_2) into a grid of potential values
2. For every possible distribution (λ, σ) and (q_1, q_2) in the grid, we draw a vector of thresholds for each individual in the village, and a realization of her adoption decision³⁵
3. We calculate the moment in the village for that realization of threshold and adoption decisions assuming that the seeds in the village were selected according to treatment t
4. We repeat this for 2000 realizations of the thresholds and adoption decisions
5. We then take the mean across villages and realizations of the simulated moment

We calculate the empirical moments for each treatment by calculating the mean outcome variable across the subset of villages assigned to that particular treatment. We then match the vector of 8 simulated moments to the 8 empirical ones. Since villages were randomly assigned to treatments, the villages assigned to that particular treatment should be a consistent estimate of these theoretical moments defined across our entire sample. In doing so, there are only 2 types of variation used to identify the underlying distribution: the mean of each of these outcome variables, across treatments, and the differences in outcome variables induced by the random assignment to treatment.

³⁵ The adoption decision is based on a realization given an individual i 's predicted probability of adopting given the sample-wide seed adoption rate.

Both the non-parametric and parametric estimates, shown in Table 8, suggest that a 28%-34% of people are “simple” types, with a threshold of 1, and that the remaining 66-72% of people are “complex” types, with a threshold larger than 1. If we further impose that the threshold is 3 if the threshold is greater than 2 in the non-parametric case, we can estimate the expected value of τ , as well. Once again, both distributional assumptions generate similar estimates, suggesting that the average value for the threshold is about 1.74-1.80.³⁶

Uncovering a value between 1 and 2 for the threshold parameter helps to explain some of the reduced form results above. There is heterogeneity in diffusion thresholds in the population, and the true model appears to lie somewhere in between the two parameterizations ($\lambda = 1$ or 2) we used in the experimental design. This may be why the performance of the simple and complex treatments were statistically indistinguishable for some outcomes.

The fact that the estimated threshold exceeds 1 for a large fraction of people implies that seed targeting strategy within a social network may be quite critical. Most people do not adopt on the basis of just a single contact with an informed person. Picking seeds at random or in different, unrelated parts of the network may lead some technologies to die off before they have a chance to diffuse. This may help explain the low take-up of some novel, productive technologies.

7. Concluding Remarks

This paper sought to understand the diffusion process for new technologies by designing a field experiment closely tied to specific formulations of social network-based diffusion theory. We develop a methodology to select seed farmers who are predicted to maximize diffusion of a productive new agricultural technique in Malawi, under specific theoretical assumptions about the

³⁶ Since the normal distribution imposes continuous thresholds, but the true parameters are ultimately discrete (i.e. a threshold of 1.01 has the same contagion properties as a threshold of 1.99), this calculation discretizes the thresholds to be comparable to the nonparametric case.

process of diffusion. By implementing the field experiments in partnership with the Ministry of Agriculture and Food Security, we were able to test: (a) whether theory-driven targeting using detailed social network data can increase technology adoption relative to the status quo approach to agricultural extension services; (b) whether a less data-intensive approach can generate similar gains; and (c) whether the learning environment for new technologies is “complex” or “simple”, in the language of the linear threshold model.

The theory and simulations provide a few specific predictions for the experiments and data: (a) targeting is only critical if the learning environment is complex: under simple learning all treatments will generate similarly large adoption gains after three years, (b) if diffusion has complex contagion properties, then it is useful to cluster seeds in the same part of the network; otherwise no one crosses the threshold and we would observe no adoption, and (c) under complex learning, multiple connections to seeds should be predictive of adoption. Our estimates suggest that most farmers are convinced to adopt a new technology only if they receive information about it from multiple sources. This implies that diffusion follows a Complex Contagion pattern. Policy makers with local contextual knowledge may be able to suggest inexpensive means of finding entry points to ensure that most farmers have multiple contacts; and future work should explore whether these inexpensive proxies are effective.

Our data show that thinking carefully about the process of diffusion is important for agricultural extension: using status quo practices creates a risk that no one adopts a productive new technology. We document that using threshold diffusion theory to target specific injection points generates significant adoption gains relative to alternative approaches, with particularly large gains on the extensive margin of whether anyone adopts at all. That is, threshold diffusion theory has predictive power for agricultural extension. More broadly, the exercise teaches us that when thresholds matter for adoption, then marketers and policymakers interested in promoting adoption

of new products and technologies should pay attention to the clustering of injection points within a network. Otherwise, productive technologies and promising, useful products may never take off or die out if no one crosses the threshold in any part of the network.

References

- Acemoglu, Ozdaglar and Yildiz (2011). "Diffusion of Innovations in Social Networks," In IEEE Conference on Decision and Control (CDC), 2011.
- Alvarez, F. E., Buera, F. J., & Lucas Jr, R. E. (2013). "Idea Flows, Economic Growth, and Trade." National Bureau of Economic Research, No. w19667.
- Anderson, J, and G Feder (2007). "Agricultural Extension." In *Handbook of Agricultural Economics* 3, pp. 2343-2378.
- Apicella, C. L., F. Marlowe, J. Fowler and N. Christakis (2012). "Social Networks and Cooperation in hunter-gatherers," *Nature* 481: 497-502.
- aWhere, Inc. 2014. WeatherTerrain® Daily Surfaced Weather Data. aWhere, Inc. 4891 Independence St. Wheat Ridge, CO 80033 USA.
- Bandiera, O., & Rasul, I. (2006). "Social networks and technology adoption in northern Mozambique." *The Economic Journal*, 116(514), 869-902.
- Banerjee, A., Chandrasekhar, A. G., Duflo, E., & Jackson, M. O. (2013). "The Diffusion of Microfinance." *Science*, 341(6144). doi: 10.1126/science.1236498
- Banerjee, A., Chandrasekhar, A. G., Duflo, E., & Jackson, M. O. (2014). "Gossip and Identifying Central Individuals in a Social Network." *Working paper*.
- Bardhan, P. and C. Udry (1999). *Development Microeconomics*. Oxford University Press.
- Beaman, L. (2012) "Social Networks and the Dynamics of Labor Market Outcomes: Evidence from Refugees Resettled in the U.S." *Review of Economic Studies*, 79(1), pp. 128-161.
- BenYishay, A. and A. M. Mobarak (2015). "Social Learning and Communication," Working Paper, Yale University.
- Bursztyjn, L., F. Ederer, B. Ferman and N. Yuchtman (2013) "Understanding Mechanisms Underlying Peer Effects: Evidence from a Field Experiment on Financial Decisions," *Econometrica*, Forthcoming.

- Carrel, S., B.I. Sacerdote, and J.E. West. (2013). "From Natural Variation to Optimal Policy? The Importance of Endogenous Peer Group Formation." *Econometrica* 81(3). pp 855-882.
- Centola, Damon and Michael Macy (2007). "Complex Contagions and the Weakness of Long Ties," *American Journal of Sociology*, 113: 702-34.
- Centola, Damon. 2010. "The Spread of Behavior in an Online Social Network Experiment," *Science* 329:1194-7.
- Chandrasekhar, A., H. Larreguy, and J.P. Xandri. (2012) "Testing Models of Learning on Social Networks: Evidence from a Framed Field Experiment" Working Paper, Stanford University.
- Chandrasekhar, A. and R. Lewis (2011) "Econometrics of Sampled Networks," Working Paper, Stanford University.
- Coleman, James, Katz, Elihu, & Menzel, Herbert. (1957). The Diffusion of an Innovation Among Physicians. *Sociometry*, 20(4), 253-270.
- Conley, T. and C. Udry (2010). "Learning about a New Technology." *American Economic Review*, 0100(1), pp 35-69.
- Derpsch, R. (2001). "Conservation tillage, no-tillage and related technologies." *Conservation Agriculture, A Worldwide Challenge* 1:161-170.
- (2004). "History of Crop Production, With & Without Tillage." *Leading Edge*. March: 150-154.
- Doumit, G., M. Gattellari, J. Grimshaw, M. A. O'Brien (2007). "Local opinion leaders: Effects on professional practice and health care outcomes." *Cochrane Database Syst. Rev.* 1 CD000125.
- Duflo, E., M. Kremer, and J. Robinson (2011). "Nudging Farmers to Use Fertilizer: Theory and Experimental Evidence from Kenya." *American Economic Review*, 101(6), pp. 2350-90.
- Duflo, Esther and Emmanuel Saez, (2003). "The Role of Information and Social Interactions in Retirement Plans Decisions: Evidence from a Randomized Experiment," *Quarterly Journal of Economics*, 118 (3), pp. 815-842.
- Easley, David, and Jon Kleinberg. "Networks, crowds, and markets: Reasoning about a highly connected world." Chapter two. Cambridge University Press, 2010.
- Foster, Andrew and Mark Rosenzweig (1995). "Learning by Doing and Learning from Others: Human Capital and Technical Change in Agriculture". *Journal of Political Economy* 103(6), pp. 1176-1209.
- Grannovetter (1978). "Threshold Models of Collective Behavior," *American Journal of Sociology*, 86 (6), 1420-1443.

