

Abstract

The increased adoption of fertiliser and improved seeds are key to raising land productivity in Ethiopian agriculture. However the adoption and diffusion of such technologies has been slow. We use data from the Ethiopia between 1999-2009 to examine the role of learning from extension agents versus neighbours for both improved seeds and fertiliser. We use the structure of spatial networks of farmers and panel data to identify these influences and find that while the initial impact of extension agents was high, the effect wore off, in contrast to learning from neighbours.

Keywords: social networks, adoption, diffusion.

JEL codes: C31 Q16

Raising agricultural productivity is seen as vital to economic growth in poor countries, particularly in sub-Saharan Africa where productivity growth has lagged behind other continents (Evenson and Gollin, 2003). Consequently, there has been enormous interest in replicating the Asian Green Revolution here. The focus has thus been on new technologies, particularly the adoption and diffusion of improved seed varieties and the increased use of fertiliser, supported by investments in effective extension services. Understanding how new technologies spread and how effective extension services are in this process remain important questions. The role of both extension services and learning from others have been explored in the literature but there are few studies that attempt to study them together despite the fact that learning takes place simultaneously from different sources (see Moser and Barrett (2006) for one attempt to do so). This is largely because of the difficulties in identification of impact in both cases.

We use longitudinal household data from rural Ethiopia to study the adoption of improved seed and fertiliser between 1999 and 2009. We contribute to the literature in three ways. First, we offer comparative estimates of the role of learning from extension services compared to learning from a network of peers. Second, our econometric techniques offer credible identification of both effects. In particular, we exploit recent techniques in the empirics of social networks (Bramoullé, Djebbari, and Fortin, 2009) to use the spatial distribution of farmers within villages to identify the impacts of neighbours on adoption. Furthermore, the longitudinal nature of the data allows us to control for the fixed sources of hetero-

geneity in the placement of extension services. Finally, we offer results of policy relevance for Ethiopia. We find that the adoption of fertiliser and especially of improved seeds is slow; that learning from adopting neighbours is mainly responsible for the spread of these technologies throughout this period, and that extension agents had a significant impact on adoption in 1999, but by 2004 and later by 2009, their role was almost irrelevant for the adoption process despite a vast increase in extension agents throughout rural Ethiopia.

Ethiopia offers a particularly interesting case study in this respect. The Ethiopian government has placed agricultural growth at the centre of its growth strategy. It has put forward ambitious targets to increase the use of chemical fertiliser and improved seeds in its recent development plans such as PASDEP (Plan for Accelerated and Sustainable Development to End Poverty) and the Growth and Transformation Plan (Government of Ethiopia, 2004, 2010), and spends 1% of GDP on extension services. However, as in much of sub-Saharan Africa, adoption of such technologies has been slow. Current levels of improved seed use in Ethiopia are around 5% of cereal area, which is double that in 1997/98, but is undoubtedly low. It is only for maize that adoption has increased, fourfold, to 20% , but still well below target (Central Statistical Authority (CSA) data, 1998, 2003 and 2008). Fertiliser is applied to 39% of the total land area cropped with cereals, an increase from 32% in 1997/98 but below levels attained in 2001/02 (CSA data, 1998, 2003 and 2008). Fertiliser use is about 25 kg per hectare of arable land (Gollin, 2011), although on fertilised land, application rates are close to 100kg per hectare or the recommended average application rate (CSA 2008). The key issue appears to be to get more farmers to use chemical fertiliser and improved seeds.

Low adoption is not unique to Ethiopia and the literature offers many reasons for low take up of new technologies (Feder, Just, and Zilberman, 1985; Doss et al., 2003). For Ethiopia in particular, there has been much discussion of constraints on adoption of new technologies. The supply of seed faces serious difficulties (Dercon and Hill, 2009; Davis et al., 2010), while fertiliser use faces heterogeneity in profits (Taffesse, 2008; Suri, 2011). Related to this are the high risks involved in taking up relatively expensive new technologies without insurance against harvest shortfalls (Dercon and Christiaensen, 2011). Alternative explanations such as lack of access to appropriate financial instruments (Duflo, Kremer,

and Robinson, 2011) seem unlikely in this case given the widespread availability of credit at least until 2009; this problem may become salient again as the supply of formal (government) credit has disappeared since.

Other plausible suspects are imperfect information about the returns to a new technology and the consequent importance of learning and this is our focus in this paper. Two mechanisms to overcome this are typically studied: learning from social networks of peers (social learning) and extension services. Most studies look at each mechanism separately. Foster and Rosenzweig (1995) offer a careful review of the current literature on the microeconomics of technology adoption and discuss the evidence on social learning. The literature as exemplified by Foster and Rosenzweig (1995), Conley and Udry (2010) and Bandiera and Rasul (2006) examine the role of learning from others without reference to institutional sources such as extension services¹. For instance, Conley and Udry (2010) examine this problem in Ghana where farmers learn from the experiences of others and the flows of information depend on the structure of social networks, with no access to extension services. Clearly, however, extension programmes may be an effective way to transmit information about modern inputs and encourage adoption. Across sub-Saharan Africa, the evidence of the effectiveness of the extension system is varied and disputed (Evenson, 1997; Bindlish and Evenson, 1997; Gautam and Anderson, 1999). Thus, our first contribution lies in the fact that we nest both mechanisms in one empirical model and derive comparative estimates on the role of extension and learning from peers. A similar attempt, by Moser and Barrett (2006) examines rice intensification in Madagascar using recall data to reconstruct adoption over time, and also allow for extension and local adoption to influence this decision. However, they are unable to control for placement of extension services or the endogeneity of learning in networks². In contrast, we are able to both nest extension and network effects, and to establish credible identification of the impact of both sources This is our second contribution as explained below.

We use data from a longitudinal survey of farm households, the Ethiopian Rural Household Survey, using data over three rounds covering a decade (1999, 2004 and 2009), and 15 communities across the country. Both extension services and neighbours appear to offer relevant information for adoption: for example, the data in 1999 suggest that most infor-

mation on fertiliser and seeds came from these two sources, with about half of all farmers who use fertiliser and two-thirds of all farmers who use new seed reporting that they get information from extension agents, and the rest from talking to friends or neighbours or, to a lesser extent, from observing early adopters. While this offers fertile ground to explore both effects, linking learning from neighbours and extension workers to adoption is not without econometric problems. First, the problem with identifying the learning links between peers is that peer decisions are contemporaneous and perhaps just correlated rather than influential in affecting own adoption. We use recent techniques from the empirics of network effects to address this source of endogeneity (Bramoullé, Djebbari, and Fortin, 2009). Second, extension services may target those with high potential to adopt, and their placement is therefore not random. We exploit the panel data nature of the data to control for fixed heterogeneity in the placement of extension services. Finally, as explained earlier, our analysis is in the context of large investments by the Ethiopian government in extension services, and as a result, we can offer an evaluation of its effectiveness in boosting adoption of seeds and fertiliser. In 1995, a first large expansion of the extension programme took place as part of the PADETES/NAEIP programme, aiming to reach about 9 million farmers, using the adapted T&V (Training and Visit) model. Bongor, Ayele, and Kuma (2004) and EEA/EEPRI (2006) have suggested that these extension programmes have been a mixed success. During the last five years, a further expansion of the extension programme has taken place, increasing the number of extension workers (locally called "development agents") threefold by 2008, and adapting the T&V system to reach a larger number of farmers. The most recent expansion of the services has yet to be evaluated, although Davis et al. (2010) provide a careful review of the current functioning, identifying a series of weaknesses. At present the extension system, measured in terms of the number of extension workers per farmer, is among the most intensive systems, with 600 farmers to a development agent at present, thus similar to China; in contrast, Tanzania has four times and India eight times as many farmers per extension worker (Davis et al., 2010).

The previous work on these issues stops short at 1999 which is the beginning of the expansion in extension services. Using the same data set, but only until 2004, Dercon et al. (2009) showed that access to extension agents can be linked to 7% higher consumption

growth in the subsequent period. Here, we examine the role of extension agents in the later period, between 1999 and 2009. We are thus able to examine the effects of this expansion in services over the next decade in comparison with the results in Dercon et al. (2009). We find evidence of the role of social learning throughout the decade: learning from neighbours is strongly significant, and stable throughout: an increase of one standard deviation in average adoption of improved seeds by neighbours (corresponding to local diffusion rates increasing by 22%) raises the probability of own adoption by 11% points. For 1999, the results by Dercon et al. (2009) are confirmed, in that extension services matter. But learning from extension ceases to be relevant after 1999, and despite further vast investment in extension by government in subsequent years, we cannot find any return. A recent paper by Bachewe et al. (2011), provides further evidence that is consistent with the results found here. . They find a significant impact of extension on output for the years 1994-99 and 1999-2004, but no effect on output post 2004. These results are consistent with our finding that extension services in 1999 produced the biggest effect on adoption (and hence potentially on output growth in the subsequent period) but this effect wears off for both adoption in our study and output growth in Bachewe et al. (2011).

