Social Learning in Health Care: Evidence from Tanzania*

Kenneth L. Leonard†

February 9, 2005

Abstract

Learning is an important force for progress in developing countries and may represent a significant underutilized resource in health care. Using two separate data sets from Tanzania, we show that households know the outcomes of large numbers of health episodes, that they value quality at health facilities, that they demonstrate increased willingness to pay for quality as the tenure of clinicians increases and that household behavior is consistent with local social networks of learning. Households gather information, change their opinions about doctors based on this information and improve their health by choosing the best doctors.

JEL Classification: I1, O1

Keywords: Social Networks, Social Learning, Health Care

*This work was funded by NSF grant number 00-95235 and an ISERP (Columbia University) seed grant. I am grateful to the careful and extensive input of David Bedoll as a research assistant, and to Tim Conley and Yaw Nyarko for their advice. In addition, this paper has benefited from the comments of Doug Golin, Benoit Dostie and participants at the development mini-conference at CIRPEE.

†2200 Symons Hall, University of Maryland College Park, MD 20742 `kleonard@arec.umd.edu`

1 Introduction

Learning is an important source of economic growth. It allows actors to improve their efficiency for existing activities and to adopt new technologies. In developing countries, social learning is seen as a particularly important source of progress. Actors observe the activities and/or outcomes of other actors and use this information to augment the information that they are able to gather on their own. Social learning may be a potent force in the health care systems of developing countries, as households are forced to make almost daily decisions about medical care that affect their health and yet have access to few forms of formal information about medical practice. Although most health systems and technologies are relatively stable, there are occasional changes in the local clinicians who practice medicine. Quality varies significantly from clinician to clinician (even within organizations and facilities) and therefore every time a new clinician arrives, patients are presented with an opportunity to learn about quality.

Using two separate data sets from Tanzania, we show that households know the outcomes of large numbers of health episodes, that they value quality at health facilities, that they demonstrate increased willingness to pay for quality as the tenure of clinicians increases, and that household behavior is consistent with local networks of learning. Households gather information, change their opinions about doctors based on this information and improve their health by choosing the best doctors. Households gather data and appear to invest more in gathering data when it is more likely to contain useful information. This data implies that, by communicating with others, the average household knows about the outcomes of at least 10 times as many health episodes as they are likely to experience themselves. Households are attracted to facilities that practice better medicine and they are more attracted to good facilities when they have had time to learn about quality. In addition, we show that households that are close to each other and are more likely to communicate exhibit strong similarities in behavior when clinicians are new. We estimate that patient behavior is largely stable four to five years after a clinician has arrived, suggesting that learning takes up to five years.

Rural areas all over the world are famous for their abilities to spread information about others. Health care in Africa is no exception. People know who among their neighbors was sick, what they were sick with, where they chose to go, and what happened in the end. We examine data in which people are asked about the health histories of randomly selected neighbors at varying distances and find that people correctly identify about 2% of the illnesses suffered by any of 600 households within their village over the course of the last year. When they correctly identify the illness, they know the medical location visited and the outcome. People know more about neighbors who are closer to them and households are more likely to know about an illness if the other household visited a newer provider. The fact that people know more when other households visit providers who are new suggests that the information shared among households may be of some use; it is not idle gossip.

In a separate data set from Tanzania, we observe the behavior of patients choosing between a set of possible providers spread over a large geographical area. When someone is sick she may choose to go the nearest provider (who may or may not be in her village), or to any of the other providers in the area. Travel is expensive relative to other costs and incomes, and the amount patients spend traveling beyond the nearest health facility for the average condition amounts to about 4% of annual per capita public expenditure on health care. We measure quality at the facilities between which patients can choose and show that patients are willing to travel further and pay greater fees for facilities with higher objectively measured quality. Demand for objectively measured quality implies knowledge of quality: learning. The willingness to pay for quality increases as the tenure of the clinician increases; patients are learning about quality and are willing to pay more when they are more confident in their assessment of quality. When a new clinician arrives at a health post, he is treated differently. As time passes, patients change their pattern of use even when the quality of that clinician has not changed.

In addition, we show that the unexplained portion of behavior is highly correlated for people who live near to each other at precisely the time when we expect people to be learning from each other. People who should learn from each other exhibit a similarity in behavior that does not exist for people who live far apart, and this tapers away as the tenure of the clinician increases. These patterns are consistent with households who learn from the experience of other households in their social network until they reach a point where quality is essentially known.

Learning and technology adoption have been central issues in development economics for many decades (see Feder, Just and Zilberman, 1985, for a review) and the role of social learning in promoting growth and technology diffusion has been featured in the endogenous growth literature (Aghion and Howitt, 1998; Lucas, 1988; Romer, 1986). Only recently, however, have economists made efforts to measure the quantitative importance of learning from others. Measuring the extent of social learning is difficult for two major reasons. First, the set of neighbors from whom an individual can learn is difficult to define. Second, even with a proper definition of this set, distinguishing learning from other phenomena that may give rise to similar observed outcomes is problematic. In the absence of learning, individuals may still act like their neighbors as a result of interdependent preferences, technologies, or because they are subject to related unobservable shocks.

This paper shows that households posses the information from which they can learn though we cannot show that they deliberately collect this information. There is some weak evidence that they choose to know more about health episodes that may be more useful in generating information. The behavior exhibited by patients shows strong spatial correlation in behavior. These patterns can be caused by many factors that are unrelated to learning. We identify the portion of correlation that is due to learning by comparing spatial correlation in behavior when a clinician is new to spatial correlation when a clinician has been present for a long time. Other contributions to spatial correlation are unlikely to systematically vary with the tenure of clinicians. Also important to identifying our story of learning is the fact that household behavior is rational (exhibits preferences for objectively measured positive features of medicine) and increasingly rational as tenure increases.

In the following section, we develop a simple model of learning that we suggest is appropriate to the health care context. That section includes the analysis of the data on knowledge about health outcomes of other households. Section 3 briefly introduces the data of provider choice (that data is discussed at greater length in Leonard, Mliga and Haile Mariam (2002)) and analyzes both the determinants of choice and the patterns of spatial correlation. Section 4 concludes.

2 A model of learning in health care

Households in rural Africa face a choice between practitioners of varying quality. A woman suffering from abdominal pain can choose to visit a provider who will give her medicine to alleviate the pain, or a provider who will give her a careful examination to see if she is suffering from a serious infection (such as Pelvic Inflammatory Disease, PID). The woman does not know if she is suffering from this condition, nor does she know that she could be suffering from this condition. When she leaves a consultation, she does not even know whether she visited the high quality provider or the low quality provider. She might be cured by the low quality provider (if it was not PID) or she might even fail to be cured by the high quality provider (if it was a serious infection, or resistant to the drugs administered). Yet if she chooses the low quality provider she might face infertility and a greatly increased chance of becoming infected with HIV/AIDS, whereas with the high quality provider she would most likely be cured and reduce her chance of getting HIV/AIDS. The choice is important, yet she possesses no information with which to guide her choice. In this case, she will turn to the advice of her friends and neighbors. Quality health care does not make bad outcomes impossible, nor does low quality health care eliminate the possibility of good outcomes. However, the probability of being cured is much higher with high quality health care. When people learn from the experience of others, they increase the number of observations from which to draw inferences. Social learning is not only important; it may be the only source of information about quality. Personal experience yields too few observations and there is no formal method of evaluating quality (such as reading newspaper articles or learning from extension officers). At any given moment in time, households may have strong opinions about the qualities of the providers from which they can choose, but when a new doctor arrives, households can choose to learn about the quality of care they might get if they visit that practitioner.

Recent work on learning in developing countries (Foster and Rosenzweig (1995) and Conley and Udry (2003), for example) is based on the target-input model represented in Jovanovic and Nyarko (1996). Learning in our context is better represented as a choice between technologies (such as Besley and Case, 1994). The literature on learning with multi-armed bandits is a more appropriate starting place. In the multi-armed bandit model (see Banks and Sundaram, 1994; Brezzi and Lai, 2002; Gittins and Jones, 1974, for example) a player chooses between two options that offer uncertain payoff. The player engages in experimentation, trying to learn about the true payoffs associated with each choice. Some set of outcomes will lead the player to choose one option and to cease experimentation. If a new option becomes available, the player must measure the expected payoff of the unknown before deciding to experiment. For the purposes of this paper, a more appropriate model is the multi-agent multi-armed bandit in which players can learn from their own experimentation as well as that of other players. Aoyagi (1998) examines learning when the actions but not payoffs of the other player are observable and Bolton and Harris (1999) examine the case in which the payoffs of others are observable.

For a variety of reasons, a simplified version of this model of behavior can represent learning about quality in a health care context. Patients learn from their own experience as well as the experience of others. Learning from others has a cost, but there is no attempt by others to hide information or deceive; there is no advantage gained when you retain information. Furthermore, as we shall show, patients continue to visit providers even when they have learned that quality is low and patients will visit new providers even if their prior expectation of quality is low. Thus, although there may be an experimentation motive, it is not the only means by which patients learn about quality. Therefore, we model learning as a myopic process in which patients will only visit a provider if the expected value of visiting that provider is higher than at any other provider; patients do not invest in information through their choice of provider.