- Griliches, Zvi (1957). "Hybrid Corn: An Exploration in the Economics of Technical Change." *Econometrica* 25(4), pp. 501-522.
- Haggblade, S. and G. Tembo (2003). "Conservation Farming in Zambia." International Food Policy Research Institute EPTD Discussion Paper #108.
- Iyengar, R., Van den Bulte, C., & Valente, T. W. (2011). Opinion leadership and social contagion in new product diffusion. *Marketing Science*, 30(2), 195-212.
- Kremer, M., E. Miguel, S. Mullainathan, C. Null and A. Zwane (2011). "Social Engineering: Evidence from a Suite of Take-up Experiments in Kenya" Mimeo, UC Berkeley.
- Kim, David A., Alison R. Hwang, Derek Stafford, D. Alex Hughes, A. James O'Malley, James H. Fowler, and Nicholas A. Christakis. (2015). "Social network targeting to maximise population behaviour change: a cluster randomised controlled trial." *The Lancet*.
- Magruder, J. (2010) "Intergenerational Networks, Unemployment, and Inequality in South Africa." *American Economic Journal: Applied Economics*, 2(1), pp 62-85.
- Manski, C. (1993). "Identification of Endogenous Social Effects: The Reflection Problem." *The Review of Economic Studies* 60(3), pp. 531-542.
- Miller, Grant and A. M. Mobarak (2014). "Learning about New Technologies through Opinion Leaders and Social Networks: Experimental Evidence on Non-Traditional Stoves in Rural Bangladesh." *Marketing Science* forthcoming.
- Munshi, Kaivan (2004). "Social learning in a heterogeneous population: Technology diffusion in the Indian Green Revolution." *Journal of Development Economics*, 73(1), pp 185-215.
- Munshi, Kaivan (2007). "Information Networks in Dynamic Agrarian Economies" Handbook of Development Economics, Vol 4, Chapter 48. Pp 3085-3113.
- Munshi, Kaivan (2008). "Social Learning and Development," in Lawrence E. Blume and Steven N. Durlauf, eds. New Palgrave Dictionary of Economics, 2nd Edition. Palgrave Macmillan.
- Mwato, I.L., Mkandawire, A.B.C. and Mughogho, S.K. (1999). "Combined Inputs of Crop Residues and Fertilizer for Smallholder Maize Production in Southern Malawi." *African Crop Science Journal* 7 (4): 365-373.
- Nkhuzenje, H. (2003). "Contribution of promiscuous and specific soybean variety residues to soil fertility improvement and maize yield under smallholder farms in Zomba District, Southern Malawi." Thesis submitted in partial fulfilment of the requirements of Master of Science degree in agronomy (Soil science).
- Nyirongo, J., S. Mughogho, and J. Kumwenda (1999). "Soil Fertility Studies with Compost and Igneous Phosphate Rock Amendments in Malawi." *African Crop Science Journal*, 7(4), pp. 415-422.

- Oster, E. and R. Thornton (2012). "Determinants of Technology Adoption: Private Value and Peer Effects in Menstrual Cup Take-Up." *Journal of the European Economic Association* 10(6):1263-1293, December.
- Perla, J., & Tonetti, C. (2014). "Equilibrium imitation and growth." *Journal of Political Economy*, 122(1), 52-76.
- Rogers, E. M. (1962). *Diffusion of Innovations*. New York, NY: The Free Press.
- Ryan, B. and N.C.Gross (1943). "The Diffusion of Hybrid Seed in Two Rural Iowa Communities" *Rural Sociology*, Vol 8(1). pp 15-24.
- Udry, Christopher (2010). "The Economics of Agriculture in Africa: Notes Toward a Research Program." *African Journal of Agricultural and Resource Economics* forthcoming.
- World Bank (2008). *World Development Report 2008: Agriculture for Development*. Washington, DC: The World Bank.

Table 1: Characteristics of the Seeds (Injection Points) Chosen by Each Treatment Arm

	Wealth Measures		Social Network Measures		
	Farm Size (1)	Total Index (PCA) (2)	Degree (3)	Betweenness Centrality (4)	Eigenvector Centrality (5)
Treatment type					
Simple Contagion	-0.152 (0.19)	0.113 (0.23)	0.455 (1.03)	156.009 (67.93) **	0.009 (0.01)
Complex Contagion	-0.037 (0.19)	0.380 (0.23)	3.725 (1.02)	*** 146.733 (67.74) **	0.064 (0.01) ***
Geographic	-0.614 (0.19) ***	-0.740 (0.23) ***	-3.616 (1.03) ***	*** -90.204 (68.04)	-0.046 (0.01) ***
p-values for tests of equality in seed characteristics					
Simple = Complex	0.335	0.067	0.000	0.815	0.000
Complex = Geographic	0.000	0.000	0.000	0.000	0.000
Simple = Complex = Geographic	0.000	0.000	0.000	0.000	0.000
N	1248	1248	1232	1232	1232
Mean Value for Seeds in Benchmark Treatment (the omitted category)	2.06	0.626	11.9	169	0.173
SD for Seeds in Benchmark Treatment	2.97	1.7	6.77	343	0.0961

Notes

1 The sample includes all seeds and shadows. The sample frame includes 100 Benchmark farmers (2 partners in 50 villages), as we only observe Benchmark farmers in Benchmark treatment villages, and up to 6 additional partner farmers (2 Simple partners, 2 Complex partners, and 2 Geo partners) in all 200 villages.

2 *** p<0.01, ** p<0.05, * p<0.1

Table 2: Adoption Rates of Actual Seeds Relative to Shadow (Counterfactual) Farmers

	Adopted Pit Planting			Adopted Crop Residue Management	
	(1)	(2)	(3)	(4)	(5)
Seed	0.258 *** (0.03)	0.230 *** (0.03)	0.183 *** (0.04)	0.137 *** (0.04)	0.047 (0.04)
N	686	672	490	686	467
Mean of Shadows	0.054	0.093	0.139	0.320	0.207
Season	1	2	3	1	2

Notes

¹ Also included are village fixed effects. Sample includes only seed farmers (chosen by the simulations and trained on the technologies through our intervention) and shadow (counterfactual) farmers with the same network characteristics but not trained by the experiment. The sample excludes Benchmark villages. Standard errors are clustered at the village level.

² *** p<0.01, ** p<0.05, * p<0.1

Table 3: Difference in Adoption Rates Across Seed Farmers chosen through Different Targeting Strategies

Treatment:	Adopted Pit Planting			Adopted Crop Residue Management	
	(1)	(2)	(3)	(4)	(5)
Simple Contagion	-0.006 (0.07)	0.129 * (0.07)	0.176 ** (0.09)	0.078 (0.08)	-0.097 (0.09)
Complex Contagion	-0.020 (0.08)	0.002 (0.07)	0.037 (0.08)	-0.001 (0.08)	-0.077 (0.09)
Geographic	-0.095 (0.08)	-0.064 (0.07)	-0.003 (0.08)	-0.011 (0.08)	-0.075 (0.10)
N	353	352	259	353	243
Mean Adoption for Seeds in Benchmark Treatment (the omitted category)	0.337	0.276	0.238	0.442	0.339
<i>p-value for tests of equality in adoption rates across treatment cells:</i>					
Simple = Complex	0.862	0.077	0.108	0.311	0.808
Complex = Geographic	0.360	0.358	0.625	0.886	0.977
Joint test of 3 treatments	0.252	0.008	0.049	0.235	0.795
Season	1	2	3	1	2

Notes

1 Also included are stratification controls (percent of village using compost at baseline; percent village using fertilizer at baseline; percent of village using pit planting at baseline); village size and its square; and district fixed effects. Only seed farmers chosen by our interventions to be trained on the technologies are included. Standard errors are clustered at the village level.

2 *** p<0.01, ** p<0.05, * p<0.1

Table 4: Conversations other farmers report having about Pit Planting with Seed and Shadow Partners

	with Simple Partner		with Complex Partner		with Geo Partner
Treatment arm	(1)		(2)		(3)
Simple Contagion	0.047 (0.015)	***	0.018 (0.012)		0.005 (0.009)
Complex Contagion	0.020 (0.010)	*	0.036 (0.014)	**	0.000 (0.008)
Geographic	0.005 (0.012)		0.004 (0.009)		0.031 (0.016)
N	3728		3677		3718
Mean of Benchmark	0.019		0.025		0.018
SD of Benchmark	0.135		0.157		0.133
<i>p-values for equality in coefficients:</i>					
Simple = Complex	0.077		0.206		0.543
Complex = Geo	0.140		0.008		0.028
Simple = Geo	0.005		0.160		0.079
Season	1		1		1

Notes

- 1 Sample excludes seed and shadow farmers. We only collected this data in 2011-2012, and therefore this analysis is restricted to Mwanza and Machinga.
- 2 Also included are stratification controls as listed in table 3; district fixed effects; and controls for the number of partner farmers (of the type asked about in the respective column) we asked about in the questionnaire by including a dummy variable for each number of partner farmers from 0 to 4.
- 3 Standard errors clustered at the village level.
- 4 *** p<0.01, ** p<0.05, * p<0.1