Given low adoption, especially for seeds, this may suggest that there is a problem with the nature of extension in recent years. However, it is also consistent with a view that after an early boost from extension, adoption will largely be through social learning. Social learning appears to be mainly about farmers identifying for themselves from own and neighbours' experience whether it is profitable to adopt new seed and use fertiliser. Since rates of adoption are stagnant, this suggests that other constraints bind. For seeds, supply constraints are very likely, while for fertiliser, profitability may be low at current prices, given both limited seed supply and concerns about quality. In the next two sections we explain the empirical approach taken to identifying social learning and outline the formal econometric approach taken here. This is followed by a summary of the data. Section 4 offers the results and Section 5 concludes.

1 Empirical Strategy

The fundamental identification problem, in the estimation of peer effects is termed the *reflection problem* by Manski (1993). In a linear-in-means model, identification of peer effects depends on the functional relationship between the variables characterizing peer groups and those directly affecting group outcomes. Manski (1993) lists three effects that need to be distinguished in the analysis of peer effects: the main focus, endogenous effects, where an individual's behaviour is influenced by her peers and learning from their actions; contextual effects, or the propensity of an individual to behave as her peers do because they face similar environmental constraints and share similar characteristics and finally, correlated effects whereby peers behave like each other simply because they have similar (unobserved) characteristics or face similar shocks. The main challenges, therefore, consist in (1) disentangling *contextual* effects, and *endogenous* effects, and (2) distinguishing between *social* effects, i.e., exogenous and endogenous effects, and *correlated* effects, i.e., household in the same peer group may behave similarly because they are alike or share a common environment. A further complication are selection effects that arise when an individual chooses her own peer/reference group; this causes a bias in the estimated endogenous peer effect due to the presence of unobservables that both influence the choice of peer group and the outcome. This is the case when group formation is endogenous, for example, when households sort themselves into a locality of their choice.

1.1 Identification of peer effects in social networks

The results presented here focus on close spatial neighbours, within a kilometer of the household, as being the source of social learning. We use the distance of 1 km because it is approximately the mean distance to the plots owned by the household. Distances to plots are not available - but the time taken to plots are available and based on this, we construct a mean distance. The advantage of using spatial neighbours is that we are implicitly accounting for the fact that learning is more likely from those with similar soils and similar exposure to the vagaries of rainfall. This is especially so in a context where returns to new technologies are difficult to ascertain, and where yields are highly variable,

even within villages (see Getachew, 2011). The terrain in many of these villages is hilly and neighbourhoods within villages vary in slope and soil. Spatial neighbourhoods allow us to take account of such variation³. However, there are a number of issues of concern. These include the vexed issue of whether farmers may not acquire information more readily from very different sources, such as relatives or people they trust in other contexts. There is also the issue of the timing of decisions: we have assumed thus far that decisions are made contemporaneously but it might be more natural to take account of previous decisions made by neighbours rather than current information. We will present results that deal with these variations in the measurement of networks but below, we discuss the econometric technique we exploit to identify the endogenous effects of social networks, using the spatial network of neighbours to illustrate the method.

Our proposed strategy to disentangle endogenous and exogenous effects, relies on Bramoullé, Djebbari, and Fortin (2009) (henceforth BDF) who show that these effects can be distinguished through a specific network structure, for example the presence of intransitive triads within a network. Intransitive triads describe a structure in which individual i interacts with individual j but not with individual k whereas j and k interact.⁴ (The intuition is that individual k , in this example, is a *non-overlapping* neighbour of j , whose characteristics and behaviour can then serve to identify the impact of j on i). BDF account for correlated effects through a local or global within transformation i.e., network fixed effects⁵.

The model can be characterised as follows. Denote the set of farmers as $i \in \{1, \dots, F\}$; y_{it} denotes the outcome of farmer i at time t (here the adoption of seed or fertiliser) and x_{it} is the farmer's exogenous characteristic⁶ at time t . Each farmer has a peer group η_i of size n_i . η_i represents the farmer's local network i.e. direct connections i to other farmers in any given network l . The network l thus consists of all the connections, i.e. both the direct connections and those that are indirect, or farmers connected to farmer i only via other farmers. By assumption farmer i is excluded from η_i , i.e., $i \ni \eta_i$. We assume that our sample of size n is i.i.d. and from a population of networks with a fixed and known structure. The assumption of a fixed network structure is made on the basis that networks are defined by the location of the farmer's household. We distinguish between three types of effects: a farmer's outcome y_{it} is affected by (i) the mean outcome of her peer group

(endogenous effects), (ii) his own characteristics and (iii) the mean characteristics of his peer group \mathbf{x}_{it} (contextual effects):

$$y_{it} = \beta \frac{\sum_{j \in \eta_i} y_{jt}}{n_i} + \gamma x_{it} + \delta \frac{\sum_{j \in \eta_i} x_{jt}}{n_i} + u_{it} \quad (1)$$

Hence, β captures endogenous effects and δ contextual effects. Correlated effects are contained in u_{it} .

Turning to the estimation of Equation (1), we first construct the matrix of neighbours (alternatively interpreted as a peer interaction matrix), \mathbf{W} , which is interacted with the outcome variable and exogenous peer characteristics to form spatial lags, where the lags refer to indirect spatial neighbours. We define \mathbf{W} using a ‘K Nearest neighbours’ (KNN) characterization. KNN is a distance-based definition of neighbours where ‘K’ refers to the number of neighbours of a farmer at a specific location. Distances are computed by the Euclidean distance between GPS locations of households. Therefore, under this approach, the set of ‘neighbours’ for household i includes the K households characterized by the shortest distance to household i within each village. In the first instance, we set $K = 5$, although we restrict this set by only considering those within a maximum distance threshold of one kilometer; in other words, of those households living within one kilometer, we pick the five nearest. One of the key reasons for doing so is that this mimics the new model of extension since 2009 which targets a model farmer and his spatial network of the 5 nearest neighbours around him who are then monitored and targeted via the model farmer. This method, using a 1 km. radius seems sensible empirically as well, since in practice, only 1 percent of neighbours lived farther away. Using this method, we drop all such households that are not a nearest neighbour to any other household in the sample.⁷ (Note that alternative definitions of K are possible, even desirable and we discuss these in the next section - however, for the sake of simplicity we confine ourselves to the 5 nearest neighbours here). Depending on the number of nearest neighbours used in our definition of \mathbf{W} , this leads us to drop a small number of households which causes slight variations in the sample size across specifications.

Under the assumption that households are a random sample of the underlying population,

dropping such ‘island’ households does not bias our results. We row normalize \mathbf{W} so that $\mathbf{W}_i \mathbf{y}_t$ represents the average outcome of the agent’s peer group excluding herself i.e. it is the same as $\frac{\sum_{j \in \eta_{it}} y_{jt}}{n_{it}}$.

Equation (1) can be now written in structural form as:

$$y_{it} = \beta \mathbf{W}_i \mathbf{y}_t + \gamma x_{it} + \delta \mathbf{W}_i \mathbf{x}_t + \zeta_t + u_{it} \quad (2)$$

where $\mathbf{W}_i \mathbf{y}_t$ represent the endogenous peer effect and $\mathbf{W}_i \mathbf{x}_t$ represents the contextual effects. Note, that we allow for intra-group variations in social interactions which are asymmetric in general since farmers are attached in varying ways to their peers⁹. The nonlinearity introduced by these asymmetric interactions provides necessary conditions for identification. This is because our chosen peer interaction structure (\mathbf{W}) induces variation in the magnitude of social interactions such that each farmer has a unique and different set of peers/neighbours. Moreover the variation in the number of indirect neighbours that results due to this assymetry of connections allows us to use the non-overlapping neighbours to identify the parameters.

The reduced form of Equation (2) is given by;

$$y_{it} = (I - \beta \mathbf{W}_i)^{-1} (\gamma I + \delta \mathbf{W}_i) \mathbf{x}_t + (I - \beta \mathbf{W}_i)^{-1} \mathbf{u}_t \quad (3)$$

Denoting the variance-covariance matrix of \mathbf{u}_t as $\psi_{\mathbf{u}_t}$, it is easy to see that,

$$E[(\mathbf{W}_i \mathbf{y}_t) \mathbf{u}_t'] = \mathbf{W}_i (I - \beta \mathbf{W}_i)^{-1} \psi_{\mathbf{u}_t} \neq 0 \quad (4)$$

Kelejian and Prucha (1998, 2001) suggest the use of a spatial two-stage least squares estimator (S2SLS), with a set of instrument matrices to instrument for $\mathbf{W}_i \mathbf{y}_t$, which also avoids computation accuracy problems in the ML approach from Equation (4), we can see that, ideally the set of instruments contains linearly independent columns of $[\mathbf{W}_i^2 \mathbf{x}_t, \mathbf{W}_i^3 \mathbf{x}_t, \mathbf{W}_i^4 \mathbf{x}_t \dots]$. The use of such instruments is possible when the matrices, \mathbf{I} , \mathbf{W} and \mathbf{W}^2 are linearly independent. This is easily violated when groups are all of similar size and everyone within a group is connected to everyone else. In this case \mathbf{I} , \mathbf{W} and \mathbf{W}^2 are linearly dependent and

$W = W \cdot W^2$ cannot be then used as an instrument.