Learning takes place when a household is trying to learn the type (quality) of a clinician. We assume there are only two types of clinicians, good ($\phi^\star$) and bad ($\phi^\emptyset$): $\phi \in \{\phi^\star, \phi^\emptyset\}$. Households have a prior belief over clinician type $\tilde{r}_t$, which varies with the number of illnesses or outcomes observed, t. The prior can be expressed as a log likelihood ratio (LLR):

$$\lambda = \log\left(\frac{\Pr(\phi^\star)}{\Pr(\phi^\emptyset)}\right) = \log\left(\frac{\tilde{r}_t}{1-\tilde{r}_t}\right) \tag{1}$$

where $\Pr\phi^\star$ is the patients belief of the probability that the clinician is good. When a household observes the outcome, h_t for a visit t they gain information about the type of clinician. The LLR ratio is updated according to Bayes rule:

$$\lambda_{t+1} = \lambda_t + \log\left(\frac{\Pr(h_t|\phi^\star)}{\Pr(h_t|\phi^\emptyset)}\right) \tag{2}$$

In this notation, the LLR is updated by an increment that is independent of the previous LLR. The outcome is binary: the patient is cured or not cured, represented as $h \in \{\bar{h}, \underline{h}\}$. If the clinician is good, the probability of a good outcome is $\rho^\star$, if the clinician is bad, the probability of a good outcome is $\rho^\emptyset$.

The good type is defined by the fact that $\rho^\star > \rho^\emptyset$, but it is not necessary that $\rho^\star > \frac{1}{2}$ nor that $\rho^\emptyset < \frac{1}{2}$. For some illnesses, even high quality clinicians may have a low cure rate and for other illnesses even low quality clinicians may have a high cure rate, but the cure rate for high quality clinicians is always higher than the cure rate for low quality clinicians.

The updating rule becomes:

$$\lambda_{t+1} = \lambda_t + \left[h_t = \bar{h}\right] \log\left(\frac{\rho^\star}{\rho^\emptyset}\right) + \left[h_t = \underline{h}\right] \log\left(\frac{1-\rho^\star}{1-\rho^\emptyset}\right) \tag{3}$$

When the true value is $\phi^\star$, the expected value of the change in the LLR can be shown to be positive.

$$E\left(\lambda_{t+1} - \lambda_t | \phi = \phi^\star\right) = \rho^\star \log\left(\frac{\rho^\star}{\rho^\emptyset}\right) + (1-\rho^\star) \log\left(\frac{1-\rho^\star}{1-\rho^\emptyset}\right) > 0 \tag{4}$$

We can express the prior as $\tilde{r}_t = \frac{e^{\lambda_t}}{1-e^{\lambda_t}}$ and since the expected value of λ_t is increasing in t (when the clinician is good), $\tilde{r}_t$ approaches 1 asymptotically: with enough observations the belief approaches the true value.

In health care, households are likely to make many observations of outcomes at a particular provider, but these observations may represent multiple illnesses. $\rho^\star$ and $\rho^\emptyset$ will not be constant across illnesses. However, so long at the type is invariant to illnesses (doctors are good for all illnesses, or bad for all illnesses) and patients know the values of $\rho^\star$ and $\rho^\emptyset$ for each illness, observation of sufficient outcomes will lead the prior to approach the true value asymptotically.[1]

When a household learns only from its own experience and is myopic when choosing a practitioner, the flow of illnesses from which information may be gathered is exogenous to the learning process. However, if a household can learn from the experience of other households and there is some small cost to gathering this information, the flow of information may be endogenous: households may choose to concentrate on information that is more useful to them. We can express $\rho^\emptyset$ as a function of $\rho^\star$ and an arbitrary k, such that $k = \frac{\rho^\star}{\rho^\emptyset}$. k is a measure of the value of quality. When k is close to 1, quality is less valuable,

[1] Das (2001) examines the case in which patients do not always know the true values of $\rho^\star$ and $\rho^\emptyset$ and show that, without knowledge of these probabilities, learning does not always lead in the right direction. Ultimately this is an empirical question and our data suggests that patients in this context do know $\rho^\star$ and $\rho^\emptyset$.

when k is large, quality is more valuable. The expected value of the updating increment $(E(\lambda_{t+1} - \lambda_t | \phi = \phi^\star))$ can be shown to be increasing in both $\rho^\star$ and k. Thus, if the cost of gathering information is constant, households are more likely to gather information when $\rho^\star$ and k are large.

2.1 Choosing a practitioner

The choice to visit a provider follows a different process than the value of the information generated from a visit to a provider. The value of visiting a provider is a function of both the probability of a cure and the value of a cure. Following Leonard and Graff Zivin (2003), the expected value of health care is a function of $\rho^\star$, $\rho^\emptyset$ and $\tilde{r}_t$ (the belief of quality):

$$EU = \tilde{r}_t \underbrace{\left(\rho^\star \bar{U} + (1-\rho^\star)\underline{U}\right)}_{\text{EU under } \phi = \phi^\star} + (1-\tilde{r}_t) \underbrace{\left(\rho^\emptyset \bar{U} + (1-\rho^\emptyset)\underline{U}\right)}_{\text{EU under } \phi = \phi^\emptyset} \tag{5}$$

$$\text{where} \qquad \bar{U} = U[\bar{h}, (I(\bar{h}) - C)] \qquad \text{and} \qquad \underline{U} = U[\underline{h}, (I(\underline{h}) - C)]$$

$\bar{U}$ is the utility if the patient is cured and $\underline{U}$ is the utility if the patient is not cured. The utility of health care is a function of the outcome $(h \in \{\bar{h}, \underline{h}\})$, the income potential at the outcome $(I \in \{I(\bar{h}), I(\underline{h})\})$ and any cash costs associated with the visit (C). We assume a separable utility function $(U(H) = V[H, I(H)] - C)$, and compare the expected utility of visiting provider i and j.

$$\Delta EU_{(i,j)t}(\tilde{r}_{it}, \tilde{r}_{jt}) = (\tilde{r}_{it} - \tilde{r}_{jt}) \cdot (\rho^\star - \rho^\emptyset) \cdot \left(\bar{V} - \underline{V}\right) - (C_i - C_j) \tag{6}$$

Note that this derivation of expected utility does not include the value of information gained from a visit.

Although the relative belief of provider type ($\tilde{r}_{it}$- $\tilde{r}_{jt}$) is an important determinant of the choice to visit a provider, it is not the only determinant. When $\tilde{r}_{it} - \tilde{r}_{jt} < 0$ (the patient believes provider i is less likely to be good than provider j) patients will still visit provider

i for some conditions, even without any experimentation motive. In the extreme case that $\tilde{r}_{it} = 0$ and $\tilde{r}_{jt} = 1$, the patient may still visit provider i given some values of $\rho^\star - \rho^\emptyset$ and the net cash costs (fees paid to the provider and travel costs.) It is common to find patients who visit a local provider (travel costs are zero) even though they know that this provider is of very low quality. They do so when the difference in the probability of cure between high and low type providers ($\rho^\star - \rho^\emptyset$) is very low.

Representing $\rho^\star - \rho^\emptyset$ as $\rho^\star(1 - 1/k)$ it should be clear that, when $\tilde{r}_{it} - \tilde{r}_{jt} < 0$, the probability of visiting provider i is decreasing in $\rho^\star$ and k. If the uninformed prior for a new clinician ($\tilde{r}_{it}$) is lower than the more informed prior of another clinician ($\tilde{r}_{jt}$), patients are less likely to visit the new clinician when they suffer exactly the types of conditions that yield the most useful information about quality. In the absence of deliberate experimentation, learning is likely, therefore, to be slow initially, but if a clinician is good, learning may accelerate with myopic behavior.

2.2 Evidence of Learning in Social Networks

Data on the degree to which households know about other households' health episodes was collected by the author in Arusha region of Tanzania in 2002. The health histories of households over the previous year were compared to what neighbors at varying distances said they knew about these households. We refer to the household discussing its knowledge of other households as the respondent household and the household about which questions were being asked as the reference. We asked respondents to list every health episode they recalled from the reference household and then we compared their answers to those given by the reference household. Starting with an illness declared by the reference household, we have two possible respondent household responses: either the respondent household knew about the illness (and the outcome) or they did not. If they did not know any members of the reference household this was coded as not knowing about any of the episodes.

We use political boundaries to proxy for social distance. Households are members of cells (15 to 20 households) and then subvillages (approximately 10 cells) and then villages (usually 4 sub-villages). Respondent households correctly identified 8.8% of illnesses when the reference household was within the cell. At the sub-village level (outside the cell) they identified 1.8% of illnesses and at the village level (outside the subvillage) they identified 1.3% of illnesses. Assuming that the expected number of illness episodes for each reference household is equal to the expected number of illness episodes for the respondent household, each household learns from over 10 times as many episodes as it can expect to experience itself, despite what may seem a low percentage of knowledge.