Table 5: Individual-level Analysis of Pit Planting Awareness and Adoption

	Season 1			Season 2			Season 3		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Aware of PP	Adopted PP		Aware of PP	Adopted PP		Aware of PP	Adopted PP	
Connected to 1 seed	-0.007 (0.023)	0.009 (0.011)	0.007 (0.011)	0.032 (0.024)	0.015 (0.015)	0.012 (0.015)	0.020 (0.030)	0.010 (0.016)	0.004 (0.017)
Connected to 2 seeds	0.071 * (0.037)	0.017 (0.014)	0.015 (0.014)	0.107 *** (0.040)	0.042 ** (0.019)	0.039 ** (0.019)	0.071 (0.064)	0.019 (0.034)	0.014 (0.035)
Within Path Length 2 of a seed			0.013 (0.008)			0.022 * (0.013)			0.037 * (0.021)
N	4150	4203	4203	4532	3931	3931	3103	2998	2998
Mean of Reference Group (No connection to any seed)	0.214	0.023	0.013	0.274	0.056	0.044	0.387	0.064	0.043
SD of Reference Group	0.410	0.151	0.113	0.446	0.230	0.206	0.487	0.244	0.203
<i>p-value</i> for 2 connections = 1 connection	0.022	0.550	0.528	0.052	0.171	0.169	0.409	0.791	0.760

Notes

- 1 Sample excludes seed and shadow farmers. Only connections to simple, complex and geo seed farmers are considered (no connections to benchmark farmers included).
- 2 The dependent variable in columns (1), (4) and (7) is an indicator for whether the respondent reported being aware of a plot preparation method other than ridging and then subsequently indicated awareness of pit planting in particular. The dependent variable in the remaining columns are an indicator for the household having adopted pit planting in that season.
- 3 In columns (1), (2), (4), (5), (8) and (9), additional controls include indicators for the respondent being connected to: one Simple partner, two Simple partners, one Complex partner, two Complex partners, one Geo partner and two Geo partners.
- 4 In columns (3), (6) and (9), additional controls include all of the odd column controls plus indicators for the respondent is: within 2 path length of a Simple partner, within 2 path length of a Complex Partner, within 2 path length of the geo partner.
- 5 Also included in both panels are village fixed effects. Standard errors clustered at the village level.
- 6 The excluded group is comprised of individuals with no connections to a seed farmer (and no 2-path-length connections in columns 3,6, and 9).
- 7 *** p<0.01, ** p<0.05, * p<0.1

Table 6: Village-Level Regressions of Adoption Outcomes Across Treatment Arms

	Adoption Rate		Number Adopters		Any Non-Seed Adopters			
	(1)	(2)	(3)	(4)	(5)	(6)		
Simple Contagion Treatment	0.036 (0.017)	** (0.022)	0.006 (0.859)	1.119 (1.561)	0.283 (2.007)	0.155 (0.100)	0.189 (0.111)	*
Complex Contagion Treatment	0.036 (0.016)	** (0.026)	0.036 (1.386)	2.876 (2.007)	** (2.007)	2.237 (0.093)	0.252 (0.101)	***
Geographic treatment	0.038 (0.027)	0.013 (0.034)	0.356 (0.823)	-1.379 (1.284)	0.107 (0.096)	0.188 (0.110)	*	
Year	2	3	2	3	2	3		
N	200	141	200	141	200	141		
Mean of Benchmark Treatment (omitted category)	0.038	0.075	2.11	4.81	0.420	0.543		
SD of Benchmark	0.073	0.109	4.17	7.81	0.499	0.505		
<i>p-values for equality in coefficients:</i>								
Simple = Complex	0.981	0.173	0.258	0.373	0.300	0.240		
Complex = Geo	0.937	0.491	0.084	0.061	0.102	0.220		
Simple = Geo	0.950	0.783	0.424	0.237	0.623	0.990		

Notes

- 1 Adoption rate and Number of adopters reflect outcomes for all sampled farmers excluding seed and shadow farmers. "Any non-seed adopters" excludes only seed farmers.
- 2 Sample for year 3 (columns 2, 4 and 6) excludes Nkhotakota district.
- 3 The excluded group is the benchmark treatment.
- 4 Columns 3-4 include sample weights for village size.
- 5 *** p<0.01, ** p<0.05, * p<0.1

Table 7: Village Level Regressions of Adoption Outcomes across Treatment Arms for Villages that had Low familiarity with Pit Planting at Baseline

	Adoption Rate		Number of Adopters		Any Non-Seed Adopters		
	(1)	(2)	(3)	(4)	(5)	(6)	
Simple Contagion Treatment	0.043 (0.028)	0.018 (0.022)	1.047 (1.504)	1.126 (1.618)	0.160 (0.153)	0.312 (0.151)	**
Complex Contagion Treatment	0.050 ** (0.025)	0.092 *** (0.032)	5.077 ** (2.460)	6.154 ** (2.729)	0.344 ** (0.132)	0.458 *** (0.131)	***
Geographic treatment	0.024 (0.026)	0.035 (0.032)	0.231 (1.349)	1.071 (1.852)	0.131 (0.144)	0.350 (0.153)	**
Season	2	3	2	3	2	3	
N	99	82	99	82	99	82	
Mean Value for Benchmark Treatment (omitted category)	0.029	0.051	1.96	3.44	0.375	0.450	
SD of Benchmark	0.074	0.081	4.28	6.68	0.495	0.510	
<i>p-values for equality in coefficients:</i>							
Simple = Complex	0.812	0.016	0.157	0.109	0.157	0.279	
Complex = Geo	0.371	0.120	0.074	0.112	0.075	0.428	
Simple = Geo	0.541	0.559	0.580	0.971	0.830	0.800	

Notes

- 1 The sample is restricted to villages where less than 4.32% of households (the median) ever tried pit planting at baseline.
- 2 Sample for year 3 (columns 2, 4 and 6) excludes Nkhotakota district.
- 3 The excluded group is the benchmark treatment
- 4 Columns 3-4 include sample weights for village size.
- 5 *** p<0.01, ** p<0.05, * p<0.1

Table 8: Structural Estimates of the Distribution of the threshold (λ) parameter

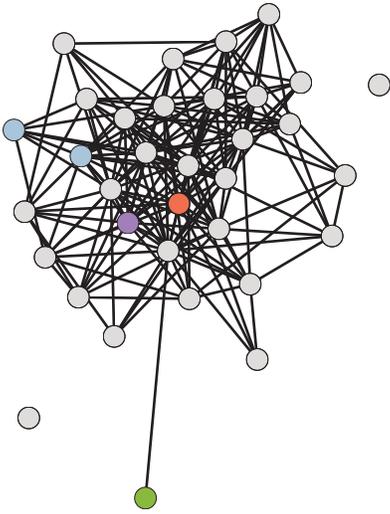
	(1)	(2)	(3)	(4)
E[λ]	1.80 (1.62,2.24)	1.74 (1.58,2.24)	1.78 (1.46,2.84)	1.77 (1.5,2.75)
Fraction with threshold = 1	0.28 (0.17,0.50)	0.30 (0.18,0.52)	0.31 (0.27,.52)	0.34 (0.21,0.53)
Fraction with threshold = 2	0.64 (0.0,0.78)	0.66 (0.0,0.78)	0.67 (.16,.74)	0.60 (0.16,0.74)
Fraction with threshold >2	0.08 (0.04,0.6)	0.04 (0,.6)	0.02 (0,.484)	0.05 (0,.47)
Moments from Benchmark Treatment	Yes	No	Yes	No
Distributional assumption	$\lambda \leq 3$	$\lambda \leq 3$	$N(\lambda, \sigma)$	$N(\lambda, \sigma)$

Notes

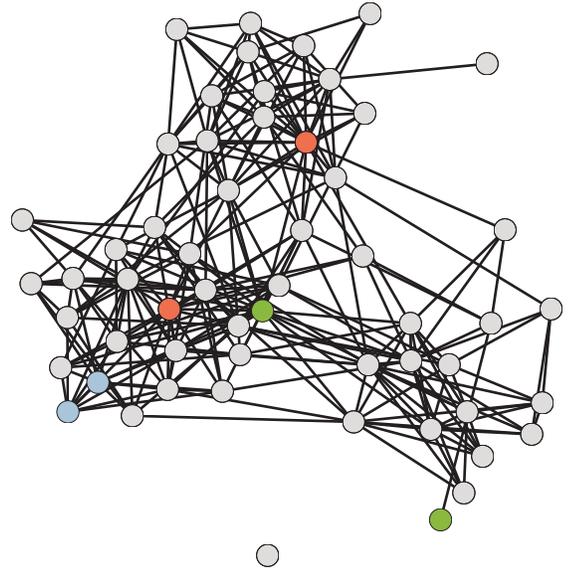
- 1 Reports results from Method of Simulated Moments estimation of mean village-level outcomes in treatment groups against simulated means across the sample. Bootstrapped confidence intervals shown in parantheses below the estimate are resampled at the village level.

