



AgEcon SEARCH

RESEARCH IN AGRICULTURAL & APPLIED ECONOMICS

The World's Largest Open Access Agricultural & Applied Economics Digital Library

This document is discoverable and free to researchers across the globe due to the work of AgEcon Search.

Help ensure our sustainability.

Give to AgEcon Search

AgEcon Search

<http://ageconsearch.umn.edu>

aesearch@umn.edu

*Papers downloaded from **AgEcon Search** may be used for non-commercial purposes and personal study only. No other use, including posting to another Internet site, is permitted without permission from the copyright owner (not AgEcon Search), or as allowed under the provisions of Fair Use, U.S. Copyright Act, Title 17 U.S.C.*

No endorsement of AgEcon Search or its fundraising activities by the author(s) of the following work or their employer(s) is intended or implied.

ECONOMIC GROWTH CENTER

YALE UNIVERSITY

P.O. Box 208629
New Haven, CT 06520-8269
<http://www.econ.yale.edu/~egcenter/>

CENTER DISCUSSION PAPER NO. 1000

**Learning, Misallocation, and Technology Adoption: Evidence from New
Malaria Therapy in Tanzania**

Achyuta Adhvaryu
Yale University

September 2011

Notes: Center discussion papers are preliminary materials circulated to stimulate discussion and critical comments.

This paper can be downloaded without charge from the Social Science Research Network electronic library at: <http://ssrn.com/abstract=1923891>

An index to papers in the Economic Growth Center Discussion Paper Series is located at: <http://www.econ.yale.edu/~egcenter/publications.html>

Learning, Misallocation, and Technology Adoption: Evidence from New Malaria Therapy in Tanzania

Achyuta R. Adhvaryu*

September 1, 2011

Abstract

I show that malaria misdiagnosis, common in resource-poor settings, decreases the expected effectiveness of an important new therapy—since only a fraction of treated individuals have malaria—and reduces the rate of learning via increased noise. Using pilot program data from Tanzania, I exploit variation in the location and timing of survey enumeration to construct reference groups composed of randomly chosen, geographically and temporally proximate acutely ill individuals. I show that learning is stronger and adoption rates are higher in villages with more misdiagnosis. Subsidizing diagnostic tools or improving initial targeting of new technologies may thus accelerate uptake through learning.

JEL codes: I15, O12, O33

Keywords: technology adoption, learning, malaria, Tanzania

*Yale University, e-mail: achyuta.adhvaryu@yale.edu, web: <http://www.yale.edu/adhvaryu>. I am indebted to S. Patrick Kachur for his hospitality at the CDC Mission in Dar es Salaam, Tanzania, and to the CDC and the Ifakara Health Institute for access to their household survey data. I thank Prashant Bharadwaj, Michael Boozer, Rahul Deb, James Fenske, Adrian de la Garza, A. Mushfiq Mobarak, Anant Nyshadham, Mark Rosenzweig, Maher Said, T. Paul Schultz, and Christopher Udry for their advice and guidance. I thank participants in seminars at Cornell, Columbia, University of Chicago, Emory, Yale and NEUDC for helpful comments. I am grateful to Rashid Khatib and Mujobu Moyo for their assistance in gathering data. All errors are my own.

1 Introduction

That learning plays an important role in the adoption of new technologies has been demonstrated in a variety of empirical contexts¹. Equally important, and perhaps less understood, is the nature of the learning process—what determines when, how well or how quickly individuals learn. For example, how do institutional arrangements, cultural norms and distribution policies change the way individuals learn? Can targeted roll-out of a new technology generate faster uptake for the population as a whole?

The answers to these questions may be particularly relevant in the developing world, where effective technologies are sometimes adopted at low rates, and where their widespread adoption could spur substantial increases in productivity and welfare. A large literature has documented this fact as it relates to agricultural innovations, and recent studies have shown that the same is true for the case of health technologies, as well. From piped water (Michael Kremer, Jessica Leino, and Alix Zwane 2010) to insecticide-treated bed nets (Jessica Cohen and Pascaline Dupas 2010; Dupas 2010) to de-worming drugs (Kremer and Edward Miguel 2007) to less polluting stoves (Mark Pitt, Mark Rosenzweig and Md. Nazmul Hassan 2006; Esther Duflo, Michael Greenstone and Rema Hanna 2008), there is a long list of health innovations which have been shown to have high returns but are adopted slowly or not adopted at all. To understand the learning processes in these contexts may be to gain some insight into—and ultimately to deliver solutions for—the reasons underlying non-adoption.