Table 1 shows the results of logit regressions of the dependent variable (knowledge of an illness episode versus no knowledge of the episode) regressed on a series of independent variables including characteristics of the patient and the illness condition. We present information on the illness condition in two different manners. In regressions I, III and V we use the symptoms and self-reported (by the reference household) severity of the illness. In regressions II and III, we use characteristics of the illness as judged by qualified medical personnel who reviewed all available information on each illness episode.[2] We use the medically evaluated probability of being cured by a good doctor ($\rho^\star$) as well as the ratio of the probability of being cured by a good doctor and the probability of being cured with no care ($k = \frac{\rho^\star}{\rho^\emptyset}$).

Households know more about other households that are within the same cell. They do know about households outside of the cell but they are much less likely to know about illnesses from a randomly selected household outside of their cell. The characteristics of the illness, and in particular the patient–declared severity, are important predictors of whether or not a respondent household will recall the details of the reference household's illness. This is not surprising and is not necessarily an indicator that learning about illnesses and

[2] 16 doctors were given files that contained the age and gender of the patient, all reported symptoms, the self-declared severity, the ability of patients to perform certain routine daily activities before and during the illness (activities of daily living: ADL) the length of time the patient had suffered and the number of days in which the patient was bed-ridden. Each case was reviewed by four doctors. We normalized the scores for each doctor and then generated an average score for each illness based on the four independent scores.

Table 1: Logit regression of illness episode being mentioned by respondent household

Regression	I		II		III		IV		V	
Illness Characteristics										
Symptoms	included				included				included	
Pat. Declared Severity*										
"very severe"	**0.956**	(0.410)			**0.744**	(0.389)			0.164	(0.808)
"severe"	**0.541**	(0.241)			0.456	(0.281)			0.113	(0.401)
"average"	0.356	(0.284)			0.314	(0.293)			-0.214	(0.454)
Med. Declared Severity										
$k = \rho^\star / \rho^\emptyset$			**0.178**	(0.076)	0.059	(0.103)				
$\rho^\star$			-0.038	(0.063)	-0.041	(0.069)				
Patient Characteristics†	included		included		included		included		included	
Match Characteristics‡										
same cell	**1.960**	(0.217)	**1.933**	(0.213)	**1.930**	(0.215)	**2.153**	(0.284)	**2.197**	(0.283)
same subvillage	0.273	(0.348)	0.235	(0.352)	0.228	(0.360)	0.586	(0.519)	0.580	(0.491)
Provider choice Chars.										
log of tenure (years)							**-0.101**	(0.049)	**-0.118**	(0.047)
constant	**-4.998**	(0.492)	**-4.244**	(0.369)	**-4.914**	(0.536)	**-4.526**	(0.321)	**-4.872**	(0.993)
Number of obs		9022		8765		8765		3653		3603
Pseudo R^2		0.13		0.10		0.13		0.09		0.15
Log pseudo-likelihood		1397		1405		1366		539		502

Logit regression: dependent variable is 1 if the respondent household correctly identified the reference household's illness and 0 otherwise. Standard errors reported in parentheses are robust to heteroscedasticity and allows for clustering of errors within each respondent cell/ reference cell match. Bold font indicates significance at the 90% level.

* Omitted category of severity is "it was not severe at all"

†Patient characteristics include gender and 5 age categories.

‡Omitted category is "same village, different subvillage"

Regressions I, II and III are for all possible matches for all illnesses. Regressions IV and V are for all possible matches for illnesses that led to a visit to a health center or clinic. Tenure is not known for traditional healers, or for individual clinicians at hospitals.

outcomes is part of a deliberate process of learning. Severe illnesses generate more interest. If we include only the medically determined characteristics of the illness ($\rho^\star$ and k), we find that households are more likely to know about illnesses when the value of quality (k) is larger. However, unlike our prediction from the Bayesian model they are not more likely to know about illnesses when the probability of a cure is higher ($\rho^\star$). When we include the symptoms and patient declared severity we find that neither k nor $\rho^\star$ remain significant.

Regressions I through III combine knowledge about illnesses that result in visits to any type of provider. Knowledge is more likely to be useful to respondent households if patients in reference households are visiting new providers. Regressions IV and V restrict the set of potential matches to those illnesses that resulted in a visit to facilities that have at most two providers (health centers and clinics.) For these facilities, we know the tenure of all clinicians and can use the log of clinician tenure (or log of average tenure for two clinicians) as an independent variable. Unconditionally (regression IV) and controlling for illness characteristics (regression V) we find that respondent households are more likely to know about an illness when patients visit newer providers. Without a formal model of the first stage (the decision to visit a new provider), we cannot draw strong conclusions from these last two regressions. However, there is some evidence that households know more about illnesses that result in visits to new providers.

These regressions suggest that patients have access to information that is useful in determining the quality of care at a provider. They know more about illnesses that would transmit more information and they know when people visit providers they should want to learn about. We can neither conclude nor exclude that patients are following a deliberate information gathering process.

2.3 Linking learning and behavior

Households have access to information about other households that is useful in learning about clinicians. In addition, households learn more from other households that are closer.

We are interested in tracking the impact of learning on the choices of households faced with uncertainty about the quality of the practitioners they visit. A simple simulation illustrates the impact of learning on behavior. In this simulation there are four village and we trace the impact of mutual learning on the behavior of two villages (A and B). These could be households, cells or villages, but within the unit there is complete information. All four villages observe four separate series of independent identically distributed outcomes from the same health care provider and each village observes one-third of the outcomes from another village. In the first scenario, there is no link between A and B. A learns from C and B learns from D. In the second scenario, there is a link between A and B. The amount of learning and the number of outcomes observed is identical in each scenario, but only in the second scenario are A and B learning from each other. This corresponds to a real world situation in which the first pair A and B are far apart and the second pair are close together.

Figure 1 illustrates the connection between learning and spatial correlation behavior. For the sake of this example, there is no choice between clinicians; in every period each village observes one outcome and uses this information to update its belief of the quality of the clinician (who in this case is of high quality).

The top third of the graph represents the stochastic nature of learning. Beliefs are converging toward the true quality, but since there are both good and bad outcomes, the convergence is not smooth. The middle axes show that the difference in A and B's priors is zero at the beginning and at the end of the process, but that the random nature of outcomes causes their priors to diverge during the learning process. With mutual learning, the divergence in the prior is less than without mutual learning. If A and B could see all of each other's outcomes there would be no difference in their beliefs. The bottom axis shows the correlation in the deviation of beliefs from the true value. This graph represents the average of 100 samples drawn from the same process as that represented in the top two axes. It shows that the correlation in deviations from the true value start out large both with and without learning, but that with learning the correlation in the deviation of A and B's beliefs