Figure 1: Network Maps

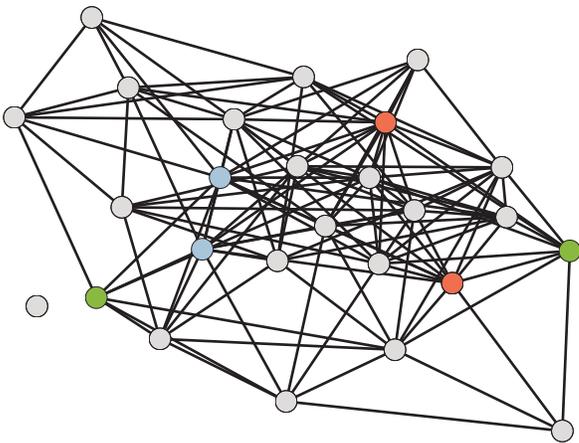
Village 45



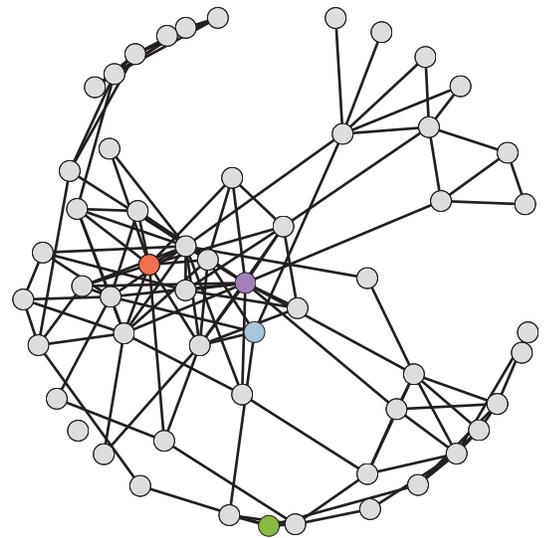
Village 63



Village 87



Village 141



Village 288

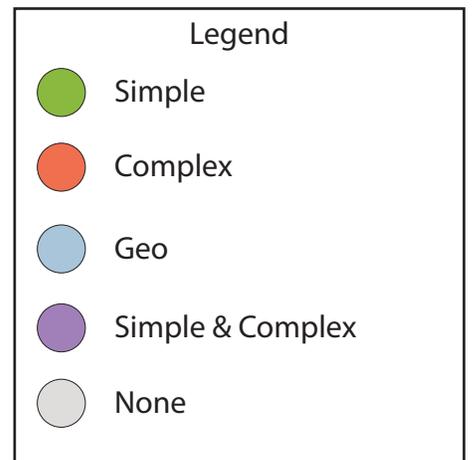
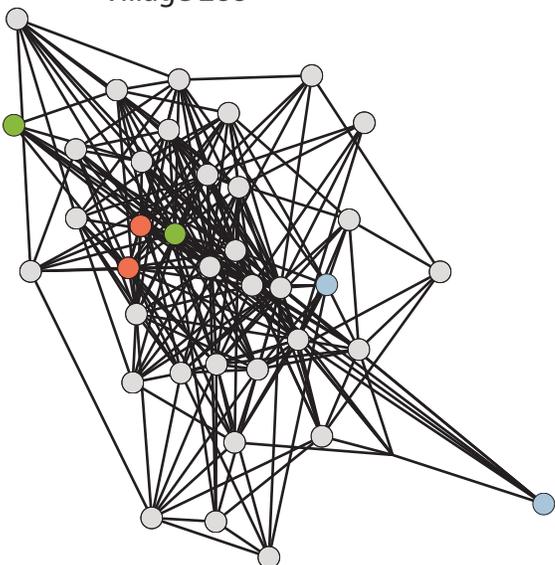
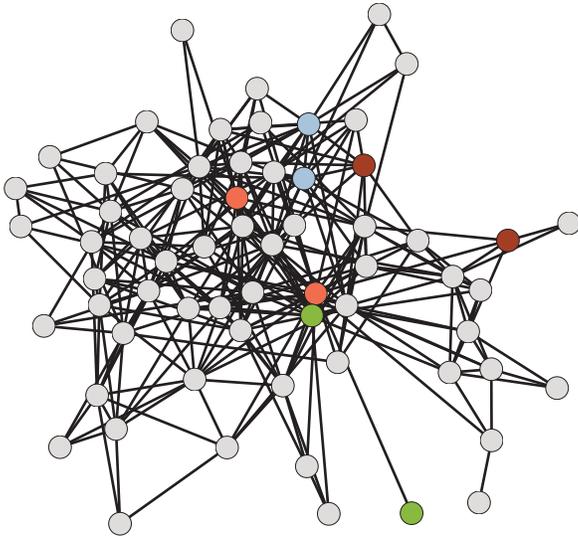
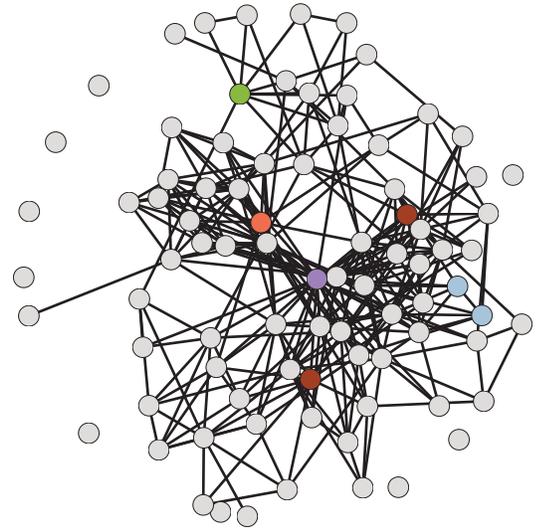


Figure 2: Network Maps of Benchmark Seed Farmers

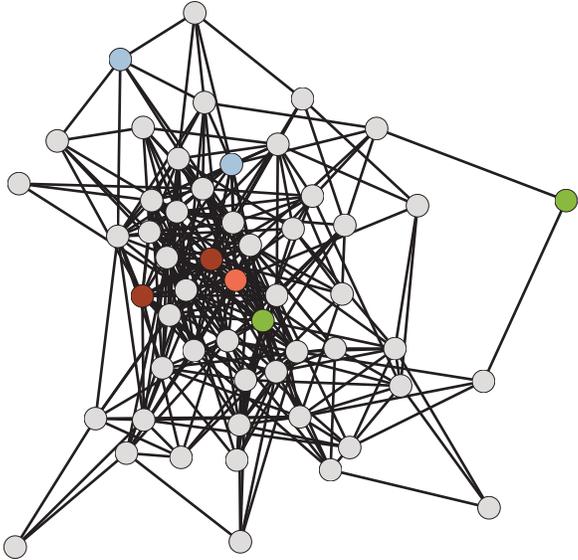
Village 33



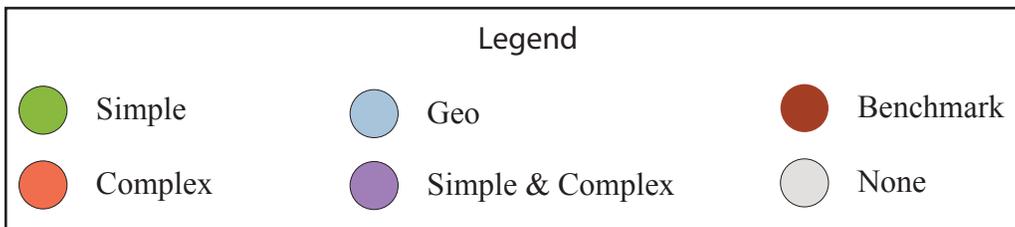
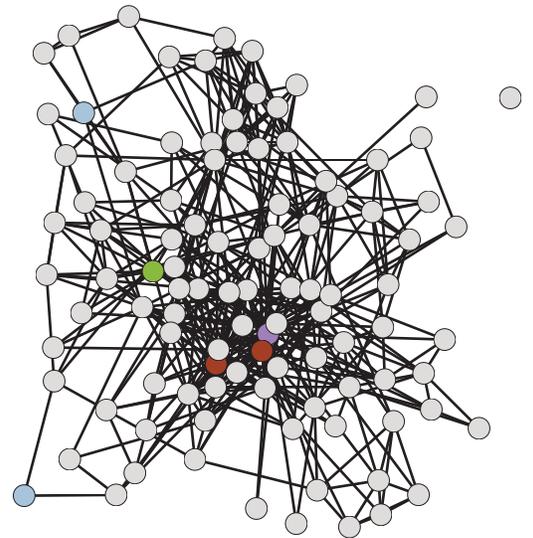
Village 42