In the case of (spatial) networks as here, identification is achieved if the network is characterized by a small degree of intransitivity e.g., farmer i is connected to farmer j and farmer j is connected to farmer k , but farmer i and farmer k are *not* connected. This produces a directed network topology which achieves identification of peer effects as shown by BDF. The networks-based intuition of this strategy is straightforward: $W^2 \mathbf{x}_t$ is an identifying instrument for $W \mathbf{y}_t$, since x_{it} affects y_{jt} (since they are connected and interact with each other) but x_{kt} can only affect y_{it} indirectly, through its effect on y_{jt} . In our particular case, the relevant instruments are then $W^2 x_{it}$, an $n \times 1$ vector of weighted averages of adoption of the neighbours of neighbours of each farmer in the village. By definition, these neighbours of neighbours are part of the overall network, but not overlapping with the direct peer group, whose effects is being identified.

While this will identify the endogenous effects, there is still an issue of correlated effects and of selection effects. In this paper, following Blume and Durlauf (2005), we employ a first-differenced specification to address the issue of correlated and selection effects. We employ differences between the three available rounds of data to account for unobservables that are constant over time. Accounting for such unobservables appears to be important in light of a large body of work suggesting that peer effect estimates are biased due to the presence of unobserved household characteristics (Evans, Oates, and Schwab, 1992). The period-difference will therefore eliminate this unobserved farmer fixed effect that could bias the peer interaction effect.

We are interested in explaining the change in adoption achieved by households between the three survey rounds. We write the change in a farmer's adoption take-up as a function of the change in a farmer's own characteristics whilst allowing for peer effects by incorporating spatial lag terms of the dependent variable. Hence, we rewrite Equation (1) as;

$$\Delta y_i = \beta \frac{\sum_{j \in \eta_i} \Delta y_j}{n_i} + \gamma \Delta x_i + \delta \frac{\sum_{j \in \eta_i} \Delta x_j}{n_i} + \Delta u_i \quad (5)$$

where $\Delta y_{it} = y_{it} - y_{it-1}$ denotes the difference in adoption levels between periods t and $t - 1$ for farmer i . $\frac{\sum_{j \in \eta_i} \Delta y_j}{n_i}$ denotes the change in farmer i 's peers' adoption status between

t and $t - 1$ and $\Delta x_{it} = x_{it} - x_{it-1}$ denotes the change in farmer i 's own characteristics while $\frac{\sum_{j \in \eta_i} \Delta x_j}{n_i}$ denotes the change in household i 's peers' characteristics between t and $t - 1$. This can be easily seen in terms of the network specification,

$$\Delta y_i = \beta \mathbf{W}_i \Delta \mathbf{y} + \gamma \Delta x_i + \delta \mathbf{W}_i \Delta \mathbf{x} + \Delta u_i \quad (6)$$

However, while we are able to difference out all the household and village level fixed effects that are constant over time, correlated effects will still matter if there are common environment related time-varying unobservables that effect both the farmers as well as their neighbour's outcomes. In the context of adoption, prices for inputs and outputs are an obvious example. We address this by including village fixed effects in our first differenced specification¹⁰. In the first differences specification the village fixed effects serve to absorb all the omitted variables at the village level that are correlated with the changes in both own and neighbour adoption.

$$\Delta y_i = \beta \mathbf{W}_i \Delta \mathbf{y} + \gamma \Delta x_i + \delta \mathbf{W}_i \Delta \mathbf{x} + \phi_v + \Delta u_i \quad (7)$$

where ϕ_v denotes indicators for the village, v , that each household belongs to. Note that our spatially motivated construction of the network implies that most peer groups are restricted to lie within villages limiting the possibility of across village interactions. Therefore we assume that conditional on village fixed effects there is strict exogeneity of \mathbf{x}_{it} with respect to u_{it} .

We estimate Equation 7 using two stage least squares. To get identifying power we use a two step method: we first regress the outcome variable on the entire set of exogenous characteristics; based on the parameters of this regression we then predict the outcome and take higher order spatial lags of this predicted outcome. These predicted values serve as our set of instruments which subsequently used to predict the endogenous variable of interest – share of neighbours adopting in each time period.

1.2 *Further empirical issues*

The identification of neighbours' adoption decisions is obtained here by using the non-overlapping sets of neighbours - or neighbours of neighbours, who can be thought of as affecting the decisions of spatial neighbours directly - but not the household's own decision. Testing the exclusion restriction is of course not directly possible; however it would appear reasonable in a spatial setting in which observing neighbours' plots matters for observing returns that one's own decision to adopt is only influenced by the neighbours of neighbours via one's direct neighbours.

Restricting the set of 5 neighbours to be within 1 km, obtains an average distance amongst them of 295 meters. We also investigate whether there are any terrain differences amongst the set of neighbours, which may impede interactions. The intra-class correlation of elevation amongst the nearest-neighbours peer groups is 0.93 (statistically significant at the 1% level). The standard deviation of elevation within nearest-neighbours peer groups is a negligible 126 meters given that the mean elevation of sample households is approximately 2000 meters. Therefore, given the high degree of terrain/elevation homogeneity within peer groups, we are confident that our Euclidean distance-based measure of neighbours is able to capture the spatial proximity of farmers.

To ensure this, we check whether there are any terrain differences amongst the set of neighbours, which may impede interactions. The intra-class correlation of elevation amongst the nearest-neighbours peer groups is 0.93. The standard deviation of elevation within nearest-neighbours peer groups is a negligible 126 meters given that the mean elevation of sample households is approximately 2000 meters. Given the high degree of terrain/elevation homogeneity within peer groups, we are confident that our Euclidean distance-based measure of neighbours is able to capture the spatial proximity of farmers.

A key issue is identifying the appropriate peer group. Neighbourhoods are not particularly small, and geographical distance is bound to matter for the extent to which farmers can observe returns to new inputs. One potential problem is that plots may be scattered, which weakens the argument for using spatial neighbours. However, in the villages studied, plots are largely consolidated or nearby, allowing us to ignore this issue.¹¹

Clearly, alternative definitions of spatial proximity might matter and we exploit two variations. We use self-reported identification of neighbourhoods (hamlets) within villages (PAs) in which the radial distance between households is 2.4 kms. As noted by BDF, peer effects are still identified since households interact in village based groups of different sizes so that the peer/neighbourhood interaction matrix, W , has block diagonal elements of varying sizes that ensures identification. A further variant as follows: two farmers A and B might reside in adjoining hamlets and hence be (spatially) close to each other even if in different hamlets, while a (non-overlapping) neighbour C might reside in the same hamlet as B but be quite far away (say over 1 km away) from either of them. In this case, we can use farmer C to identify the effect of farmer B on farmer A. We find that the results from these alternative methods are remarkably similar to the simple radial measure of 1km.

While this discussion handles the identification of social learning, we also aim to contrast it with the impact of extension visits. We use farm household fixed effects to address the potential difficulty that visits of extension agents might be related to unobservable farm household characteristics. Furthermore, the inclusion of village fixed effects in the first-differenced equation will capture trends in the village-level placement of extension as well. All estimates include (time-varying) controls as well for other farmer characteristics, such as wealth and educational levels of the household. Finally, even though data on adoption are not available at the plot level, we have information on (self-reported) quality of plots, controlling further for a possible source of targeting by extension workers and demand for modern inputs.

In the results reported below, results are shown for the basic 'five (or fewer) neighbours within 1 km' definition of peer groups¹². We report the IV cross-section results, as well as the IV first-difference results. For further robustness, we also estimate all equations accounting lagged adoption by the farmer and lagged adoption by the peer group.

2 Data and Descriptives

The data are from the Ethiopian Rural Household Survey, and in particular its rounds 5, 6 and 7, i.e. 1999, 2004 and 2009. These rounds are particularly suitable as they have

details on extension and modern input adoption, with improved seeds for crops such as wheat and maize only becoming more systematically available since 2000. This survey has been running since 1994, covering 19 Peasant Associations (PA) across the four main regions. While the sample is not nationally representative (it does not include pastoral households or urban areas for instance) Dercon et al. (2009), show that the survey is broadly representative for the diversity of the main crop farming systems in the country in population shares and attrition is low at 1-2 percent per year. Further details of the survey, including a description of farm characteristics of each sample site can be found in Dercon et al. (2009).

Additionally, in this paper, we focus only on those sample households who are present throughout the three rounds of survey and who are involved in cereal production, making 87% of the sample. We also drop those sample households for whom GPS information was missing¹³.