Figure 1: Bayesian learning within two villages with and without mutual learning

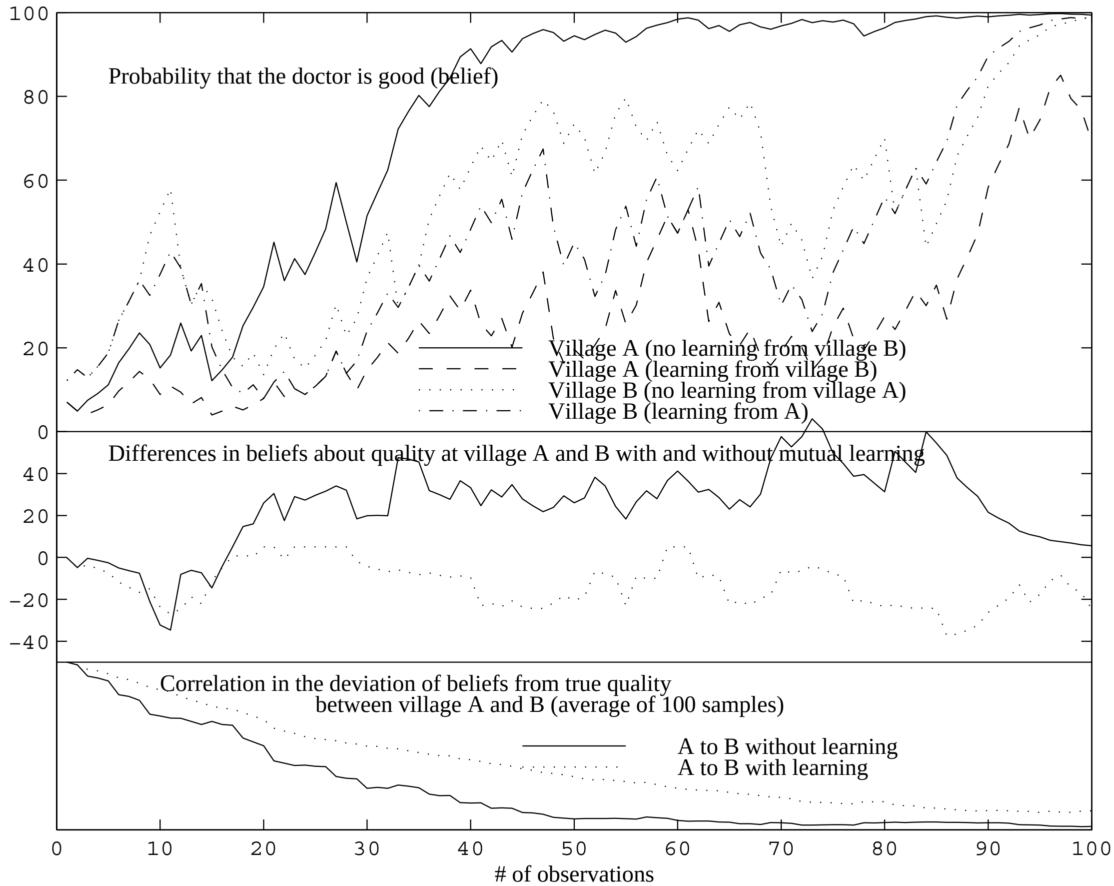

Village A and B observe 100 outcomes of visits to the same (good) doctor. In addition, each village observes one-third of the outcomes of visits to the same doctor from another village. Using these observations, each village learns the quality of the doctor, beginning from a prior ($\tilde{r}_0$) of 10% and approaching the true value ($\tilde{r}_{100} \approx 1$). Under one scenario A and B learn from each other and under the other they learn from other independent villages. In both cases they are exposed to the same number of observations. The development of the priors is plotted on three vertical axes (from top to bottom): probability that doctor is good (belief), difference between beliefs of the two villages with and without mutual learning, and the correlation in the deviation of beliefs from true quality.
The upper graph shows that village A and B approach the true value. For the random distribution graphed, village A approaches the true value faster when it is not learning, but this is only a function of this random draw.
The second graph shows that although A and B start at the same prior and arrive at a point close to the true value, their priors diverge in between these two points. Importantly, the difference in priors is larger when A and B are not learning from each other.
The bottom graph is the average of 100 repetitions of the stochastic process underlying the first two graphs. It shows that the correlation in deviation from true value falls over the number of observations both with and without learning. Importantly, the correlation in deviation is higher when there is mutual learning.
$\phi = \phi^{\star}$, $\rho^{\star} = 70\%$, $\rho^{\emptyset} = 50\%$ and $\tilde{r}_0 = 10\%$.

is higher than without learning.

When village know nothing but have the same priors, they behave in very similar ways that are not well predicted by the true value. When villages know everything with certainty they behave in similar ways that are well predicted by the true value. Neither of these facts indicates any communication between villages; it only shows that they learn from the same process. However, if villages are learning from each other, the evolution of their beliefs from uniformed to informed will follow similar paths and their behavior will therefore be more highly correlated than if they are not learning from each other.

The correlation in deviation from the true value is one of the statistics that we use to indicate the presence of learning in this paper. Correlation in behavior that is caused by spatially distributed factors (people who face similar circumstances behave in a similar manner) will not vary with experience at a particular clinician. On the other hand, correlation in behavior that is caused by peer pressure (people imitate others without improving outcomes) should not show any convergence to the true value of quality. Thus, correlation in behavior (after controlling for observable characteristics) that increases with the probability of communication and decreases over time is an indicator of social learning.

We know apply this insight to a separately generated data set.

3 Choosing Practitioners

Using data from Iringa rural district in Tanzania, we model the choice between a series of possible providers as a function of distance, expected costs and objectively evaluated measures of quality. The objective measures of quality were obtained using a medical team that evaluated facilities. These measures of quality are not the same as patient measures of quality, but if there is learning, medical– and patient–evaluated quality should be correlated.

Using this same data, Leonard et al. (2002) show that patient choice of provider is a function of medically evaluated quality, suggesting that patients have access to information that is at least correlated with medical quality. In this paper, we take advantage of information on the tenure of clinicians to show that medically evaluated quality is more important as tenure increases. In addition, the residuals of this model show strong patterns of spatial correlation that is highest when clinicians are new and tapers off as tenure increases. Both findings are evidence of learning.

3.1 Choice of provider and practitioner quality

Patients in Iringa can choose between government and nongovernmental (NGO) health care providers at clinics. The data include 4,644 patients in 90 villages choosing between 46 facilities. All of the facilities are clinics and are similar in terms of capacity but differ markedly on other measures of quality. Quality scores were gathered by other doctors and nurses on the research team and include the following:

CONSULT: A clinician on our research team observed clinicians consulting patients. For each condition presented, there is a list of expected history taking and physical examination questions. Clinicians who provide more of these expected inputs score higher. We determine a score for each clinician using the clincian level fixed effect after controlling for the order of observations.

PRESCRIP: Clinicians read randomly selected prescriptions from the files at each facility visited. This score reflects the appropriateness of the prescriptions given to patients. The 'perfect' prescription involves *only* the necessary drugs and *all* the necessary drugs.

INJECT: By examining randomly selected records of treatments given to patients we recorded the percentage of non-infant prescriptions for malaria that were given by injection. This is a per-facility average intended to reflect a proclivity to injections in general. We reverse this variable so that larger values represent facilities that give fewer injections than average, a good quality.

N DRUGS: From the same records we determine the number of drugs prescribed for the average patient. Unlike PRESCRIP, this is not necessarily in conflict with the health of patients, but inflicts unnecessary costs. This is a per-facility average. We reverse this variable so that larger values represent facilities that prescribe fewer drugs than average, a good quality.

INFRASTR: A composite score based on construction and cleanliness of the grounds, availability of important medicines and the presence of important tools such as microscopes, etc

All scores are normalized so that they have a mean of 0 and a standard deviation of 1. Thus, a positive value for a score represents above average quality. Each facility in the sample has a unique mixture of these scores and we can observe patients choosing between facilities.

We use a random utility model based in part on the model introduced in the previous section, but with a more flexible view of quality. Taking Equation 6 we express the net utility for individual n suffering from illness k of visiting provider j compared to not visiting any provider ($\tilde{r}_{it} = 0, C_i = 0$).

$$\begin{aligned} \Delta EU_{njk} &= \tilde{r}_{jt} \cdot (\rho^{\star} - \rho^{\emptyset}) \cdot \left(\bar{V} - \underline{V}\right) - C_{jk} + \varepsilon_{ijk} \\ &= \tilde{r}_{jt} G_{nk} - C_{jk} + \varepsilon_{ijk} \end{aligned}$$

The net utility of visiting provider j with illness k is a function of the household's belief of the quality of provider j, $\tilde{r}_{jt}$, the individual/illness specific value of a cure G_{nk} and the costs of care. We do not have a single binary measure of quality. Instead, we have 4 continuous measures of quality. We represent these in the following manner:

$$\Delta EU_{jk} = \sum_{l} (r_{jl} \cdot G_{kl}) - C_{jk} + \varepsilon_{ijk} \tag{7}$$

where l is an index of inputs. We estimate C_{jk} by exploiting the fact that fees differ only by the owner of the institution and although there are many facilities, there are only three owners (for details see Leonard et al., 2002). We have objective estimates of r_{jl} and we estimate G_{kl}.

The model is estimated using multinomial probit estimation allowing for correlation in the errors according to provider chosen. Each household is seen as facing a choice between the two closest government providers and the three closest NGO providers.[3] The actual *order* of distance to a facility (closest, next closest, etc.) is not part of the estimation procedure, but only *absolute* distance to each facility. Patients are choosing between 5 providers and each provider has a different set of characteristics that can impact the choice of provider. Since there are 46 total facilities not all patients face the same actual choice set. We order the choices as $j \in \{G_1, G_2, NGO_1, NGO_2, NGO_3\}$. We allow for correlation in errors according to whether a facility is government or NGO. The vector of errors $\varepsilon'_n = \langle \varepsilon_{i1}, \varepsilon_{i2}, \varepsilon_{i3}, \varepsilon_{i4}, \varepsilon_{i5} \rangle$ is normally distributed with mean of zero and covariance matrix Ω as follows:

$$\Omega = \begin{bmatrix} \alpha_1 & \alpha_3 & \alpha_4 & \alpha_4 & \alpha_4 \\ & \alpha_1 & \alpha_4 & \alpha_4 & \alpha_4 \\ & & \alpha_2 & \alpha_5 & \alpha_5 \\ & & & \alpha_2 & \alpha_5 \\ & & & & \alpha_2 \end{bmatrix} \tag{8}$$

We reduce the problem to a 4 dimensional choice by subtracting the utility of the first choice (G_1) from all other utilities, and set the scale of utilities. The covariance matrix is reduced to:

$$\Omega^* = \begin{bmatrix} 1 & \alpha_2^* & \alpha_2^* & \alpha_2^* \\ & \alpha_1^* & \alpha_3^* & \alpha_3^* \\ & & \alpha_1^* & \alpha_3^* \\ & & & \alpha_1^* \end{bmatrix} \tag{9}$$

[3] 0.1% of observed visits were to a government provider beyond the closest two and 0.2% of observed visits were to an NGO provider beyond the closest three. These observations were dropped from the data.