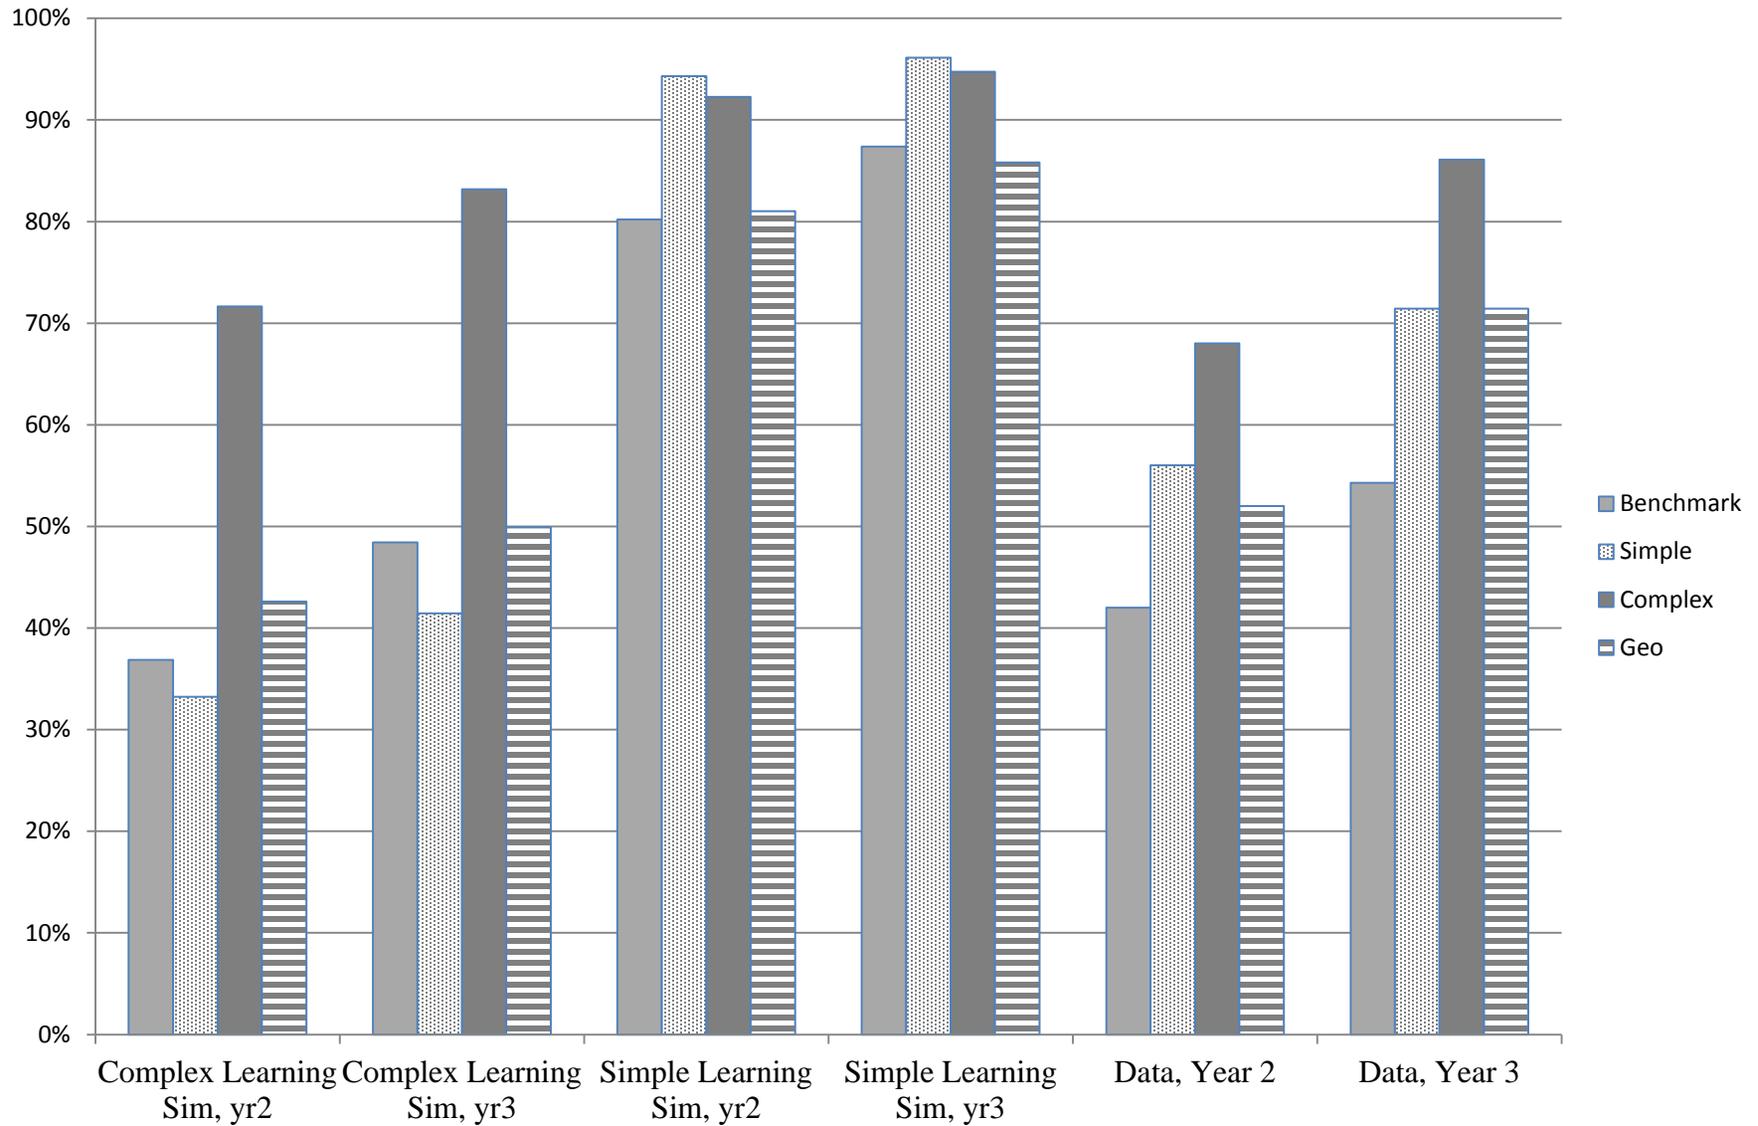
Village 43



Village 136



**Figure 3: Fraction of Villages where at least some non-seeds Adopted:
Comparing the Data to Simple or Complex Contagion Simulations**



Online Appendices for Beaman et al, "Can Network Theory Based Targeting Increase Technology Adoption?"

Appendix

A.1. Simulations of adoption regressions in section 5.2

Appendix table A3 table reflects the adoption patterns we should observe under the parameterizations of the simple and complex learning environments we used in identifying our treatment partners. Table A3 presents these simulation results for three different measures of technology adoption: the adoption rate, the total number of adopters, and an indicator for villages with any non-seed adopters. We predict these outcomes for all four experimental arms that were implemented in the field. Panel A shows what we should expect to observe across treatments based on simulations of the model with $\lambda=1$ (Simple contagion), and Panel B reports predictions under $\lambda=2$ (Complex contagion).³⁷

Columns (1)-(2) show the results for adoption rate outcomes. Complex partners initially maximize adoption in year 2 even if the learning environment is simple, but in year 3 adoption rate is highest when the simple seeds are trained. However, the effects of training simple and complex seeds are not statistically distinguishable ($p=.8$) for these outcomes simulated under simple contagion. Under simple contagion, villages where the Geo seeds are trained exhibit the lowest adoption rates. Columns (3)-(4) show a very similar set of results for the number of adopters under simple contagion. Taken together, these results indicate that the simple treatment is not expected to dominate alternative targeting strategies even if the contagion process is simple. This reinforces the intuition that if farmers truly have a low threshold for adoption, the diffusion process is not likely to be particularly sensitive to who is initially targeted with information.

In contrast, when we conduct simulations assuming the complex contagion model is correct, the complex treatment is predicted to increase adoption significantly more than all other treatments (Panel B of Table 6). The Complex treatment out-performs the simple, Geo and Benchmark

³⁷ The table differs from Figure 1 in two key dimensions: (1) this uses the realized randomization and not all villages as in figure 1, and (2) includes additional stratification control variables as in the empirical analysis.

treatments in terms of all adoption outcomes during both years (with statistical tests for differential effects producing p-values below 0.001 for every comparison).

A.2 Alternative diffusion models

There are many diffusion models which would generate similar predictions on targeting. Our analysis does not allow us to test which of these models generates the best fit to our data but only to test whether a particular parameterization of the threshold model can predict extension partners who improve adoption outcomes.

Nonetheless, we may wonder how the results compare to other well-known diffusion models. SI, or Susceptible-Infected models, are a widely discussed class of models in the diffusion literature (Jackson 2008). Under an SI model, when a susceptible node is connected to an infected node, it becomes infected with some probability ρ . In contrast to the threshold model, in an SI model the marginal impact of the n^{th} connection to an infected node is $\rho(1 - \rho)^{n-1}$ for all n , whereas in the threshold model it is either 0 (if n is less or more than the threshold) or 1 (if n is the threshold value).

In our data, we can use simulation to demonstrate the patterns we should have expected to see if the true model were an SI model for different values of ρ .³⁸ Appendix Figures A2 and A3 present bar charts indicating the average predicted values of our binary any adoption and mean adoption rate measures for different values of ρ across all villages after two rounds in the data, corresponding roughly to season 2 in our data. Both figures are quite clear: we can clearly reject the hypothesis that the true model is an SI model if ρ is much different from 0.1: for a $\rho < 0.1$ we should have observed hardly any adoption at all in year 2; and for $\rho > 0.1$ we should have observed nearly all villages having at least one adopter in our data. In fact, the estimate is pretty sharp in these figures as the slope of the dependent variable is quite steep with respect to ρ : we can say, fairly

³⁸ In these simulations, we use the actual trained seeds in the 188 study villages where we observe both seed farmers.

confidently, that the SI model that best fits our dependent variables is one where the probability of infecting a contact is pretty close to 0.1. In other words, under the SI model, we would estimate that being connected to a trained seed results in a 10% chance of adopting pit planting³⁹.

We can also explore how different we would anticipate treatment effect estimates looking under this formulation of the SI model. Appendix Figures A4 and A5 present bar charts indicating mean any adoption and adoption rate measures across treatment groups under an SI model with $\rho = 0.1$. Just as in our simulated results under our baseline complex contagion parameterization, the complex contagion treatment sees somewhat higher any adoption and mean adoption rate than the other three groups. While there are some distinct patterns in the SI vs Complex model, such as the relative performance of Geo, the broad measures are pretty similar and given the statistical power in our study we wouldn't be able to reject a difference between the two models. Going a step farther, we can examine whether village-to-village variation looks similar between these two modeling assumptions too. The correlation in the simulated binary any adoption variable between the two models predictions across villages is 0.82; while the correlation of simulated adoption rates is 0.92 across villages. The models suggest similar levels of the dependent variables, differences across treatment groups, and differences across villages as a whole.