In this study, I examine how the rates of learning and adoption are linked to the way in which new technology is allocated. The question of the allocation of new technology is of particular relevance to medical innovations, which are often distributed on the basis of a diagnosis: healthcare professionals—doctors, nurses and health workers—are often tasked with allocating treatment to individuals according to the results of disease diagnosis. I show that the extent to which a new technology is *misallocated*—that is, the extent to which it is given to individuals for whom its use is inappropriate—is negatively related to the rates of learning and adoption.

To formalize this notion, I introduce misdiagnosis into an otherwise standard social learning model, in which individuals learn over time about the effectiveness of a new therapy from the outcomes of past adopters. In this context, I show that misdiagnosis affects learning and adoption behavior in two ways. First, misdiagnosis scales down the expected benefits of adoption, since even if the new therapy were fully effective, individuals would only realize its benefits if they really had the disease. Second, misdiagnosis, as it generates noise, makes it more difficult for individuals to extract information about the new therapy's effectiveness from past adopters' outcomes.

¹See, for example, Andrew Foster and Mark Rosenzweig (1995), Kaivan Munshi (2004), Oriana Bandiera and Imran Rasul (2006) and Timothy Conley and Christopher Udry (2010).

I test the model's predictions empirically using data on the adoption of a new malaria therapy in Tanzania. Studying the link between diagnosis and learning in this context is fitting for at least two reasons. First, malaria has large economic consequences (Jeffrey Sachs and Pia Malaney 2002; Sok Chul Hong 2007; Hoyt Bleakley 2010; Adrienne Lucas 2010; David Cutler et al. 2010). Studies have highlighted the central role of effective malaria treatment in alleviating the loss of life and productivity due to malaria, yet in much of Africa and Southeast Asia, existing malaria therapies are ineffective due to the development of parasitic resistance (Baird 2005). New, effective therapy thus has the potential to generate large returns if it is used appropriately on a wide scale.

Second, misdiagnosis of malaria (which spurs the misallocation of treatment) is overwhelmingly common in the developing world, particularly in contexts in which appropriate diagnostic tools are not available (Reyburn et al. 2007). Studies of this phenomenon thus far have focused on documenting two negative consequences of misdiagnosis: 1) patients who have other diseases (which require specific treatment, e.g. antibiotics for pneumonia) are often inappropriately treated with antimalarials, leading to prolonged bouts of illness (Amexo et al. 2004); and 2) that the misallocation of therapy leads to the rapid spread of resistance (Kenneth Arrow, Claire Panosian and Hellen Gelband 2004). This study is the first to my knowledge to identify a *behavioral* channel by which misdiagnosis adversely affects the usefulness of new therapies: it decreases the speed of learning and thus discourages adoption.

Artemisinin-based combination therapy (ACT), the new treatment I study, is a highly important innovation for malaria control, because it is the most effective treatment available for the prevalent type of malaria parasite in many parts of the developing world (Arrow, Panosian and Gelband 2004). The allocation of ACT should ideally be based on an accurate diagnosis of malaria. If an individual tests positive for malaria, she should receive the therapy; otherwise, she should receive alternate care for the underlying cause of her symptoms. However, in areas of the developing world where malarial prevalence is high and adequate diagnostic tools are inaccessible, it is common for health professionals to allocate therapy based solely on the presence of fever, the primary symptom of malaria. Presumptive diagnosis, as this allocation policy is termed, has been shown to lead to a high rate of over-diagnosis of malaria (Reyburn et al. 2004).

Using household survey data from a pilot program in Tanzania, through which ACT was distributed at health facilities, I develop an empirical strategy to test the model's predictions. My strategy does two things to disentangle social learning from other sources of correlation between current health outcomes and future adoption choices. First, I exploit the plausibly exogenous timing of survey enumeration to construct reference groups for learning based on the geographic and temporal proximity of self-reported acute illnesses. Second, I compare the correlation in current outcomes and future healthcare choices across sick individuals in treatment and comparison districts, and before and after the new therapy's introduction. The results show that

the probability of future adoption increases by about 5 percentage points when current adopters are sick for 1 less day than current non-adopters.