We solve the multinomial probit model (including estimates for $\alpha_1^*, \alpha_2^*, \alpha_3^*$) using the GHK (Geweke–Hajivassiliou–Keane) solution algorithm.[4]

We estimate the model using two specifications. In the first specification, the four quality scores and the tenure of the clinician are entered directly. This corresponds to a model of learning such that the prior over input l for provider j at time t is $\tilde{r}_{ljt} = \beta_l \cdot r_{jl} + \beta_t \cdot t$, where r_{jl} is the true value of input l. This is not really a specification of learning, rather it states that patients know the quality of a provider but have a taste or distaste for tenure. In the second specification, we estimate an interaction between tenure and quality. This corresponds to a model of learning such that the prior over input l at time t for provider j is:

$$\tilde{r}_{ljt} = r_{jl} \cdot \frac{\beta_l}{1 + \exp(\gamma_l - t)} \tag{10}$$

This specification allows for a pattern in which the prior for input l at provider j is zero (average) when the provider is new and evolves towards the true value with t (tenure increases), following a two parameter logistic function. There are other possible patterns of learning, but this specification is general enough to include this reasonable form as well as no learning.

Table 2 shows the coefficients for some of the key variables from the two regression specifications. The variables used to estimate fees and the illness condition interactions with quality variables are shown in Table 3, 4 and 5.[5] Travel is a very important determinant of the choice of provider, but patients are willing to travel to further facilities to find higher quality. In the first specification, patients are willing to travel from 2 to 10 kms further for 1 standard deviation of quality. Since travel is round-trip, the actual cost of travel is twice this amount, and 5 kms of travel represents a cost to the patient that is approximately equivalent

[4] Geweke (1989, 1991); Hajivassiliou and McFadden (1998); Keane (1990, 1994) as expounded in Train (2003).

[5] The value of quality varies with patient characteristics as well as illness characteristics. To allow for this we use dummy variables for 23 different illnesses interacted with each of the four quality characteristics. The value of quality discussed in this paper is the value of average quality. The differences from this average for each of the illnesses tracked and each of the qualities measured are shown in Table 3 and 5 and the findings of these tables are discussed in Leonard et al. (2002)

Table 2: Coefficients for Probit model of choice of provider: key variables

	1		2	
variable	Coef.	std err	Coef.	std err
DIST($>$ 6km) †	**-1.000**	0.105	**-1.000**	0.132
RAINY DIST($>$ 6km)	**-0.269**	0.064	**-0.190**	0.082
CLOSE	**14.843**	1.478	**18.353**	2.341
DIST($\leq$ 6km)	0.409	0.333	1.130	0.531
INFRASTR	**8.603**	1.023	**9.321**	1.752
CONSULT	**1.971**	0.892	**9.033**	1.636
INJECT	**8.620**	1.420	**10.886**	1.961
PRESCRIP	**9.934**	1.397	**22.098**	2.931
N DRUGS	**7.004**	1.273	339.592	287.790
TENURE	**-1.339**	0.197	**-6.728**	0.757
γ (CONSULT)			1.552	0.739‡
γ (INJECT)			$-\infty$	$\star$
γ (PRESCRIP)			-0.510	0.304‡
γ (N DRUGS)			11.218	0.981‡
log likelihood	-2491.15		-1884.04	

Each regression is a multinomial probit regression over 5 choices of provider with Specification 1 representing no learning, but including a variable representing a taste or distaste for tenure. Specification 2 represents learning following a two–parameter logistic model, with a taste or distaste for tenure. All input scores (INFRASTR, CONSULT, INJECT, PRESCRIP, and N DRUGS) are normalized so that increasing values represent higher qualities.

Distance is specified through 4 variables and coefficients. Distance when total distance is greater than 6 kms, distance during the rainy season when total distance is greater than 6 kms, distance when total distance is less than 6 kms and a dummy variable indicating that the clinic is in the same village. 6 kms is approximately the distance patients would walk; trips beyond 6 kms require public transportation. Rainy season travel on public transportation is more expensive than dry season travel.

†The coefficient for distance traveled (over 6 kilometers) is set equal to -1 for both specifications.

‡With the logistic specification, a null hypothesis that the coefficient estimating the inflection point (γ_l) is equal to zero has no direct interpretation. Each of these estimates is significantly different from values that represent no learning over the range 0 to 10 years.

$\star$ The inflection point for INJECT is inderteminate but sufficiently negative that the slope of the logistic function is completely flat (see text for explanation.) We set this variable to $-\infty$. A test that it is significantly different from any relevant inflection point rejects with p-value of 0.000

to 4% of total per capita government expenditure on health care (Leonard et al., 2002, pp. 462). These are significant amounts of money that patients are willing to spend on quality.

Patients appear to be attracted to newer clinicians. This is not evidence of an experimentation motive because there are many possible reasons for the arrival of new clinicians. This may reflect the fact that new clinicians are younger, on average, and newly minted doctors may be better than older doctors. We were careful to investigate the procedure for replacing a clinician and found no evidence that this is ever correlated with the quality of the outgoing clinician.

The coefficients for the logistic in the second specification can be better interpreted graphically. We do not specify a learning regression for overall infrastructure quality (INFRASTR) because this does not vary over time with tenure. Figure 2 shows the predicted paths of learning based on the second specification. The coefficient for learning with INJECT suggest that there is no learning over time about the number of injections that a clinician is likely to give. In addition, the pattern of learning exhibited for N DRUGS (over use of drugs) suggests that there is no learning for 6 or 7 years and then a very large gain over the next few years, except that the coefficient for N DRUGS is not significant. Even though these two variables exhibit extreme differences in their patterns, we suggest they are due to the same underlying factor: both N DRUGS and INJECT should be readily observable qualities. A patient should be able to learn about how many drugs a new clinician is likely to prescribe and whether or not he tends to use injections after one or two visits at the most. On the other hand, the quality of consultation and the prescription (CONSULT and PRESCRIP) are not observable and we expect learning for these scores. Figure 2 shows that learning is faster for prescription quality but that both exhibit increasing preferences for constant quality, a pattern consistent with learning.

Regression model 1 is a restriction of regression model 2 in which the inflection point for the logistic function is set equal to a very small number and the logistic function is flat over

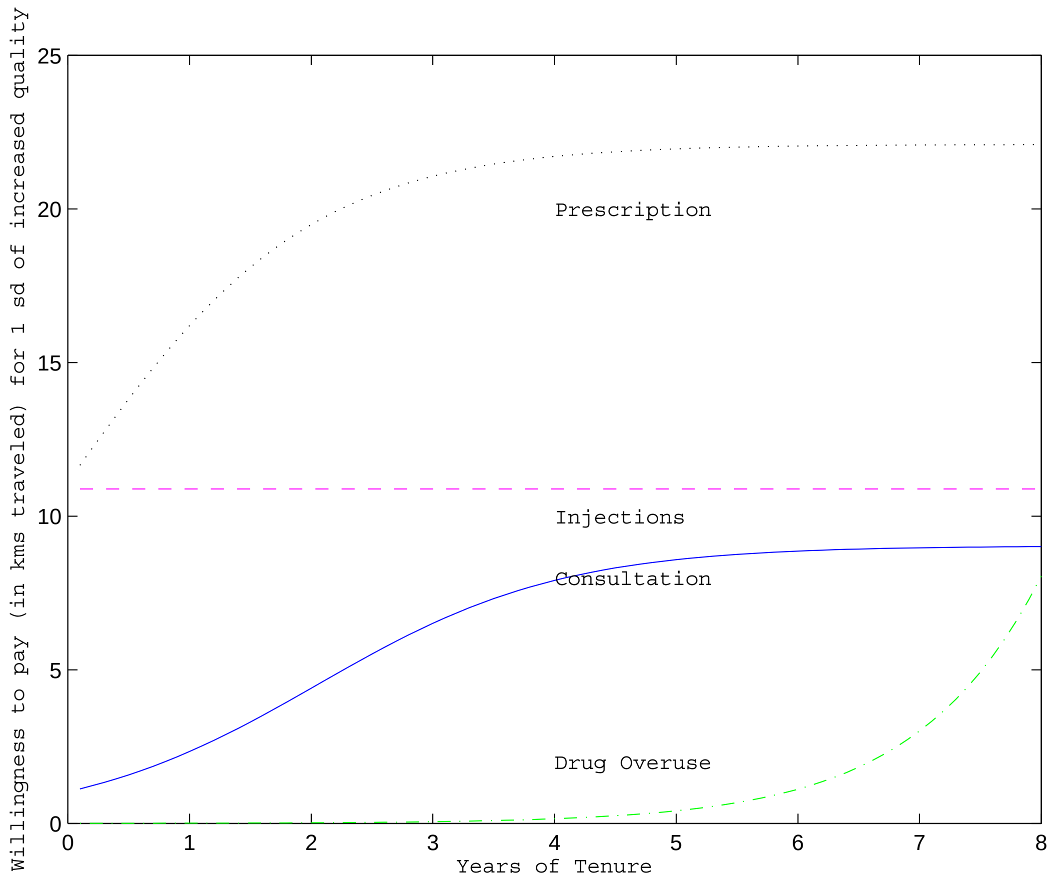

Figure 2: Graphical representation of the value of varying qualities as tenure increases

the relevant range.

$$\lim_{\gamma_l \to -\infty} \frac{\beta_l}{1 + \exp(\gamma_l - t)} = 1 \qquad \text{for } t \geq 0.$$

The likelihood ratio test of the joint hypothesis that these four restrictions are valid rejects the null hypothesis with a p-value of 0.000.