With this in mind, it is worth asking whether it would be useful or possible to separate the two models even in a higher-powered experiment. The SI model has different predictions for some

³⁹ These simulations take our network graph as given and assume links are homogeneous. Of course, our network graph may be measured with error and links may be either underreported (if respondent suffer from recall bias or survey fatigue), or overreported (if respondents list irrelevant or unconnected links), and links are likely to be heterogeneous (where some links are activated more frequently than others). If links are underreported, the similar SI model would be associated with an even lower level of ρ while overreporting would suggest a higher level. While our survey efforts attempted to solicit as many agricultural learning links as possible and restrict attention only to plausible agricultural learning links, either of these cases remains plausible. If there is measurement error in the network graph, or if links are heterogeneous, the association between the SI model with $\rho=0.1$ and the complex contagion model can still be interpreted as the expected value of adoption being transmitted along a reported link. With measurement error or heterogeneity, that expected value simply now incorporates the probability that a link does or does not get activated (possibly as it may or may not exist). Given that the SI model remains agnostic about the source of the probabilistic dissemination, the probability that a link on a network graph gets activated may be a sensible justification for this model in any event.

theoretically possible networks. Consider, for example, a network where there are two dense clusters separated by a single bridge. Under an SI model with a low transmission rate, one would plant a seed farmer in each cluster. Under the threshold model, one would plant both seeds in the same cluster to guarantee some mutual ties. However, such a model looks very different from the networks seen in Figures 1 and 2, and observed more generally in the data: the networks in our data, like most empirical networks, feature a single giant component.⁴⁰ Both an SI model where transmission is difficult and the complex contagion model would suggest finding well-connected members of that component; in many if not most networks, those would be very similar people, which explains the similarity of predictions under the alternate models. There may exist social networks where the top two central individuals do not share common links – the necessary condition for complex contagion – but this type of network structure is going to be quite rare.

As a result, we conclude that most policy prescriptions stemming from our study would hold in most real world networks under either modeling assumption, which both explains our empirical inability to separate the two and lessens the urgency of doing so.

⁴⁰ See Easley and Kleinberg (2010) for a general discussion. The villages in Banerjee et al (2013) – for example – exhibit this same property.

Table A1: Test of Balance across Randomized Treatment Arms

	Housing	Assets	Livestock	Basal fertiliser (kg)	Top dressing fertiliser (kg)	# of Adults	# of Children	Farm size (acres)	Own land	Yields	Provided Ganyu	Used Ganyu
	(1)	(2)	(3)	(4)	(5)	(7)	(8)	(9)	(10)	(11)	(12)	(13)
Benchmark	0.133 (0.031)	0.0327 (0.027)	0.0251 (0.028)	52.12 (2.699)	53.13 (2.154)	2.323 (0.016)	2.617 (0.033)	1.825 (0.040)	0.912 (0.005)	302.6 (5.469)	0.230 (0.007)	0.145 (0.006)
Simple Treatment	0.142 (0.029)	-0.035 (0.025)	0.023 (0.026)	52.57 (2.522)	51.53 (1.999)	2.313 (0.015)	2.625 (0.031)	1.628 (0.037)	0.905 (0.004)	307.3 (5.122)	0.252 (0.007)	0.129 (0.005)
Complex Treatment	-0.009 (0.032)	-0.010 (0.027)	0.032 (0.028)	53.84 (2.732)	51.51 (2.179)	2.327 (0.016)	2.676 (0.034)	1.690 (0.040)	0.908 (0.005)	295.7 (5.547)	0.248 (0.007)	0.138 (0.006)
Geo Treatment	0.027 (0.031)	-0.023 (0.026)	-0.082 (0.027)	51.60 (2.635)	50.97 (2.102)	2.309 (0.015)	2.643 (0.033)	1.771 (0.039)	0.905 (0.005)	304.3 (5.369)	0.244 (0.007)	0.151 (0.006)
Observations	14,027	14,284	14,284	10,374	10,475	14,041	14,284	14,021	14,284	13,438	14,016	14,016
<i>p-value for...</i>												
Control = Simple	0.004	0.131	0.457	0.998	0.441	0.965	0.335	0.011	0.838	0.915	0.215	0.033
Control = Complex	0.198	0.213	0.586	0.657	0.327	0.718	0.240	0.072	0.953	0.399	0.302	0.212
Control = Geo	0.548	0.327	0.065	0.901	0.629	0.980	0.519	0.618	0.812	0.811	0.759	0.575
Simple = Complex	0.150	0.696	0.776	0.708	0.897	0.613	0.664	0.554	0.856	0.389	0.944	0.280
Simple = Geo	0.325	0.624	0.273	0.925	0.705	0.989	0.802	0.124	0.924	0.871	0.496	0.026
Complex = Geo	0.808	0.875	0.105	0.573	0.570	0.711	0.542	0.322	0.840	0.480	0.562	0.121
Joint	0.034	0.446	0.210	0.949	0.775	0.959	0.662	0.058	0.993	0.810	0.598	0.076

Notes

1 Housing, assets and livestock are pca scores. Housing includes information on: materials walls are made of, roof materials, floor materials and whether the household has a toilet. Assets includes the number of bicycles, radios and cell phones the household owns. Livestock is an index including the number of sheep, goats, chickens, cows, pigs guinea fowl, and doves.

2 Standard errors in parentheses

3 *** p<0.01, ** p<0.05, * p<0.1

Table A2: Effect of Pit Planting on Agricultural Yields

	(1)	(2)
Estimation	OLS	IV
Adopted PP		0.443 ** (0.210)
Seed	0.126 ** (0.061)	
Observations	959	959

Notes

- 1 All columns include district and season FE and controls for total farm size, village size, and village baseline usage of fertilizer, composting and pit planting. The sample includes only seeds and shadows and excludes Benchmark villages.
- 2 Standard errors clustered by village in parentheses.

Table A3: Frequency with which Different Targeting Strategies
Select the same Individuals as Seeds

	Complex	Geo	Benchmark
Simple	25.51%	3.83%	9.50%
Complex		6.00%	11.22%
Geo			6.12%

Notes

- 1 Each row indicates the probability that a potential seed of the targeting strategy indicated in the row is also a seed of the targeting strategy indicated in the column.
- 2 The Benchmark column examines only villages assigned to the benchmark treatment.

Table A4: Simulation of Village Level Adoption Outcomes across all treatment cells, assuming Diffusion follows either Complex or Simple Contagion Pattern

	Simulated Adoption Rate				Simulated Number of Adopters				Simulated Any Adopters	
	(1)		(2)		(3)		(4)		(5)	(6)
Panel A: Simulations Assuming Farmers learn by Simple Contagion										
Simple Treatment	0.026	0.090	*	0.903	5.082	0.095	**	0.036		
	(0.024)	(0.052)		(1.280)	(6.036)	(0.043)		(0.037)		
Complex Treatment	0.087	***	0.072	3.570	**	4.061	0.060	0.013		
	(0.029)	(0.063)		(1.553)	(6.106)	(0.048)		(0.045)		
Geo treatment	-0.022	-0.113	**	-3.711	***	-10.200	*	-0.050	-0.070	
	(0.027)	(0.057)		(1.193)	(5.368)	(0.053)		(0.054)		
Year	2	3		2	3	2		3		
N	187	138		187	138	187		138		
Mean Benchmark Partners	0.182	0.504		11	38.1	0.845		0.927		
SD Benchmark Partners	0.149	0.306		6.03	24	0.258		0.186		
Test: Simple = Complex	0.013	0.733		0.057	0.868	0.384		0.559		
Test: Complex = Geo	0.000	0.001		0.000	0.012	0.026		0.129		
Test: Simple = Geo	0.035	0.000		0.000	0.005	0.001		0.030		
Panel B: Simulations Assuming Farmers Learn by Complex Contagion										
Simple Treatment	0.001	-0.022		-0.833	**	-4.895	***	-0.092	-0.109	
	(0.012)	(0.040)		(0.338)	(1.830)	(0.056)		(0.077)		
Complex Treatment	0.047	***	0.162	***	1.967	***	8.189	***	0.257	***
	(0.012)	(0.046)		(0.494)	(2.602)	(0.061)		(0.081)		
Geo treatment	0.008	-0.032		-0.598	*	-3.771	*	-0.028	-0.048	
	(0.011)	(0.038)		(0.342)	(1.922)	(0.060)		(0.083)		
Season	2	3		2	3	2		3		
N	187	138		187	138	187		138		
Mean Benchmark Partners	0.038	0.138		2.04	8.21	0.436		0.541		
SD Benchmark Partners	0.0479	0.194		1.79	9.37	0.341		0.39		
Test: Simple = Complex	0.000	0.000		0.000	0.000	0.000		0.000		
Test: Complex = Geo	0.001	0.000		0.000	0.000	0.000		0.000		
Test: Simple = Geo	0.533	0.777		0.323	0.304	0.192		0.370		

Notes

1 Simulations only include control villages where we had both seeds in social network census.

Table A5: Individual-level analysis of Crop Residue Management Familiarity and Adoption

	Season 1		Season 2			
	(1)	(2)	(3)	(4)		
	Adopted CRM	Aware of CRM	Adopted CRM	Aware of CRM		
Connected to one seed	-0.026 (0.020)	0.045 (0.022)	**	0.005 (0.028)	-0.042 (0.023)	*
Connections to two seeds	-0.011 (0.034)	0.096 (0.037)	**	0.057 (0.047)	-0.021 (0.038)	
N	4203	4149		2683	4531	
Mean of Excluded Group	0.239	0.295		0.159	0.366	
SD of Excluded Group	0.427	0.456		0.366	0.482	
<i>p</i> -value of Test: 2 connections = 1 connection	0.661	0.143		0.189	0.593	

Notes

- 1 Sample excludes seed and shadow farmers in all villages, and excludes control villages. Seed farmers are either simple, control or geo (no benchmark farmers included).
- 2 Additional controls include indicators for the respondent being connected to: one Simple partner, two Simple partners, one Complex partner, two Complex partners, one Geo partner and two Geo partners.
- 3 Also included in both panels are village fixed effects.
- 4 The excluded group is comprised of individuals with no connections to a seed farmer.
- 5 Columns 1, 2 and 4 include data on Mwanza, Machinga and Nkhotakota. Column 3 includes only Mwanza and Machinga, as the final round of data collection in Nkhotakota was prior to the decision on whether to use CRM.