To test for the role of misdiagnosis, I categorize villages in the sample as high- and low-misdiagnosis locales using village-level baseline malarial prevalence data. This categorization is based on the observation that since misdiagnosis (in the context of this pilot program) can be measured by the fraction of fevers which are *not malarial*, places with lower prevalence will misdiagnose fevers more often. I show that apart from prevalence of malaria, these two groups are not statistically different from one another on observable dimensions. I find, as the theory predicts, that the rate of adoption and rate of learning are lower in villages with greater misdiagnosis.

This paper makes three main contributions. First, while empirical investigations of learning are increasingly common in the economics literature, few studies have sought to examine the determinants of the rate of learning. Understanding what drives variation in the magnitude of the learning effect across contexts is an important endeavor as policymakers seek new ways to encourage the rapid take-up of effective technologies.

Second, the results of this study have implications for policies related to the distribution of new medical technologies. Especially in parts of the world with a deficit of accurate diagnostic tools, my findings suggest that improving the quality of diagnosis can generate faster take-up and acceptance of new and useful treatments. This implication has particular relevance for the distribution of ACT in Africa and Southeast Asia. Ensuring that ACT reaches at-risk populations and is accepted as an effective therapy is a virtual necessity for adequate malaria control. This study's results suggest that subsidies for rapid diagnostic tests for malaria may be effective in promoting the widespread acceptance of ACT.

Finally, this study emphasizes the optimal *initial allocation* of technology as a mechanism to promote adoption via learning, which has thus far been ignored. The basic insight that misallocation and adoption are linked through learning may be applicable to a variety of technological innovations. For example, the adoption of high-yielding variety (HYV) maize has been studied extensively, particularly since the new varieties have still not been adopted on a wide scale in many developing countries. Since HYV maize grows best only in certain types of soil, allocation rules which are not based on soil type may generate significant heterogeneity in the returns to the new variety. When soil type is unobserved, this heterogeneity makes learning about the quality of the new variety noisier, and thus the speed of learning decreases. Similarly, if farmers themselves do not fully know if their soil type will be compatible with the new HYV, the heterogeneity in returns will lower the expected benefit of adopting. Designing better allocation policies based on testing for soil type could thus improve the speed of learning as well as the rate of adoption.

The rest of the paper is organized as follows. Section 2 develops the model. Section 3 de-

scribes the pilot program, the data, and the context. Section 4 develops an empirical strategy to test the model’s predictions. Section 5 reports the results, and section 6 concludes.

2 Model

In this section, I develop a simple social learning model of adoption behavior in which individuals learn about the effectiveness of new malarial treatment from their own health outcomes and those of their neighbors who have adopted in the past. In each period, acutely ill individuals make adoption choices based on the common prior on the new treatment’s effectiveness and the costs and (potentially heterogeneous) returns to adoption. Part of this return depends on the misdiagnosis of malaria, which occurs with known probability. I show that in this context, misdiagnosis negatively affects the rate of adoption and the speed of learning.

2.1 Setup

The model is in discrete time, and time periods are indexed by t . Consider a village composed of a set N of individuals, indexed by i , with $|N| = n$. In each period, a randomly chosen subset $N_t \subseteq N$ of individuals ($|N_t| = n_t \leq n$) falls acutely ill, and each individual in N_t makes an adoption decision. Each acutely ill individual in each period realizes a health outcome, which is observed, along with his adoption choice, by all villagers.

The model begins in period 0, at which time a new therapy of unknown quality $\theta \in \{0, 1\}$ is introduced. Individuals learn about this quality (or alternatively, effectiveness) parameter over time by observing the history of adoption choices and realized health outcomes. Since information is observed perfectly, all individuals update in the same way.² For every period $t > 0$, the timing of the model is as follows:

1. All individuals enter period t with a common belief distribution, summarized by q_t , over quality.
2. A subset N_t of villagers fall acutely ill, and each draws an unobserved malarial status M_{it} .
3. Each acutely ill individual makes an adoption choice h_{it} .
4. The resulting outcomes and adoption choices $\{D_{it}, h_{it} | i \in N_t\}$ are observed by all individuals.
5. The common belief distribution is updated, and a posterior belief q_{t+1} on the probability of effectiveness is formed.
6. Period $t + 1$ begins, and the process repeats.