Table 1 offers summary evidence on the importance of the different sources of information, obtained from the 1999 round of the Ethiopian Rural Household Survey, concentrating on those households that grow cereals. It suggests that while both neighbours and extension agents are important in transmitting information, extension agents were the primary source of information for both new seed and fertiliser in 1999.

Table 2 offers summary evidence on the average adoption rates in the three years¹⁴. The adoption rate stayed much the same across years: for new seed, rising slightly from 18% in 1999 to 23% in 2009; for fertiliser, growing from 62% in 1999 to 64% in 2009, with a sharp dip in 2004 to 25%. It should be noted that in 2002-3 there was a serious drought and hence 2004 represents a sharp response to this: there was a significant drop in the percentage of farmers using fertiliser. In this year, we are unable to pin down the use of improved seed accurately: we are able only to obtain whether seed was purchased (which includes bought local seed) and this share is far higher at 31%. Admittedly, local seed is more likely to be saved than bought, while the contrary is true for improved seed and this pattern can be seen in both 1999 and 2009.

It should be noted that the use of improved seed is difficult to establish with precision.

The questions posed in both 1999 and 2009 ask whether the farmer uses local or improved varieties - and within each of these, whether the seed is saved or bought (or also exchanged in the case of local varieties). The measure used here is that of such self-reported use of improved seed, whether saved or bought. However, in 2004, this question was simply phrased as whether any seed was bought and consequently, the measure of seed use here includes all seed bought including local, non-improved varieties. A further complication is that the use of improved seed varies by crop: while improved seed can be saved for use in the case of wheat and teff for instance, improved seed for maize cannot be saved thus. For maize (as for rice, millet and sorghum), the seeds are obtained as hybrids and are effectively incapable of regeneration the following year. It is also thought that returns from improved seed are much enhanced if used with fertiliser and most households do use fertiliser with new seed, with 96% doing so in 1999 and about 82% in 2009. In 2004, given the sharp fall in fertiliser use, only 9% used (bought) seed with fertiliser. Given the measurement error in the use of improved seed in 2004, we present the cross-section results by year and also the change in use between 2009 and 1999 as well as the comparison across all three years for seed alone. The figures for fertiliser are less prone to such error for the questions were asked in a similar fashion across all three years and the recorded use of DAP and Urea is easier to establish.

We also note that these figures are higher than national averages, suggesting we have areas with higher potential for agriculture in the sample than on average in the country¹⁵. At the same time, these figures are well below what potentially could be obtained as we are focusing on farmers involved in cereal production. Seed adoption seemed to respond slightly more to neighbour's use of seed, relative to visits by extension agents with the correlation between own and neighbours' adoption higher in both years, at 0.47 in 1999 and 0.29 in 2009. Note also the striking increase in the average number of extension visits, going from about 0.3 to 5.5 visits per farmer, reflecting the vast expansion of the supply of development agents or extension agents in this period. In sum, adoption of seed has increased very little over the decade and the use of new seed remains rather low but the use of fertiliser has remained relatively high and steady. These figures rely on panel data - and it might well be the case that with heterogenous returns, only those farmers who expect to profit take this

up in the first instance and hence it is unsurprising to see little change.

However, there is a lot of churning which the seeming stability of figures disguises as is evident in figures 1 and 2. The figures are a sequence index plot in which the y-axis represents the ID number of each household, ordered by their (1999) adoption status, while the x-axis plots the year of adoption. Each horizontal line represents a household/farmer and the colour of the line (from light grey to dark grey) changes according to the adoption status at each survey round from year 1999 to year 2009. In this way, adoption histories of all farmers can be represented and major trajectories identified. In particular, the figures shows the divergence in adoption trajectories and how some farmers continue to use new seed or use fertiliser over time. Note that only 7.5% of new seed users (as opposed to 45.5% of fertiliser users) continued to use new seed in 2009, once taken up in 1999. Only 4% of farmers continued to use new seeds and (19% of fertiliser users) through the period, once adopted in 1999. Overall the figures illustrate the complex dynamics in the take-up of seed and fertiliser adoption amongst farmers in Ethiopia.

Clearly, such churning demands explanation. A first step is to examine the characteristics of adopters and non-adopters in each period: are there clear correlates of adoption? Table 3 offers a summary of the differences in characteristics between adopters and non-adopters of new seed. The main difference, in 1999, between adopters and non-adopters of in the first year is the visit of extension services over the previous 5 years, where over two thirds of adopters report being visited at least once relative to about 5% of non-adopters. This difference falls in 2009, with 47% of non-adopters being visited. Recall that this is in a context of a vast increase in the numbers of extension agents and corresponding visits as shown in table 2. Adopters in all three years are slightly more educated, have slightly better quality land but do not differ significantly in terms of assets measured as livestock. However, the key difference across all years appears to be that adopters are more likely to have neighbours who are adopters too.

Fertiliser adoption is confined to the wealthy farmers (table 4). They also have more and slightly better land and are better educated, with the differences narrowing by 2009. Again, the visit of extension agents seems to define the adopters particularly in 1999 - but again, the key difference across the decade is that adopters had neighbours who also adopt. This

sets the stage for the results in the regressions below, where we control for the endogeneity of the decision of neighbours to adopt new seeds or use fertiliser.

3 Results: Neighbours' adoption and extension agents

The sample is restricted to cereal farmers and households with identified GPS locations. We use a sample of 954 households across the three years for whom we have consistent panel data. Recall that we construct spatial neighbours based on a distance of 1 km from the household¹⁶. We instrument for the average neighbour's decision to adopt by using the non-overlapping sets of neighbours - or neighbours of neighbours, who can be thought of as affecting the decisions of spatial neighbours directly - but not the household's own decision. It might be argued that extension visits¹⁷ are also endogenous and ought to be instrumented for. However, in this context, it appears that village level variables (distance to the nearest extension office) are critical in explaining extension visits¹⁸. We take the view in what follows that the visit of extension agents can be regarded as largely exogenous and control for both own characteristics (in the form of all these variables) and village-level fixed effects. In addition, we also present the results in first differences: these in turn allow us to look at the robustness of these results in the presence of unobserved fixed factors at the household level that might bias the estimates for each year, including those linked to the placement of extension services.

The estimates below are based on the specification given in Equation 1. For clarity, we reproduce it below with the the names of the variables used:

$$y_{itk} = \alpha + \beta_1(\text{Extension visits})_{itk} + \beta_2(\text{share of neighbours adopting})_{itk} + \gamma X_{itk} + v_i + \epsilon_{itk}$$

where: y_{itk} is a discrete variable denoting whether household i , adopted technology k at time t , X_{itk} denotes a vector of individual and household characteristics, including the characteristics of the plots on which cereals are grown and v_{itk} denotes village level fixed effects. The first-differenced specification retains the village fixed effects to account for different trends in placement of extension services at the village level. The share of (direct) neighbours who also adopt the technology is instrumented for using a first-stage regression

where their average decision to adopt is predicted using their own characteristics and the share of their direct but non-overlapping neighbours who adopt new technology. The full specification of these regressions is available on request but we offer summary tables that focus on the impacts of the key variables below.

3.1 *Peer Effects in Adoption*

We begin with estimates of the likelihood of adoption using each year as a cross-section and abstracting from problems of selection and placement. Table 5 presents the uninstrumented and instrumented probit estimates of the effects of extension services and neighbours' adoption decisions on one's own probability of adopting new seeds in both 1999 and 2009. All estimates control for a wide variety of household and farm level variables (including land, land quality, livestock, household composition, education, attitudes to risk, whether identified as a model farmer in previous decade) and community fixed effects. The results are reported as marginal effects, i.e., the impact of each variable on the probability to adopt evaluated at the average of all variables. Cragg-Donald F-test statistics are offered, and throughout we can reject the hypothesis the decisions of the neighbours' neighbours are weak instruments.