Table A6: Village Level Adoption Outcomes for Crop Residue Management

	Adoption Rate	Number Adopters	Any Non-Seed Adopters
	(1)	(2)	(3)
Simple Contagion Treatment	-0.037 (0.027)	-3.333 (1.934)	* -0.083 (0.062)
Complex Contagion Treatment	-0.026 (0.027)	-1.198 (2.408)	-0.064 (0.060)
Geographic treatment	-0.054 (0.029)	* -3.677 (2.167)	* -0.152 (0.070)
Year	2	2	2
N	141	141	141
Mean of Benchmark Treatment (omitted category)	0.204	13.70	0.971
SD of Benchmark	0.109	15.30	0.169
<i>p</i> -values for tests of equality of coefficients...			
Test: Simple = Complex	0.680	0.380	0.794
Test: Complex = Geo	0.366	0.332	0.258
Test: Simple = Geo	0.583	0.870	0.336

Notes

- 1 Adoption rate and Number of adopters reflect outcomes for all sampled farmers excluding seed and shadow farmers. "Any non-seed adopters" excludes only seed farmers.
- 2 Analysis only includes Mwanza and Machinga.
- 3 The excluded group is the benchmark treatment.
- 4 Column 2 includes sample weights for village size.
- 5 *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Figure A1: Project Timeline