Our two specifications suggest that patients know about quality on average and that they increase their willingness to pay for quality as the tenure of clinicians rises. We suggest that this means they learn about quality gradually and increase their willingness to pay as their beliefs of quality approach certainty.

3.2 Examination of Residuals

The coefficients from the regression models above represent the best possible prediction (given the model) of behavior. However, there are errors in the predictions. Regression specification 1 represents the predicted behavior if everyone knew the true value at all times. The simple simulation represented in Figure 1 suggests that the deviation in behavior from the behavior expected if everyone had access to the true information should follow certain predictable patterns. To test this we generate residuals as the probability of choosing a practitioner minus the observed probability (0 or 1). Each observation of a patient choosing a practitioner generates 5 residuals associated with 5 different providers. This residual is the product of a number of different processes including, but not limited to, individual idiosyncrasies, measurement error in quality, and spatially correlated unobserved factors that influence the choice of provider (such as difference in income, livelihood, road condition, etc.). In addition to these factors, the residual can be influenced by local information about quality. Unlike the other factors influencing the residual, local information and the impact of local information should change with tenure.

We generate a two-dimensional smoothed autocorrelation function from these residuals. The correlation in residuals is calculated according to the distance between actors as well as the tenure of the clinician for each patient, facility by facility.[6] Figure 3 is the result of a local area regression of the correlation between residuals for three types of clinicians at distances varying from 1 to 42 kilometers. Three lines are drawn: the correlation in residuals for pairs of household/choices that share a choice over a particular clinician with less than two years of tenure, pairs for which one choice is a clinician with more than 5 years of tenure and all possible pairs.

The correlation pattern for all clinicians shows strong correlation in the behavior of

[6] 4,644 patients choosing between 5 facilities generates 539 million pairs. However, since patients are choosing between 46 different facilities, the set of pairs for which patients might choose the exact same facility represents only 3.8 million pairs. In addition, we restrict our attention to pairs between 0 and 42 kilometers apart.

Figure 3: Spatial correlation over distance for three types of clinician

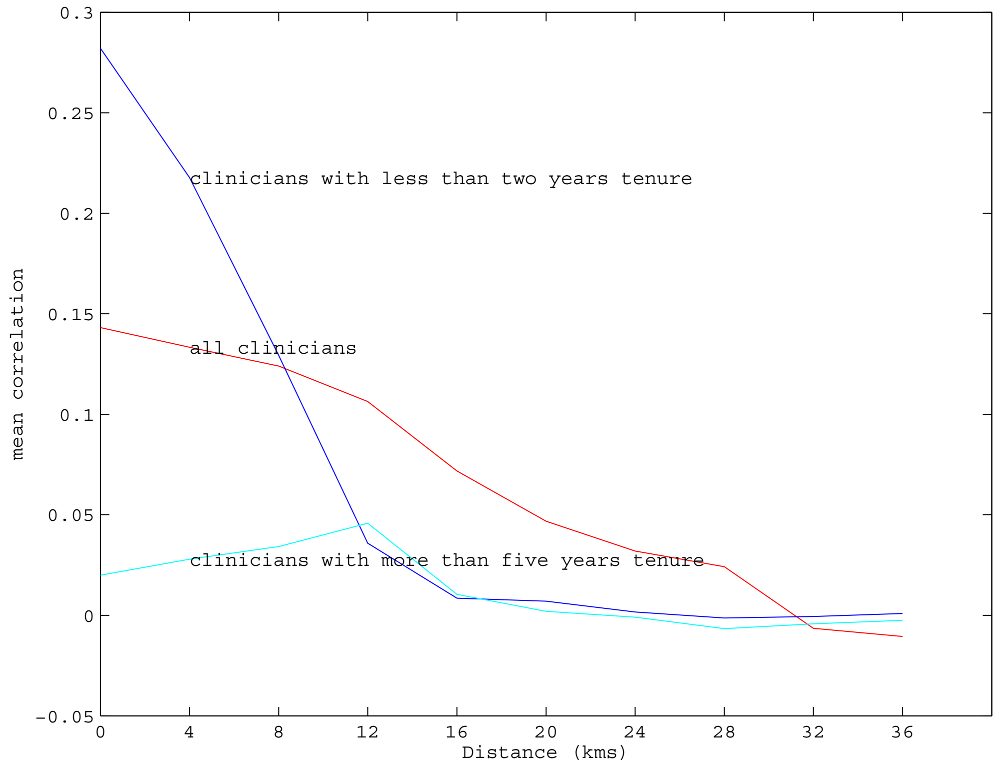

Correlations in the residuals for pairs of individuals located at varying distances from zero (same village) to 42 kilometers shown for clinicians with less than 2 years of tenure, greater than 5 years of tenure and for all clinicians combined. The lines shown are the points obtained from a local area regression with a uniform kernel with a bandwidth of ±6 kms centered on the distances listed on the x-axis.

patients who live close to each other, and much less correlation in the behavior of patients who are far apart. It is not surprising to find spatial correlation in the behavior of patients and the pattern could be due to any of the causes mentioned above. For example, two individuals 5 kms apart may share the same full set of possible facilities (the two nearest government and three nearest NGO facilities). Their preferences over one of these clinics might therefore have a lot in common. However, for two individuals 42 kms apart it is unlikely that the set of 5 providers between which they choose is exactly the same. The fact that the correlation pair exists means that they must have at least one facility in common, but not necessarily any other facilities. The reasons they visit or do not visit this provider are therefore much less likely to be similar. Thus, the pattern shown for all clinicians is not a function of learning, it is a function of environmental similarity; similarity of choice sets and other spatially correlated unobserved factors. The impact of learning is observed in the difference in correlation between new clinicians and established clinicians.

Figure 4 is a more complete picture of how correlation varies with tenure and distance. The pattern shown is almost exactly the pattern implied by Figure 1. Spatial correlation is highly correlated with the distance between actors and it varies with tenure. At close distances correlation starts above zero but increases to a peak at about 3 years and then decreases towards zero at 5 years and remains flat after that point. This is consistent with learning that begins almost immediately but that accumulates a stock of knowledge only gradually. Between 1 and 5 years a stock of knowledge is created that is local. Even though people are all learning about the quality of the same clinician, there is a stochastic component to their experience that means the knowledge created will be different. As knowledge improves beliefs become more similar across geographically isolated areas eliminating the difference across space, and approache the actual level of quality, reducing the size of the residual.

One aspect of the pattern of learning that we expect but do not find in this figure is strong correlation in behavior for individuals who live far apart when a clinician is new. If individuals have the same priors, their behavior should deviate from predicted behavior in approximately the same manner even if they live very far apart. The correlations shown in the region from 12 to 36 kms and for less than one year of tenure are significantly different from zero, but not very large. The impact of similar priors is most likely overwhelmed by the variance associated with environmental dissimilarity.

Figure 2 suggests that an overall picture of quality is developing between one and a half and 4 years, and that a good portion of learning is complete between 3 and 4 years. Figure 4 largely suggests the same pattern: local information about quality is causing large deviations in behavior between 2 and 3 years and deviations in behavior loose their strong spatial pattern after 4 years. These two patterns are derived from different features of the data, but show the same overall pattern of learning.