The results suggest that there is a strong relationship between the adoption decisions of neighbours and one's own decision to adopt new seed, with a strong and significant coefficient of approximately 0.46 in both 1999 and 2009 (with a slightly higher estimate of 0.68 in 2004) in the IV regressions. In brief, an increase of one standard deviation in the average neighbours' adoption raises the probability of own adoption by about 11% in 1999, by 19% in 2004 and 12% in 2009. Average adoption rates range from 0.18-0.23, so this is large - more than double current levels. An increase of one standard deviation in extension visits (by 1.3 visits in 1999) raises the probability of own adoption by 3.7%, falling to 1.3% in 2004; while in 2009 this effect is at 2.9% (where 1 sd which is now 10 visits). These results also show that there is a clear collapse in the return to one extra visit, for those not yet visited: the increased probability of adopting in 2009 is only one-tenth what it was in 1999. Clearly the impact of neighbours' decisions drowns out any impact through

extension. These effects correct for endogeneity - but might still be contaminated by the changing environment over time that is unaccounted for in each regression. To examine the robustness of these estimates, we estimate a regression differenced between the three periods (this time using a linear probability model) and examining the robustness of the basic specification to including controls for previous years and lagged adoption. The estimates are presented in table 6. For robustness we also present household fixed effects estimates in Column (3), and controlling for survey-round specific village fixed effects (Column (4)). These results confirm the importance of neighbours' adoption on own adoption, with results suggesting much higher impacts of neighbour adoption, controlling for household fixed effects. The effects of neighbour adoption are stable at 0.9 and as column three indicates, the impact of previous rounds is negligible. This is distinct from the effect of extension visits: here, the effect is large in 1999, but collapses in 2004 and 2009. In fact, the low and significant average effect of 0.003 is the direct consequence of this pattern. All effects are marginal effects at the mean from a probit model, while in table 6, they are from a linear regression model. Hence, a comparison between table 6 and table 5 is only suggestive (and valid around the mean of all variables). Nevertheless, it is striking that they mimic the findings in table 5 for 2009, including that of the impact of extension which is very small but significant. This suggests that by 2009, extension services are widespread, and perhaps not targeting particular farmers. The fact that the coefficient is only one-tenth of the effect in 1999 in table 5 suggests that initially farmers more likely to adopt were targeted. In any case, they confirm the result for 2009: the impact of extension is small and the role of neighbours' decisions is relatively strong.

The next two tables, tables 7 and 8 display the results of a similar analysis for the use of fertiliser over time.

The results tell us that a one standard deviation increase in the average fertiliser adoption of neighbours (0.35) raises own probabilities of adoption of fertiliser by 19%, in both 1999 and 2009¹⁹. This is a substantial effect given that adoption is already at 62% in the survey areas. The effect of extension visits in 1999 is large and significant. The impact of an extra extension visit in 1999 (1 sd) is to add 22% to the probability of adopting fertiliser. But by 2009, and in a pattern similar to seed adoption, the effect becomes negligible and

insignificant. Table 8 offers the results using first differences, thereby controlling for household fixed effects. As before we also report household fixed effects estimates (Column (3)) controlling for survey-round specific village fixed effects (Column (4)). Again, the results look more like the 2009 effects than the 1999 effects, with a collapse of the extension coefficient. It is likely that in 1999 extension agents targeted farmers who were likely to adopt fertiliser. The 'true' impact of extension services on the typical farmer was small and possibly negligible after all. The impact of neighbours adopting is again high as compared to the cross-section results in table 7, but significant and as large as the effects for the adoption of seed. The impact of a one standard deviation increase (0.25) is about 10%.

A final question centres around the impact of extension through diffusion: given that extension visits start a cycle of learning, should not a proper assessment of the impact ought to include the indirect effects that such initial impacts generate? To examine this, we use the initial results from table 5 and simulate the impact on learning, accounting for fixed effects, using a simple adaptive model of learning (hog-cycle model). Figure 3 plots the results of the simulation. First, note that in the absence of either extension or learning from neighbours, adoption is determined entirely by own, fixed characteristics, which implies an adoption rate of 1.3%. Sans any learning from extension but allowing learning from neighbours implies a long run equilibrium level of adoption of 17%. Further adoption requires an injection from other sources: the initial level of average visits (0.27), extension visits, while potent, adds an extra percentage point to the initial level of 1.3%. Of course, it gets a learning cycle going via social learning, but it would take many iterations before we would get higher adoption – after 10 iterations, only 6% extra is added to the initial 17%, with a long set of iterations taking us to the equilibrium of an additional 14% impact. The biggest part of the variation is explained by learning – the independent effect of extension is really small in terms of economic significance. But this is because the number of visits is small. With the same learning technology and same marginal effect of extension, boosting extension visits to 1.06 on average (as in 2004) would have a substantial impact. Here, after one iteration, 4 percent would have been added, and after 10 iterations half the sample would be adopting, or 32% more adoption added on. At levels of 5.5 visits, (the average in 2009), one would get full adoption after only 4 iterations!

But the regressions also show that this did not happen. By 2004, the marginal return to extension had dropped off and so there was no independent source of boosting adoption. The learning cycle was very slow with only 1% take-up added after a year. We know that this round is less reliable. But by 2009, with massive boosting of visits, we get the same low outcome; the boost to 5.5 visits means that something is added in each iteration, but now after 10 iterations, it would only have boosted overall adoption by about 6%, barely different from the adoption rate from 1999, with the crucial difference that the results obtained in 1999 were achieved with far fewer but seemingly more efficient extension visits. In short, the return from the expansion of extension has had no impact on the speed of adoption, and adoption is still dramatically low. Furthermore, though social learning is crucial, it is also not high enough to sustain itself either.

In summary, we find that in the initial period both neighbours and extension agents mattered for adoption of new seed - but that the impact of adoption by neighbours is about three times as high, with an increase of one standard deviation in average adoption of improved seeds by neighbours. This is large, corresponding to local diffusion rates increasing by 22% and raising the probability of own adoption by 11% points. While the impact of raising the number of visits by 1 standard deviation (1.3 more visits) is about 4% points. In 2009, these impacts are similar for improved seeds. However, the impact of extension services by 2009 fell to a return in uptake of modern seeds per visit of only one-tenth of what it was in 1999, as there are far more extension visits in 2009 compared to 1999 (with a mean number of 0.3 in 1999 and a mean of 5.5 visits in 2009), so that one standard deviation corresponds to 10 more visits. The impact on adoption of fertiliser is mixed, with a large impact of extension agents in the initial period and a substantial impact of neighbours. By 2009, both wear off, but for diffusion via neighbours, this appears to be a problem of precision of estimation, while for extension, the return per visit collapses to near zero. We also ask if these effects can be robustly identified in impacts over time, focusing on *changes* in technology use. While this allows us to control for time-invariant factors, including the extent to which placement of extension services is driven by concerns about current yield, this analysis comes with the qualification that identification may not be easy as growth in adoption of both seed and fertiliser has been rather low over the decade. Nevertheless, we

confirm our earlier results: we find that neighbours seem to matter for adoption of both seed and fertiliser but the effect of extension agents is non-existent. This suggests that our results are robust, pointing evidence of social learning rather than any sharp impact of extension workers in this period.

3.2 Robustness: Alternative social networks and timing of adoption

As mentioned earlier, there are a number of issues of concern. These include the issue of whether farmers may not acquire information more readily from very different sources, such as relatives or people they trust in other contexts. There is also the vexed issue of the timing of decisions: we have assumed thus far that decisions are made contemporaneously but it might be more natural to take account of previous decisions made by neighbours (lagged information) rather than current information. We explore these issues below.

The supplementary tables A.1 and A.2 explore the possibility that proximity does not define the relevant peer groups for farmer interaction. Here, we exploit a network survey carried out in Round 6 of the ERHS. This survey identified for each farmer the relevant networks involved in labour sharing, credit and relatives and mutual insurance. These networks overlay substantially with our measure of spatial proximity. We construct an alternative (weighted) network, by weighting equally the type of link shared between any given pair of farmers (including spatial) and estimated the effects of peer interaction arising from this weighted network. The tables show that coefficients increase slightly but remain positive and significant throughout. Therefore if anything, the results obtained in our paper using spatial networks are lower bound estimates of the actual peer effect.

Thus far, we have assumed that the contemporaneous decisions of neighbours affect own decisions to adopt new seed or use fertiliser. Given the enormous variation in rainfall season after season and the particular vagaries of yield variation even within small neighbourhoods, this is a natural starting point. We have assumed therefore that farmers engage in discussions about fertiliser use and seed adoption at the beginning of each season and our attempt to deal with the endogeneity (and reflection effects) is based on this notion. It may well be the case however that farmers base their decisions upon observation of the decisions

of neighbours in past seasons and if so, this is easier to pin down since it would imply using lagged information of the decisions of peers in making one's own decision. We examine whether adoption decisions by neighbours in the previous season affect own adoption in the next season with the proviso that crops and thus adoption decisions differ between seasons and hence the impact of lagged adoption is based on a much smaller sample of farmers who grow the same crops in both seasons. The results are presented in the supplementary table 11. The inclusion of own adoption makes little difference here, even when significant as in take-up of fertiliser.

3.3 Neighbours and extension visits: Graphical analysis

As the proportion of farmers who have adopted new technologies increases, the impact of additional extension visits and neighbors' experience might be expected to fall. For instance, in a community where half of all farmers are using fertilizer, a non-adopter gains very little information when a neighbor adopts.