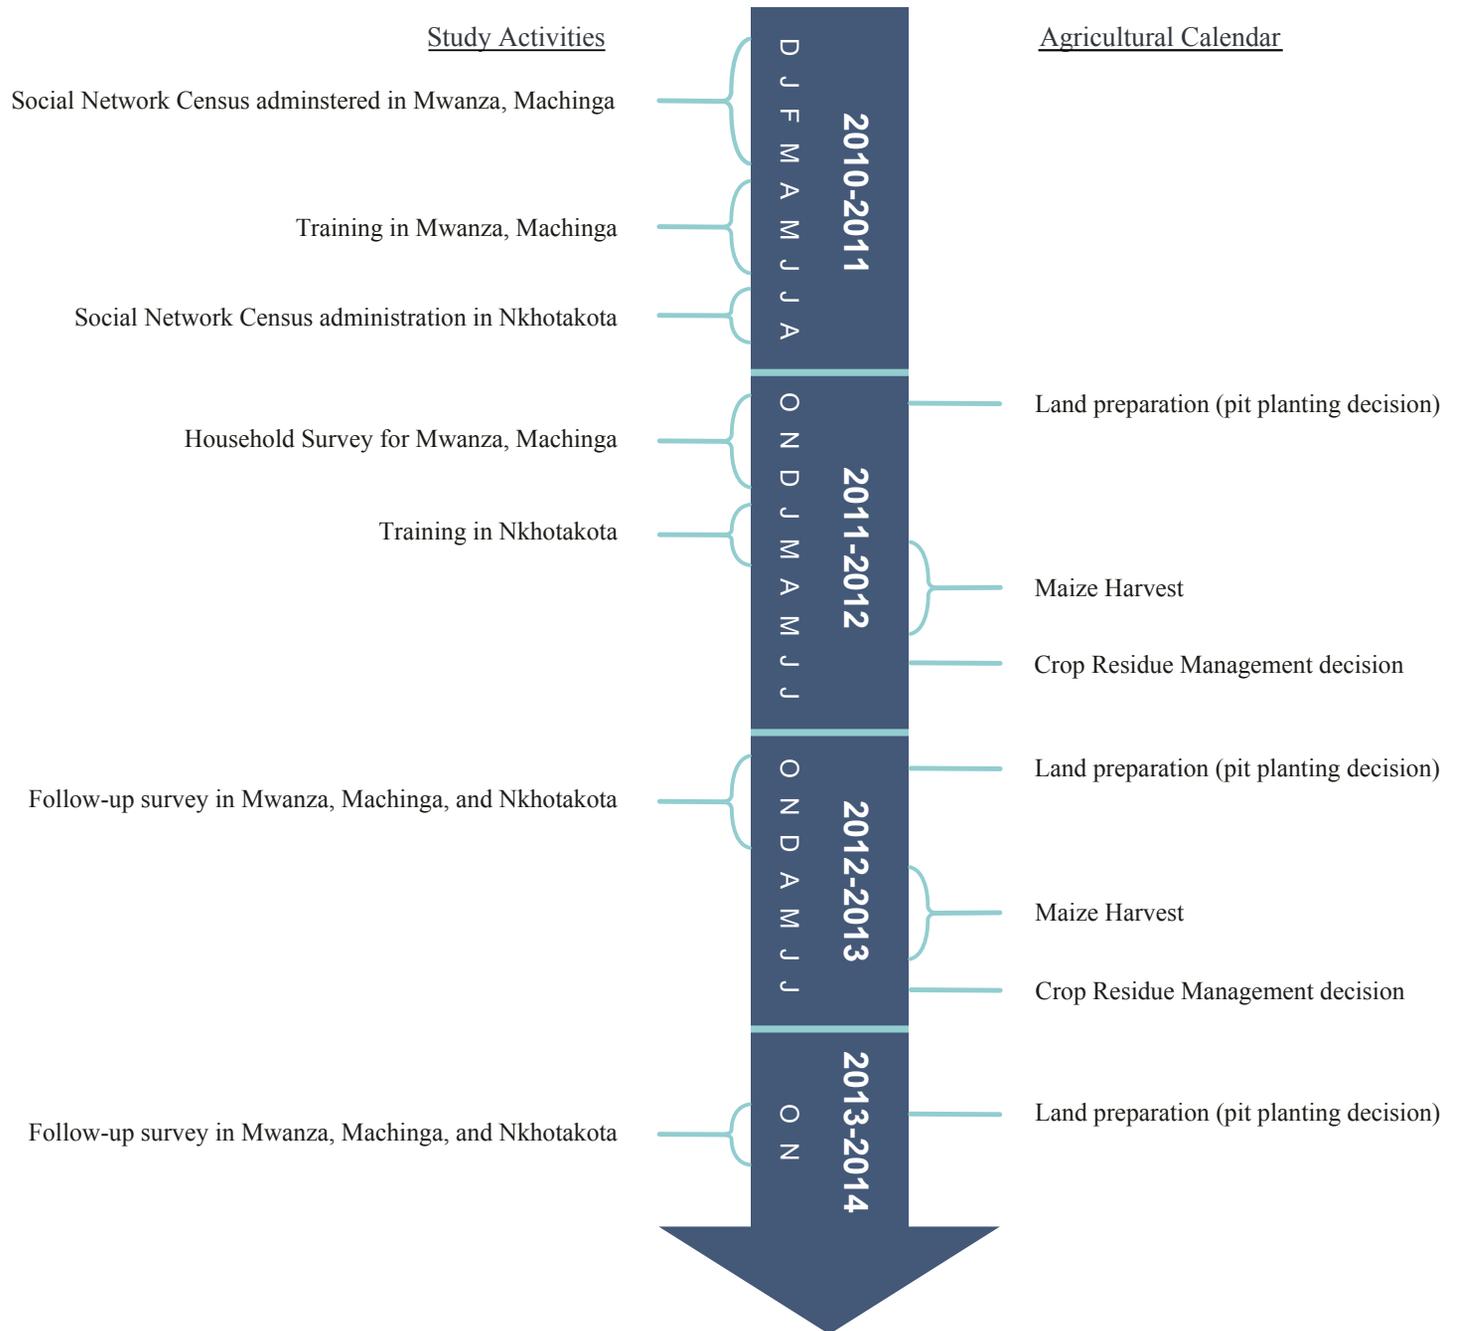


Figure A2: Simulation of SI Model, Any Adoption

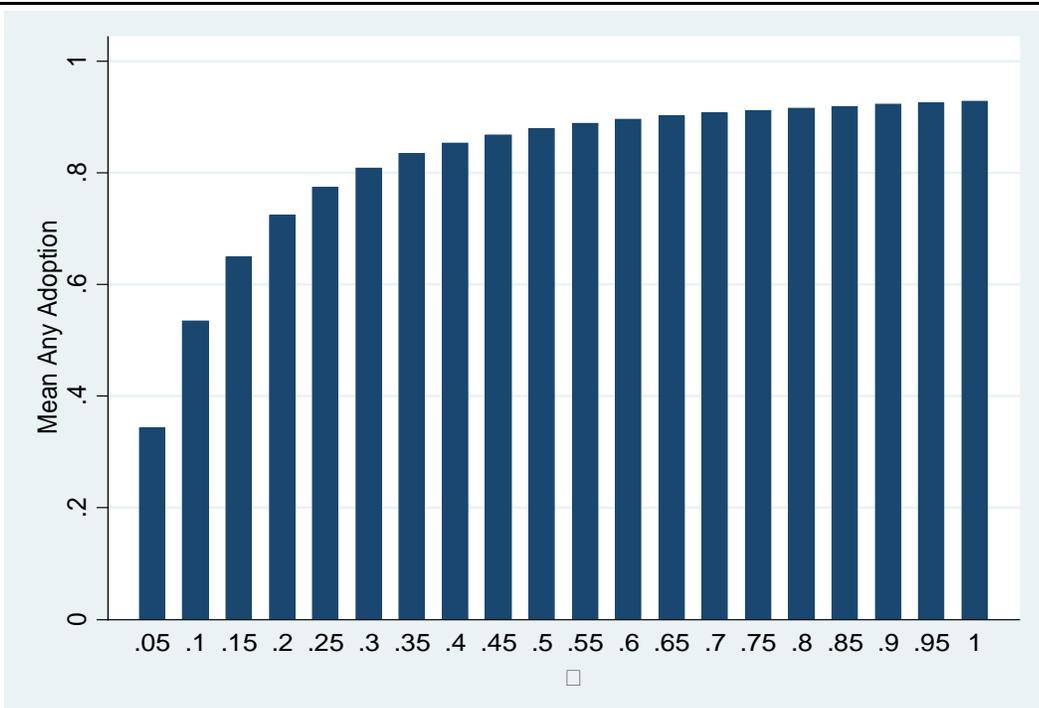


Figure A3: Simulation of SI Model, Mean Adoption Rate

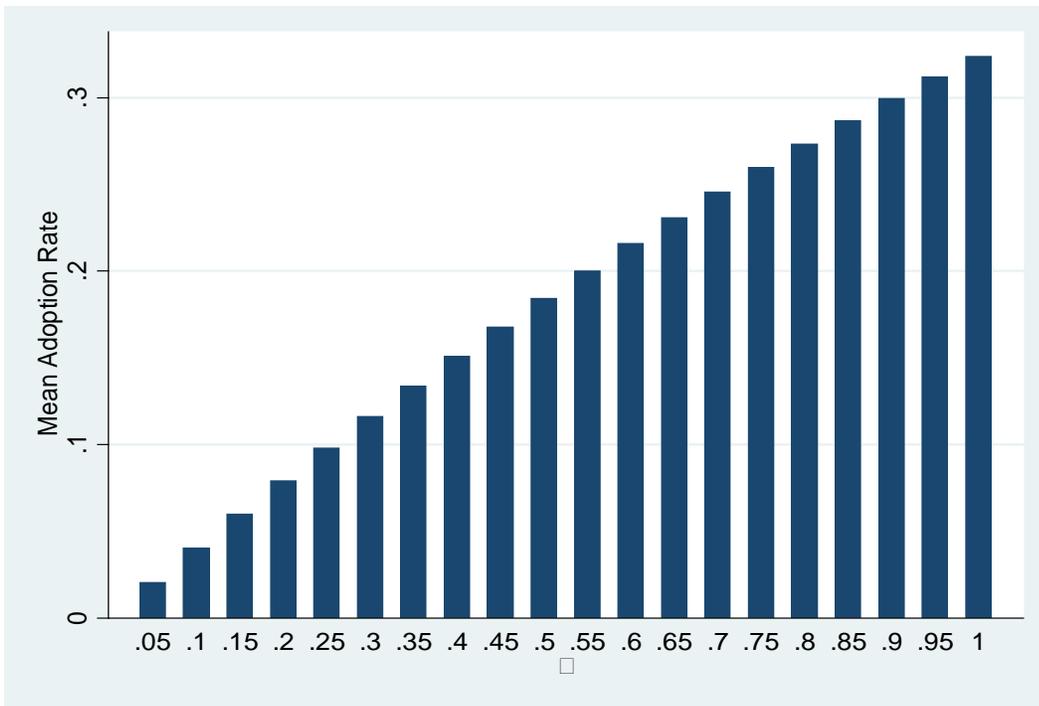
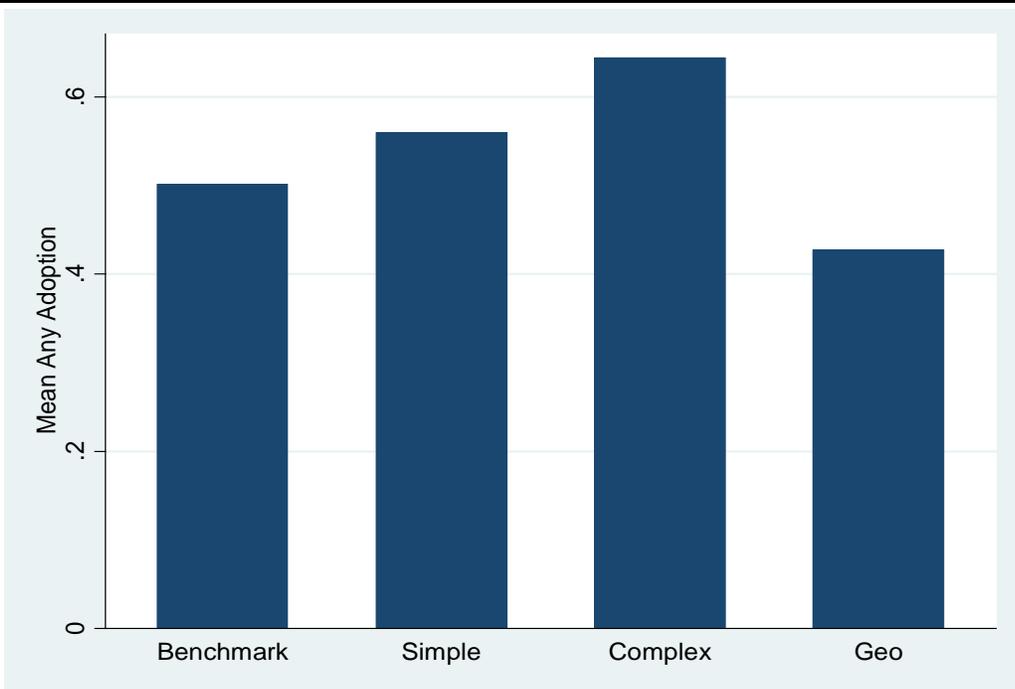
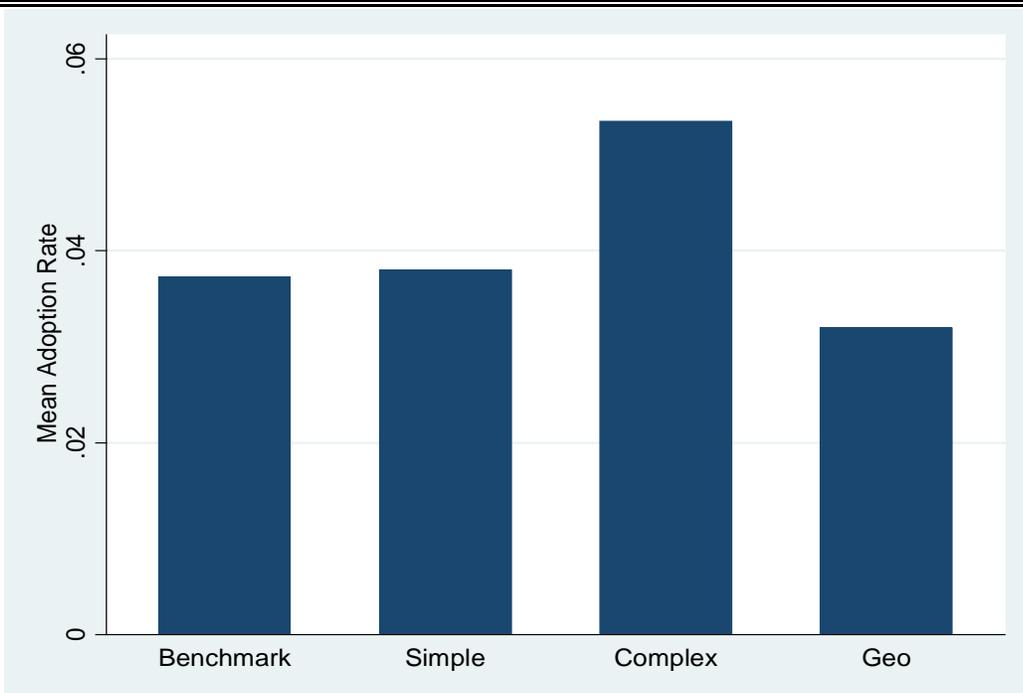


Figure A4: Simulation of SI Model by Treatment Group, Any Adoption



Rho = .1

Figure A5: Simulation of SI Model by Treatment Group, Mean Adoption Rate



Rho = .1