Figure 4: Spatial autocorrelation with respect to distance and tenure of the clinician

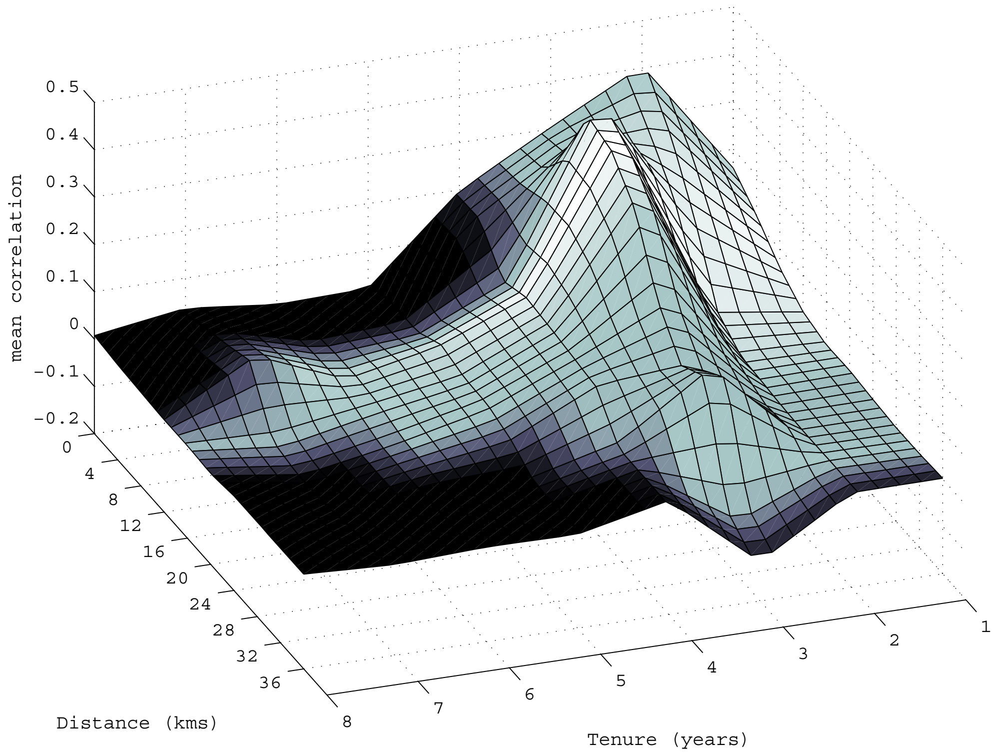

Correlations in the residuals for pairs of individuals located at varying distances from zero (same village) to 42 kilometers shown for tenure from 1 to 9 years. The surface represents a local area regression at the points shown with a uniform kernel and bandwidths of ±1 year of tenure and ±6 kilometers of distance. The shading is a function of the standard error of the point estimates. The light areas shown represent points that are significantly different from zero, whereas the darker shading is not significantly different from zero.

The data show that patients value objectively measured quality for new clinicians differently than they value objectively measured quality for established clinicians. The patterns are consistent with a learning story, but since we do not observe the patterns over a long period, it is hard to conclusively dismiss other possible patterns. Patient visits to each clinician are observed over a 12–month period, so the tenure of studied clinicians varies by up to one year; however we are trying to identify differences that are unlikely to be very large over such a small interval.[7] Another possibility is that the quality of clinicians could vary systematically over time. If clinicians improve their quality from 1 to 5 years that would explain the increased willingness to pay for quality and the deviations from predicted be-

[7] For the examination of residuals, patients are paired over the whole period observed, so the short panel variation in quality is effectively ignored.

havior. However, it would also require that quality be observable or that learning be nearly instantaneous, because patient behavior would have to follow changes in quality. In a separate research project in Arusha region of Tanzania, we have panel data on medical quality and have found no evidence of changes in quality over time. Therefore it seems unlikely that true quality is increasing in the manner shown and much more likely that perceptions of quality are following the patterns shown. Nonetheless, the short time duration of our panel does limit our confidence in these results.

4 Conclusion

Social learning is an important source of progress in developing countries and it should be an important element in health seeking behavior. Outcomes matter, are variable and are partially determined by unobservable quality. In addition, it is not likely that information about health seeking will suffer from strategic behavior; there may be costs to gathering and processing information, but it is not clear that there are gains to withholding information.

We have shown that patients have access to knowledge about outcomes that represents over 10 times the amount of information they could gather from their own experience; that patients demand quality; that their willingness to pay for quality increases with the tenure of clinicians and that local patterns in behavior are consistent with networks of social learning. Patients learn and use this information to improve their health.

These results do not tell us how households learn. Our model suggests that each household processes information about illnesses and maintains their own ideas about quality. It is far more likely that networks of social learning or villages mimic each others' behavior. Perhaps households do not rationally process information about quality, but villages or village elders process this information. The author's focus group interviews and interviews with elders in various African countries[8] certainly suggest that the latter is a more likely explanation. This does not mean that everyone in a village or network visits the same provider

[8] Tanzania, Gabon, Cameroun and Ethiopia.

(they do not), but rather that, given the same illness or circumstances, everyone would visit the same provider.

Health care is a rather unique situation and it is not clear that the results found here would generalize to other subjects. Outcomes are more readily observable because of both the public nature of illness and the discrete nature of outcomes. The culture of rural life in Africa (with relatively smaller variations in income or wealth within the village than might be found in Asia, for example) makes it highly unlikely that social exclusion is practiced within communities. There is no reason to hide information about quality from other households. In addition, almost all health seeking decisions — and certainly those for children — involve women at some point. Even if women are not the decision makers, they will possess all of the information about choices and outcomes. This may significantly impact the degree of communication.

The existence of social networks for learning in health care has important policy implications. The results suggest that households believe quality is important and that patients are willing to bypass poor quality facilities in order to find quality. It means that households can process information about quality. However, the average tenure of clinicians in our sample is about 5 years. If it takes four to 5 years to learn about the quality of clinicians, inefficient outcomes from uniformed priors are highly likely. Policy measures that seek to disseminate information may have positive impacts.

All health care services with which the author is familiar seek to obscure differences in training and cadre within their institution rather than to accentuate them. Throughout Africa, in an effort to show progress in service delivery many public services organizations, governments in particular, encourage the practice of obscuring titles. Patients are encouraged to refer to all medical personnel as ‘doctor’ and not to question the skills, training and ability of these clinicians. Whether patients know the difference between clinical assistants and doctors, they do appear to know the difference between good and bad practitioners. Rather than impede learning, government services in particular could benefit from paying more attention to the quality of care and the way in which information about the quality of care is disseminated.

References

Aghion, P. and P. Howitt, *Endogenous Growth Theory*, Cambridge: MIT Press, 1998.

Aoyagi, M., "Mutual Observability and the Convergence of Actions in a Multi-person Two-Armed Bandit Model," *Journal of Economic Theory*, 1998, *82* (2), 405–24.

Banks, J. S. and R. K Sundaram, "Switching Costs and the Gittins Index," *Econometrica*, 1994, *62* (3), 687–694.

Besley, T. and A. Case, "Diffusion as a learning process: evidence from HYV cotton," Working Paper, Princeton 1994.

Bolton, P. and C. Harris, "Strategic Experimentation," *Econometrica*, 1999, *67* (2), 349–74.

Brezzi, M. and T. L. Lai, "Optimal Learning and Experimentation in Bandit Problems," *Journal of Economic Dynamics and Control*, 2002, *27* (1), 87–108.

Conley, T. G. and C. Udry, "Learning about a new technology: Pineapple in Ghana," mimeo, under review, GSB, University of Chicago 2003.

Das, J., "Do Patients Learn about Doctor Quality? Theory and an Application to India," mimeo, Harvard University 2001.

Feder, G., R. Just, and D. Zilberman, "Adoption of Agricultural Innovations in Developing Countries: A Survey," *Economic Development and Cultural Change*, 1985, *33* (2), 255–298.

Foster, A. and M. Rosenzweig, "Learning by doing and learning from others: Human capital and technical change in agriculture," *Journal of Political Economy*, 1995, *103* (6), 1176–1209.

Geweke, J., "Bayesian inference in econometric models using Monte Carlo integration," *Econometrica*, 1989, *57*, 1317–1339.

_ , "Efficient simulation from the multivariate normal and Student-t distributions subject to linear constraints," in E. M. Keramidas, ed., *Computer Science and Statistics: Proceedings of the Twenty-Third Symposium on the Interface*, Fairfax: Interface Foundation of North America, Inc, 1991, pp. 571–578.

Gittins, J. C. and D. M. Jones, "A Dynamic Allocation Index for the Sequential Allocation of Experiments," in J. Gani et al., eds., *Progress in Statistics*, Amsterdam: North-Holland, 1974, pp. 242–266.

Hajivassiliou, V. and D. McFadden, "The method of simulated scores for the estimation of LDV models," *Econometrica*, 1998, *66*, 863–896.

Jovanovic, Boyan and Yaw Nyarko, "Learning by Doing and the Choice of Technology," *Econometrica*, 1996, *64* (6), 1299–1310. JSTOR link.

Keane, M., "Four essays in empirical macro and labor economics." PhD dissertation, Brown University 1990.

_ , "A computationally practical simulation estimator for panel data," *Econometrica*, 1994, *62*, 95–116.

Leonard, Kenneth L. and Joshua Graff Zivin, "Outcome versus service based payments in health care: lessons from African traditional healers," Working Paper W9797, National Bureau of Economic Research 2003. Forthcoming, *Health Economics*.

_ **, Gilbert Mliga, and Damen Haile Mariam**, "Bypassing Health Facilities in Tanzania: Revealed Preferences for Observable and Unobservable Quality," *Journal of African Economies*, 2002, *11* (4), 441–471.

Lucas, R. E. Jr., "On the Mechanics of Economic Development," *Journal of Monetary Economics*, 1988, *22*, 3–42.

Romer, Paul, "Increasing Returns and Long-Run Growth," *Journal of Political Economy*, 1986, *94* (5), 1002–1037.

Train, K., *Discrete Choice Methods with Simulation*, Cambridge: Cambridge University Press, 2003.