To assess the relationship between community diffusion levels and the impact of extension visits, we use the regression results in tables 7 and 8 to ask how an individual farmer's probability of seed adoption and fertilizer use is related to the initial diffusion level of the technology. In other words, we calculate the marginal effects of the diffusion level on the probability of farmer adoption. Figure 4 shows how a 10-percent change in the current level of neighborhood diffusion will affect an individual farmer's probability of adoption; it then plots this marginal effect across different levels of diffusion. The figure shows that this neighborhood spillover increases until local diffusion reaches a level of about 70 percent, beyond which the effect falls.. For fertilizer, the effect of local diffusion peaks at around 30 percent. In both cases, these effects are relevant in size: an increase by 10 percent in diffusion in the neighborhood increases the probability of adopting by about 5 percent for individual farmers, at current levels of diffusion in these villages for seeds and fertilizers. (The figures for 2004 are an anomaly here.)

Parallel results for extension are presented in figure 5, which plots the marginal effects of additional extension visits against different levels of extension activity. It shows that the

marginal effects of additional extension visits fell close to zero by 2009, as some farmers had effectively reached saturation with extension learning. Extension visits had a positive association with fertilizer use in 1999 due in part to deliberate selection of the farmers receiving visits, or to the roles of extension agents in supplying fertilizer and seeds. But by 2009, the average number of extension visits per farmer had increased dramatically (from 0.3 per farmer in 1999 to 5.5 in 2009, averaged across the sample), yet the marginal effect had fallen close to zero²⁰.

4 Conclusions

This paper contributes to the literature on learning and technology adoption in agriculture by examining and identifying the impact of both learning from extension and learning from neighbours using a combination of panel data and exploiting recent techniques in the identification of peer effects in social networks. We do so in a setting where the investment in extension services has increased largely through an expansion in the number of agents over time and where the explicit aim was to increase the take-up of new seed and raise use of fertiliser.

The traditional explanation for the observed differences in the adoption of new technology is heterogeneity in characteristics - some farmers are simply more receptive or entrepreneurial than others. More recent explanations centre around the notion that returns are both heterogeneous and uncertain, making social learning important for the gradual adoption of new technology even in a homogeneous population. In this study we find evidence that social learning is a powerful force for adoption of new technologies and is far more persistent than learning from extension services in this period. We find that the returns to extension may have been high in Ethiopia in 1999, but by 2009 they appear to have collapsed to very low levels. The extension model of this period, and intensity of visits may transmit useful information to the farmers, but as a model to encourage modern input adoption, it does not appear to be very effective. This is not inconsistent with the general evidence on extension which suggests that extension services have an important role in raising awareness in the early stages of adoption but the impact on diffusion falls over time. In

brief, it is clear that the simple expansion of extension services that has been seen recently, (at a cost of over 1% of GDP) has simply not paid off.

This conclusion is important for policy makers and offers further credence to evidence of offering other approaches to evaluating extension services both in Ethiopia and elsewhere in sub-Saharan Africa. Davis (2008) offers an overview of the evidence and suggests that the impact of extension services has been mixed. Other evaluations cited there suggest that while the Ethiopia's Participatory Demonstration and Training Extension System (PADETES), based on Sasakawa Global 2000's (SG-2000) approach to extension did raise adoption initially, farmers also stopped using new seed and fertiliser packages (see Bonger, Ayele, and Kuma (2004)). Spielman, Kelemwork, and Alemu (2011) summarise four recent studies on the impact of extension services and conclude that: "Nonetheless, the entire body of evidence on agricultural extension suggests that the impact on productivity and poverty has been a mixed experience to date. Although many farmers seem to have adopted the packages promoted by the extension system, up to a third of the farmers who have tried a package had discontinued its use (Bonger, Ayele, and Kuma, 2004; EEA/EEPRI, 2006). Indeed, Bonger, Ayele, and Kuma (2004) also find that poor extension services were ranked as the top reason for non-adoption."

The expansion of extension services has been an important plank in the agricultural strategy of the Ethiopian government over the past decade. Since 2011, they have attempted to restructure the role of extension services and the new model appears to concentrate on targeting a farmer and his closest spatial neighbours - which is a mirror of the identification strategy that we have pursued here. There is also an intent to identify extension packages that are more specific to their settings and attempt to transmit information that covers a range of management practices. The evidence in the current paper suggests that this is likely to be a move in the right direction but the impact of this strategy is yet to be seen.

Notes

¹Foster and Rosenzweig (1995) examine the adoption of high-yielding varieties in India during the Green Revolution. They find that imperfect knowledge about the management of the new seeds was a significant barrier to adoption; this barrier diminished as farmers increased their use of the new seed and watched their neighbours' experience with HYVs. Conley and Udry (2010) examine pineapple cultivation in Ghana. They find that farmers do learn (about optimal input use: in particular, the use of fertiliser) from their neighbours in social networks.

²Gautam and Anderson (1999), for example, conclude that early studies overstated the impact of extension and pinning down its impact involves difficult issues of attribution and identification; they concluded that the data for Kenya simply do not suggest a discernible impact. They argue that panel data are required to allow more accurate identification.

³A recent study of model farmers and their neighbours found large differences which are not readily attributable to observable factors. The evolution of land fertility offers one of the factors which farmers find hard to handle. For example, in 1999, a third of farmers in the ERHS found that yields are stable, but 58% reported declining yields while 10% reported increasing yields. With limited experience of modern inputs and changing land fertility, information about new technologies is hard to be sure about, and could make learning about new technologies difficult and adoption a slow process.

⁴This particular network structure produces exclusion restrictions which achieve identification in the same way as exclusion restrictions achieve identification in a system of simultaneous equations.

⁵In a local transformation the model is written as a deviation from the mean equation of the individual's peers and in a global transformation it is written as a deviation from an individual's network. Note that in the presence of correlated effects, the distance between individuals within the network needs to be ≥ 3 . Distance in this context is defined as the shortest directed path between two nodes in a given network.

⁶For ease of notation, in this section, we represent only one exogenous characteristic but the empirics take into account many exogenous characteristics that are described later.

⁷In fact, for such 'island' households, column sums of the spatial weight matrix \mathbf{W} are zero.

⁸ \mathbf{W}_i is the i^{th} row of the $n \times n$ matrix \mathbf{W} . When post multiplied by \mathbf{y}_i whose dimension is $n \times 1$, it produces a 1×1 firm specific peer average.

⁹This is in contrast to those studies that use the entire reference group such as a village, where it is assumed that individuals within the peer group are all fully connected and have the same level of social interactions; variation in social interaction in this case are brought about only due to across group variation

¹⁰We also use a fixed-effects specification where instead of differencing the model we account for and estimate separately, the household fixed effects. The drawback of this approach is that we are unable to explicitly account

of village effects since they get absorbed in the household fixed effects. We report results from the fixed-effects specification together with that from the first-differenced model.

¹¹This may reflect the fact that after 1976, land reform had ensured that all land is owned by the state and allocated by the peasant association, with an aim to improving productivity by reducing the farming of scattered plots.

¹²*Overlap Statistics:* To see how close the neighbours' peer group is to the neighbourhood group we compute distance based summary measures for neighbourhood peer groups. The following statistics are averaged across all villages. The mean distance between any two households in a neighbourhood is 2:04 km with a standard deviation of 1:67 km. The maximum and minimum distance in a neighbourhood are 5:60 km and 0:11 km respectively. To see the extent of overlap, we note that the spatial peer group imposed a cutoff of 1 km. In this case, the average percentage of household pairs within neighbourhoods, with a distance within the 1 km radius is 67%. The average number of neighbours within this definition of neighbourhoods is about 6.

¹³However the proportion of such households in our sample is quite low (54 households or 5% of the sample). A simple t-test for differences in basic demographics revealed no statistically significant differences between households with and without GPS data, at least in terms of plot soil type (lemshare, meddashare) plot/land area, amount of irrigated land, household head education, livestock value and total oxen.

¹⁴A companion piece by Getachew (2011), using these data examines the role of new seeds and fertiliser in yield growth. Cereal yield grew by 21% over the decade (lower than the national average) while input use is far higher than the national average. The paper finds that there is a significant response of yield to the use of improved seed and fertiliser.

¹⁵The figure A.1 , in the supplementary figures shows the location of the villages, adoption rates and extension visits. The number of extension visits in the village is shown by the size of the circle; green circles show the number of non-adopters and red circles show the number of adopters.

¹⁶Note that two alternative definitions of neighbours based on self-reported neighbourhoods as well as the overlap between distance and self-reported neighbourhoods was used. The estimators here rely on the fact that such spatial neighbours vary in number (as opposed to the simpler definition using the five closest neighbours). However, estimates remain unaffected by such considerations.

¹⁷Evenson and Mwabu (2001) argue that this is a relevant measure because this variable "captures both agriculture-specific human capital embodied in extension workers as well as the amount of it that the extension workers transmit to farm people." (Evenson and Mwabu (2001, 5)).

¹⁸Farmers were also more likely to be visited by extension agents in both years if they had more and better land, had some irrigated plots, and possessed more assets in livestock.

¹⁹The effect in 2004 is negligible, which is unsurprising given the fall in credit facilities in this period.