Table 3: Coefficients for Probit model of choice of provider: illness/resource interactions

	CONSULT		INJECT		PRESCRIP		N DRUGS	
variable	Coef	z-test	Coef	z-test	Coef	z-test	Coef	z-test
Overall (avg)	1.971	2.21	8.620	6.07	9.934	7.11	7.004	5.50
	Infants							
malaria	-0.769	-0.51	1.252	0.83	-3.595	-1.82	2.731	1.46
upper rti	-1.322	-0.40	-3.066	-0.96	1.529	0.37	0.161	0.03
pneumonia	-3.409	-1.88	-1.037	-0.19	-26.777	-4.76	-13.843	-1.57
cough	9.347	4.70	-1.277	-0.58	12.004	2.40	-8.165	-1.75
diarrhea	-1.950	-0.73	3.037	0.76	73.872	5.01	3.781	0.97
eye problem	-2.281	-0.53	-10.620	-2.40	0.657	0.10	-5.293	-1.01
injury	-0.008	0.00	-8.545	-0.82	37.921	1.13	-15.388	-0.84
dysentery	-1.029	-0.22	1.455	0.18	5.988	0.99	-5.942	-0.92
scabies	3.810	0.81	-3.669	-0.82	8.521	1.43	-1.938	-0.26
worms	0.413	0.04	-1.302	-0.21	-1.347	-0.14	4.021	0.45
	Non-Infants							
malaria	-1.095	-1.12	0.554	0.44	-0.694	-0.46	4.810	3.47
upper rti	-4.835	-3.82	-4.463	-2.51	1.591	0.75	2.000	1.05
injury	-2.467	-1.09	1.599	0.61	0.888	0.31	4.211	1.64
pneumonia	0.161	0.11	-0.116	-0.06	8.715	2.95	-2.508	-1.11
cough	3.083	1.91	2.752	1.60	5.475	2.16	-6.483	-2.94
dysentery	2.843	1.45	9.483	2.99	5.182	1.66	-2.512	-0.86
diarrhea	-5.459	-2.14	-11.070	-3.27	8.314	1.73	2.409	0.60
worms	3.012	1.37	-4.190	-0.82	6.450	1.34	-3.112	-0.76
severe abd pain	3.198	1.58	6.736	2.45	2.855	0.67	-4.930	-1.26
eye problem	-4.198	-1.37	3.112	0.79	-8.269	-1.86	0.607	0.14
pelvic infl dis	1.178	0.35	-5.450	-1.20	5.306	0.97	-9.169	-0.84
skin inf	-0.074	-0.03	-1.351	-0.62	-8.690	-1.61	9.037	1.64
STD	2.793	0.44	-3.835	-1.07	23.826	1.05	7.934	0.74

Table 4: Coefficients for Probit model of choice of provider: estimated drug costs

	Regression I				Regression I			
	NGO I		NGO II		NGO I		NGO II	
variable	Coef.	z-test	Coef.	z-test	Coef.	z-test	Coef.	z-test
constant	-1.616	-0.39	28.959	5.34	4.337	0.93	28.702	5.12
age (years)	-0.083	-0.99	-0.153	-1.60	-0.048	-0.56	-0.159	-1.63
infant (0-5)	11.624	2.58	10.504	1.88	9.716	2.17	6.466	1.14
child (5-10)	6.905	2.24	1.029	0.22	5.697	1.81	-1.428	-0.30
youth (10-16)	-0.021	-0.01	-6.657	-1.47	-1.240	-0.44	-8.202	-1.86
female	1.037	0.72	-0.855	-0.48	0.870	0.59	-0.423	-0.24
malaria	-12.062	-2.22	4.302	0.75	-9.063	-1.56	4.245	0.71
upper rti	16.628	2.36	-3.174	-0.31	12.087	1.60	-3.587	-0.35
pneumonia	-72.157	-4.16	-71.393	-3.52	-63.712	-3.79	-71.298	-3.37
cough	6.736	0.52	22.041	2.02	11.981	0.87	20.654	1.78
diarrhea	107.298	5.58	155.483	3.91	124.889	6.63	177.048	4.65
eye problem	-19.174	-1.59	-49.603	-3.91	-15.288	-1.24	-45.558	-3.26
injury	0.244	0.01	14.566	0.21	0.385	0.01	3.846	0.06
dysentery	16.116	1.23	-1.798	-0.07	17.385	1.18	0.106	0.00
scabies	3.736	0.24	-4.215	-0.27	1.014	0.06	-5.245	-0.40
malaria	-8.446	-2.31	7.248	1.60	-9.163	-2.30	2.683	0.59
upper rti	-13.455	-2.58	-22.727	-3.43	-16.135	-2.74	-27.870	-4.00
injury	-11.107	-1.77	-3.471	-0.41	-10.582	-1.55	-7.817	-0.94
pneumonia	-14.017	-2.48	-15.872	-2.00	-14.159	-2.42	-15.572	-1.99
cough	5.525	0.79	12.773	1.68	5.452	0.71	7.764	0.97
dysentery	24.649	3.58	32.430	2.92	25.442	3.09	34.563	2.98
diarrhea	4.074	0.39	-1.587	-0.15	0.476	0.05	-8.947	-0.77
worms	0.731	0.07	1.186	0.09	-0.952	-0.10	-3.091	-0.27
severe abd pain	-1.781	-0.18	-11.802	-0.42	-3.822	-0.33	-15.240	-0.65
eye problem	-14.409	-1.32	-8.318	-0.57	-13.413	-1.13	-11.553	-0.81
Pelvic Infl. Dis.	16.140	1.61	5.500	0.34	12.284	1.19	2.048	0.13
skin inf.	-20.161	-1.72	-30.194	-1.89	-24.245	-1.60	-33.572	-2.25
STD	13.740	0.49	8.079	0.18	17.034	0.67	9.394	0.24

Coefficients for government normalized to zero. Regression I is with no learning specification. Regression II is learning specification.

Table 5: Coefficients for Probit model of choice of provider with learning: illness/resource interactions

	CONSULT		INJECT		PRESCRIP		N DRUGS	
variable	Coef	z-test	Coef	z-test	Coef	z-test	Coef	z-test
Overall (avg)	2.861	2.27	16.750	8.06	8.760	6.11	10.389	7.11
	Infants							
malaria	-1.909	-1.25	-2.904	-1.26	1.260	0.82	1.279	0.66
upper rti	-3.442	-0.89	-1.293	-0.28	-2.264	-0.68	1.990	0.32
pneumonia	-4.591	-2.35	-23.048	-4.09	-1.139	-0.23	-20.255	-2.25
cough	9.788	4.38	16.668	3.04	-1.563	-0.66	-13.111	-2.82
diarrhea	-3.188	-1.22	81.697	5.96	2.649	0.72	2.918	0.66
eye problem	-2.582	-0.6	0.096	0.01	-10.317	-2.53	-5.803	-1.21
injury	-1.750	-0.24	33.168	1.10	-8.130	-0.84	-18.404	-0.87
dysentery	-2.072	-0.44	3.976	0.59	0.760	0.09	-7.192	-1.11
scabies	2.716	0.60	5.712	0.91	-4.019	-0.84	-2.476	-0.36
worms	-0.938	-0.07	-2.644	-0.29	-1.774	-0.24	6.014	0.54
	Non-Infants							
malaria	-2.207	-2.19	-2.120	-1.21	-0.024	-0.02	3.663	2.61
upper rti	-5.822	-4.14	-0.947	-0.38	-4.745	-2.4	1.014	0.51
injury	-3.611	-1.65	-0.274	-0.09	0.986	0.41	2.178	0.9
pneumonia	-0.534	-0.35	6.928	2.21	-0.288	-0.14	-4.423	-1.9
cough	2.139	1.23	6.077	2.13	2.764	1.64	-10.183	-4.11
dysentery	2.250	1.09	5.293	1.43	10.606	3.04	-3.120	-1.00
diarrhea	-6.067	-2.15	5.163	1.00	-9.813	-2.73	2.053	0.49
worms	2.111	0.99	4.255	0.87	-4.481	-0.99	-4.514	-1.23
severe abd pain	2.192	0.93	1.760	0.32	7.197	2.47	-8.462	-1.99
eye problem	-6.087	-2.15	-9.380	-2.01	2.625	0.72	-2.423	-0.57
pelvic infl dis	0.168	0.05	1.341	0.25	-5.413	-1.30	-8.447	-0.79
skin inf	-1.447	-0.52	-11.221	-1.68	-2.351	-1.00	9.413	1.54
STD	0.918	0.17	21.490	1.06	-3.688	-1.14	5.048	0.59

Glossary

ΔEU_{ijk} Net Expected Utility. Expected gain in utility from seeking health care, for individual i seeking care at provider j for illness condition k.

ε_{lk} Elasticity of healthiness for a particular illness condition k to resource l.

F_{jk} Expected fee and drug costs of illness condition j at provider k.

r_l The quantity of resource l provided.

$\hat{r}_{s(t)}$ The quantity of skill (resource) provided at an untrained provider (or a traditional healer).

T_{ij} Travel cost of individual i to provider j.