²⁰The supplementary figure A.2, illustrates the limited effect of extension visits for a number of villages near

Debre Berhan. In this figure, each circle represents a household, and the size of the circle shows the number of visits by extension agents. A red circle is a household that adopts improved seeds, while a green circle shows a household that did not adopt. As can be seen, there was little adoption despite high numbers of visits.

References

- Bachewe, F., G. Berhane, S. Asrat, G. Getachew, A. Taffesse, and H. J. 2011. "Sources of Inefficiency and Growth in Agricultural Output in Subsistence Agriculture: A Stochastic Frontier Analysis." Working paper, Development Strategy and Governance Division, International Food Policy Research Institute - Ethiopia Strategy Support Program II (ESSP II).
- Bandiera, O., and I. Rasul. 2006. "Social networks and technology adoption in northern Mozambique." *The Economic Journal* 116:869–902.
- Bindlish, V., and R. Evenson. 1997. "The impact of T&V extension in Africa: The experience of Kenya and Burkina Faso." *The World Bank Research Observer*, pp. 183–201.
- Blume, L., and S. Durlauf. 2005. "Identifying social interactions: A review." Working paper.
- Bonger, T., G. Ayele, and T. Kuma. 2004. "Agricultural extension, adoption and diffusion in Ethiopia." Working paper No. 1, Ethiopian Development Research Institute (EDRI) Research Report.
- Bramoullé, Y., H. Djebbari, and B. Fortin. 2009. "Identification of peer effects through social networks." *Journal of econometrics* 150:41–55.
- Conley, T., and C. Udry. 2010. "Learning about a new technology: Pineapple in Ghana." *The American Economic Review* 100:35–69.
- Davis, K. 2008. "Extension in sub-Saharan Africa: Overview and assessment of past and current models and future prospects." *Journal of International Agricultural and Extension Education* 15:15–28.
- Davis, K., B. Swanson, D. Amudavi, D. Mekonnen, A. Flohrs, J. Riese, C. Lamb, and E. Zerfu. 2010. "In-depth assessment of the public agricultural extension system of Ethiopia and recommendations for improvement." Working paper, International Food Policy Research Institute (IFPRI) Discussion Paper 01041.
- Dercon, S., and L. Christiaensen. 2011. "Consumption risk, technology adoption and poverty traps: evidence from Ethiopia." *Journal of Development Economics* 96:159–173.

- Dercon, S., D. Gilligan, J. Hoddinott, and T. Woldehanna. 2009. "The impact of agricultural extension and roads on poverty and consumption growth in fifteen Ethiopian villages." *American Journal of Agricultural Economics* 91:1007–1021.
- Dercon, S., and R. Hill. 2009. "Growth from agriculture in Ethiopia: Identifying key constraints." Working paper, Prepared for a study on Agriculture and Growth in Ethiopia, University of Oxford.
- Doss, C., W. Mwangi, H. Verkuil, and H. De Groote. 2003. *Adoption of maize and wheat technologies in eastern Africa: A synthesis of the findings of 22 case studies*. International Maize and Wheat Improvement Center.
- Duflo, E., M. Kremer, and J. Robinson. 2011. "Nudging Farmers to Use Fertilizer: Theory and Experimental Evidence from Kenya." *American Economic Review* 101:2350–90.
- EEA/EEPRI. 2006. *Evaluation of the Ethiopian agricultural extension with particular emphasis on the Participatory Demonstration and Training Extension System (PADETES)*. Ethiopian Economic Policy Research Institute, Addis Ababa.
- Evans, W., W. Oates, and R. Schwab. 1992. "Measuring peer group effects: A study of teenage behavior." *Journal of Political Economy*, pp. 966–991.
- Evenson, R. 1997. *The economic contributions of agricultural extension to agricultural and rural development*.
- Evenson, R., and D. Gollin. 2003. "Assessing the impact of the Green Revolution, 1960 to 2000." *Science* 300:758–762.
- Evenson, R., and G. Mwabu. 2001. "The effect of agricultural extension on farm yields in Kenya." *African Development Review* 13:1–23.
- Feder, G., R. Just, and D. Zilberman. 1985. "Adoption of agricultural innovations in developing countries: A survey." *Economic development and cultural change* 33:255–298.
- Foster, A., and M. Rosenzweig. 1995. "Learning by doing and learning from others: Human capital and technical change in agriculture." *Journal of political Economy*, pp. 1176–1209.

- Gautam, M., and J. Anderson. 1999. "Reconsidering the evidence on returns to T&V extension in Kenya." Working paper, WPS 2098, World Bank, Washington D.C.
- Getachew, A.A. 2011. "Cereal Productivity in Ethiopia: An Analysis Based on ERHS Data 1999 -2009." Working paper, International Growth Centre, London.
- Gollin, D. 2011. "Crop Production in Ethiopia: Assessing the Evidence for a Green Revolution." Working paper, International Growth Centre, London.
- Kelejian, H., and I. Prucha. 2001. "A generalized moments estimator for the autoregressive parameter in a spatial model." *International economic review* 40:509–533.
- . 1998. "A generalized spatial two-stage least squares procedure for estimating a spatial autoregressive model with autoregressive disturbances." *The Journal of Real Estate Finance and Economics* 17:99–121.
- Manski, C. 1993. "Identification of endogenous social effects: The reflection problem." *The review of economic studies* 60:531–542.
- Moser, C., and C. Barrett. 2006. "The complex dynamics of smallholder technology adoption: The case of SRI in Madagascar." *Agricultural Economics* 35:373–388.
- Spielman, D., D. Kelemwork, and D. Alemu. 2011. "Seed, Fertilizer, and Agricultural Extension in Ethiopia." Working paper, Ethiopia Strategy Support Program II (ESSP II) ESSP II.
- Suri, T. 2011. "Selection and comparative advantage in technology adoption." *Econometrica* 79:159–209.
- Taffesse, A. 2008. "Decomposition of growth in cereal production in Ethiopia." Working paper, Prepared for a study on Agriculture and Growth in Ethiopia.

Figure 1: Index Sequence Plot: Seed Adoption




Figure 2: Index Sequence Plot: Fertiliser Adoption




Figure 3: Hog-Cycle Model Estimates




Figure 4: Probability of Adoption given Neighbours' Adoption




Figure 5: Probability of Adoption given Extension Visits




Table 1: Main sources of information for fertiliser/seed

Source of information	Fertiliser	Seed
extension agents (%)	50	68
friends/neighbours (%)	36	17
observed early adopters (%)	14	15
Average number of people discussed adoption with	4	3

Source: ERHS 1999

Table 2: Rates of adoption: 1999 - 2009

	1999	2004	2009
Adopt new seed %	18	31 ⁿ	23
Use fertiliser %	62	25	64
Neighbours adopting new seed %	17	31 ⁿ	21
Neighbours using fertiliser %	59	26	63
Correlation: own and neighbour seed adoption	0.47	0.37	0.29
Correlation: Own and neighbour fertiliser adoption	0.59	0.55	0.57
Correlation: seed adoption and extension visits	0.29	0.04	0.16
Correlation: fertiliser adoption and extension visits	0.19	0.04	0.11
Number of extension visits in past 5 seasons	0.29	1.06	5.5

ⁿNote: Adoption of new seed is 31% but those using both seed & fertiliser is 9%

Table 3: Differences between adopters and non-adopters of new seed

Years	1999		2004		2009	
	Adopt seed	Not adopt	Adopt seed	Not adopt	Adopt seed	Not adopt
Sample size	151	803	300	654	225	729
Extension visits	0.64	0.05 *	0.28	0.23	0.62	0.47 *
Seed (N'bours)	0.46	0.11 *	0.47	0.24*	0.34	0.17*
Extension (N'bours)	0.37	0.10*	0.81	0.7	0.52	0.47
Land (hectares)	0.80	1.23*	1.72	1.67	1.55	1.48
Irrigated plot	0.19	0.10*	0.24	0.25	0.41	0.35
Share lem land	0.66	0.49*	0.62	0.53	0.69	0.48*
Value of livestock	2090	2284	2697	3121	9621	8961
Oxen	0.96	1.28*	0.92	1	1.2	1.05
Male-headed	0.84	0.77*	0.71	0.68	0.69	0.64
Some schooling	0.36	0.29*	0.42	0.27*	0.60	0.50

* Indicates significant differences across adopters and non-adopters

Table 4: Differences between adopters and non-adopters of fertiliser

	1999		2004		2009	
	Adopt fert	Not adopt	Adopt fert	Not adopt	Adopt fert	Not adopt
Sample size	526	428	240	714	612	342
Extension visits	0.23	0.03 *	0.28	0.23	0.54	0.43 *
Fertiliser (N'bours)	0.73	0.36 *	0.35	0.29	0.78	0.36*
Extension (N'bours)	0.18	0.10	0.78	0.73	0.50	0.45
Land (hectares)	1.31	0.90*	2.34	1.46*	1.73	1.08*
Irrigated plot	0.12	0.10	0.29	0.23	0.45	0.22
Share lem land	0.61	0.38*	0.52	0.57	0.57	0.45*
Livestock value	2763	1394*	4543	2464*	11623	4679*
Oxen	1.39	0.94*	1.1	0.8	1.3	0.69*
Male-headed	0.83	0.70*	0.72	0.68	0.69	0.58*
Some schooling	0.34	0.23*	0.38	0.29	0.58	0.43*

*Indicates significant differences across adopters and non-adopters

Table 5: Neighbours' influence and extension agents in seed adoption

	1999		2004		2009	
	Probit	Probit IV	Probit	Probit IV	Probit	Probit IV
Extension (s.e.)	0.03*** (0.01)	0.03*** (0.01)	-0.00 (0.00)	-0.002 (0.003)	0.003** (0.001)	0.003** (0.001)
Neighbours adopt (s.e.)	-0.15** (0.07)	0.46** (0.20)	-0.17** (0.09)	0.68*** (0.32)	-0.11 (0.07)	0.47** (0.25)
Cragg-Donald F		136.47		89.47		57.94
Sample Size	954					

Notes:

1. Standard errors, robust to heteroscedasticity at the HH level (Huber-White sandwich estimator) in parentheses
2. All specifications control for plot soil type (lemshare, meddashare) plot/land area, amount of irrigated land, household head education, livestock value and total oxen.)
3. * indicates significance at 10%; * at 5%; *** at 1%.

Table 6: Panel estimates, adoption of seed: 1999-2009

	OLS (FD)	IV (FD)	IV (round/FE)	IV (round/FD)
neighbours adopt	0.354** (0.061)	0.930** (0.061)	0.905** (0.092)	0.847** (0.182)
neighbours adopt × Round 6			0.014 (0.101)	0.115 (0.282)
neighbours adopt × Round 7			-0.002 (0.102)	
extension visits	0.003 (0.002)	0.004* (0.002)	0.051** (0.017)	0.004* (0.002)
extension × Round 6			-0.051** (0.017)	-0.001 (0.006)
extension × Round 7			-0.047** (0.017)	
Village F.E.	Yes	No	No	Yes
Observations	1908	1908	2862	1908
Cragg-Donald F		2588.401	545.142	198.167

Notes:

1. Standard errors, robust to heteroscedasticity at the HH level (Huber-White sandwich estimator) in parentheses
3. Columns (1), (2) and (4) report First-Differenced estimates; Columns (3) account for household fixed effects
3. All specifications control for plot soil type (lemshare, meddashare) plot/land area, amount of irrigated land, household head education, livestock value and total oxen.)
4. * indicates significance at 10%; * at 5%; *** at 1%.

Table 7: Neighbours' influence and extension agents in fertiliser adoption

	1999		2004		2009	
	Probit	Probit IV	Probit	Probit IV	Probit	Probit IV
Extension (s.e.)	0.22*** (0.05)	0.22*** (0.048)	0.003 (0.003)	0.003 (0.004)	0.002 (0.002)	0.002 (0.002)
Neighbours adopt (s.e.)	0.18* (0.09)	0.53** (0.21)	0.06 (0.09)	0.13 (0.33)	0.10 (0.11)	0.53 (0.36)
Cragg-Donald F		22.98		79.48		57.94
Sample Size	954					

Notes:

1. Standard errors, robust to heteroscedasticity at the HH level (Huber-White sandwich estimator) in parentheses
2. All specifications control for plot soil type (lemshare, meddashare) plot/land area, amount of irrigated land, household head education, livestock value and total oxen.)
3. * indicates significance at 10%; * at 5%; *** at 1%.

Table 8: Panel estimates, adoption of fertiliser: 1999-2009

	OLS (FD)	IV (FD)	IV (round/FE)	IV (round/FD)
neighbours adopt	0.519** (0.039)	0.972** (0.050)	0.931** (0.057)	0.949** (0.136)
neighbours adopt × Round 6			0.018 (0.057)	
neighbours adopt × Round 7			-0.007 (0.050)	0.058 (0.249)
extension visits	0.002 (0.002)	0.002 (0.002)	0.028** (0.011)	0.001 (0.002)
extension × Round 6			-0.027** (0.012)	0.004 (0.004)
extension × Round 7			-0.028** (0.011)	
Village F.E.	Yes	No	No	Yes
Observations	1908	1908	2862	2862
Cragg-Donald F		2556.335	672.345	207.146

Notes:

1. Standard errors, robust to heteroscedasticity at the HH level (Huber-White sandwich estimator) in parentheses
2. Columns (1), (2) and (4) report First-Differenced estimates; Columns (3) account for household fixed effects
3. All specifications control for plot soil type (lemshare, meddashare) plot/land area, amount of irrigated land, household head education, livestock value and total oxen.)
4. * indicates significance at 10%; * at 5%; *** at 1%.

SUPPLEMENTARY APPENDIX

Title: “AJAE appendix for Neighbours and Extension Agents in Ethiopia: Who Matters More for Technology Adoption”

Authors: Pramila Krishnan, Manasa Patnam

Date: February 2013

Note: The material contained herein is supplementary to the article named in the title and published in the American Journal of Agricultural Economics (AJAE).

Figure A.1: Location of the Villages, Adoption Rates and Extension Visits 2009




Figure A.2: Adoption and Extension visits in Villages near Debre Berhan



Table A.1: Panel estimates, adoption of seed (alternative networks): 1999-2009

	OLS (FD)	IV (FD)	IV (round/FE)	IV (round/FD)
neighbours adopt	0.531** (0.069)	0.936** (0.059)	0.899** (0.088)	1.005** (0.144)
neighbours adopt × Round 6			0.033 (0.094)	
neighbours adopt × Round 7			0.049 (0.101)	-0.205 (0.316)
extension visits	0.004 (0.002)	0.004** (0.002)	0.055** (0.016)	0.004* (0.002)
extension × Round 6			-0.054** (0.016)	0.000 (0.007)
extension × Round 7			-0.052** (0.016)	
Village F.E.	Yes	No	No	Yes
Observations	1908	1908	2862	2862
Cragg-Donald F		4798.867	1168.641	206.855

Notes:

1. Standard errors, robust to heteroscedasticity at the HH level (Huber-White sandwich estimator) in parentheses
2. Columns (1), (2) and (4) report First-Differenced estimates; Columns (3) account for household fixed effects
3. All specifications control for plot soil type (lemshare, meddashare) plot/land area, amount of irrigated land, household head education, livestock value and total oxen.)
4. * indicates significance at 10%; * at 5%; *** at 1%.

Table A.2: Panel estimates, adoption of fertiliser (alternative networks): 1999-2009

	OLS (FD)	IV (FD)	IV (round/FE)	IV (round/FD)
neighbours adopt	0.717** (0.045)	0.998** (0.049)	0.975** (0.057)	1.052** (0.141)
neighbours adopt × Round 6			0.027 (0.056)	-0.103 (0.265)
neighbours adopt × Round 7			-0.005 (0.049)	
extension visits	0.002 (0.001)	0.002 (0.001)	0.040** (0.010)	0.001 (0.002)
extension × Round 6			-0.039** (0.011)	0.005 (0.004)
extension × Round 7			-0.040** (0.010)	
Village F.E.	Yes	No	No	Yes
Observations	1908	1908	2862	2862
Cragg-Donald F		6575.838	1561.941	310.987

Notes:

1. Standard errors, robust to heteroscedasticity at the HH level (Huber-White sandwich estimator) in parentheses
2. Columns (1), (2) and (4) report First-Differenced estimates; Columns (3) account for household fixed effects
3. All specifications control for plot soil type (lemshare, meddshare) plot/land area, amount of irrigated land, household head education, livestock value and total oxen.)
4. * indicates significance at 10%; * at 5%; *** at 1%.

Table A.3: Dynamic adoption (dep. variable: adoption in mehr)

	1999		2004		2004 (dynamic)	
	Seeds	Fertiliser	Seeds	Fertiliser	Seeds	Fertiliser
Extension (s.e.)	0.018*** (0.006)	0.104*** (0.026)	-0.000 (0.002)	0.004 (0.003)	0.000 (0.003)	0.003 (0.005)
Neighbours adopt in Belg (s.e.)	0.921 (0.698)	0.606* (0.381)	0.501* (0.313)	-0.685 (0.553)	0.672** (0.313)	-0.695 (0.477)
Own Lagged Adoption (1999) (s.e.)					-0.151 (0.207)	0.330*** (0.121)
Cragg-Donald F	110.200	347.756	146.54	208.300	9.029	7.339
Sample Size	805	891	954	789	954	954

Notes:

1. Standard errors, robust to heteroscedasticity at the HH level (Huber-White sandwich estimator) in parentheses
2. All specifications control for plot soil type (lemshare, meddshare) plot/land area, amount of irrigated land, household head education, livestock value and total oxen.)
3. * indicates significance at 10%; * at 5%; *** at 1%.