²For simplicity of exposition, we also assume a common initial belief q_0 .

2.2 Definitions

Let h_{it} denote the binary adoption choice, where $h_{it} = 1$ denotes adoption, and $h_{it} = 0$ denotes non-adoption.

For simplicity, consider only two possible health outcomes: good ($D_{it} = D^g$) and bad ($D_{it} = D^b$). If we think of D as the length of illness in days, then a reasonable ordering would be $D^g < D^b$.

We denote the unknown quality (or effectiveness) of the new treatment as $\theta \in \{0, 1\}$. The common period- t belief about the new therapy's effectiveness is $q_t = \Pr(\theta = 1 | q_{t-1}, \{h_{it-1}, D_{it-1} | i \in N_{t-1}\})$.

$M_{it} \in \{0, 1\}$ is unobserved malarial status: $M_{it} = 1$ indicates the presence of malaria, and $M_{it} = 0$ indicates no malaria, and $m = \Pr(M_{it} = 1) \in (0, 1)$, which does not vary across i or t .

The probability of good and bad health outcomes being realized depends on the effectiveness of treatment, adoption choice, and malarial status:

$$D_{it} = h_{it} \left(M_{it} \left(\theta D^g + (1 - \theta)(pD^g + (1 - p)D^b) \right) + (1 - M_{it}) \left(\tilde{p}D^g + (1 - \tilde{p})D^b \right) \right) + (1 - h_{it}) \left(M_{it} \left(pD^g + (1 - p)D^b \right) + (1 - M_{it}) \left(\tilde{p}D^g + (1 - \tilde{p})D^b \right) \right). \quad (1)$$

If the individual adopts the new therapy ($h_{it} = 1$) and has malaria, he will recover quickly if $\theta = 1$ (i.e., if the new therapy is effective). If it is ineffective, he will recover with probability p , capturing the possibility that even ineffective therapy works some of the time.

If the individual does not adopt ($h_{it} = 0$), he recovers quickly with probability $p < 1$, reflecting the fact that alternative antimalarial treatments are relatively ineffective.³

If the individual does not have malaria, regardless of adoption, he will recover quickly (despite having adopted the wrong therapy) with probability \tilde{p} . This parameter captures the fact that of acutely ill individuals without malaria, some may recover regardless of intervention—for example, those who caught a common cold—while some may need specific treatment for the underlying causes of their fevers, e.g. in the case of pneumonia.

The impact of misdiagnosis on the learning process, as it turns out, depends on the relative magnitudes of p and \tilde{p} . Note from above that the lower is the effectiveness of the outside option (i.e., of existing therapy), the smaller p will be. The magnitude of \tilde{p} depends on the most prevalent causes of non-malarial fevers. This may differ significantly depending on geography, climate, demographic characteristics and baseline health of the population in question.

Evaluated over the distribution of M_{it} , the probabilities of receiving good or bad outcomes

³Note that for simplicity, we equate the effectiveness of non-adoption conditional on malaria with the effectiveness of adoption when $\theta = 0$ conditional on malaria.

when the therapy is of high or low are the following:

$$P(D_{it} = D^g | \theta = 1) = m + (1 - m)\tilde{p} \quad (2)$$

$$P(D_{it} = D^g | \theta = 0) = mp + (1 - m)\tilde{p} \quad (3)$$

$$P(D_{it} = D^b | \theta = 1) = (1 - m)(1 - \tilde{p}) \quad (4)$$

$$P(D_{it} = D^b | \theta = 0) = m(1 - p) + (1 - m)(1 - \tilde{p}). \quad (5)$$

2.3 Expected utility maximization

Utility is given as $u_i(C) - P(h)$, where the function u_i is increasing in consumption, C , and varies across individuals i . $P(h)$ is the price of health care at option h , and is measured in utils. The budget constraint is $C = w_i(\Omega_i - D)$, where w_i is the individual's wage rate and Ω_i is the amount of time he would work if fully healthy. This individual-level heterogeneity is perfectly observed by all individuals. The individual's expected utility maximization problem is thus $\max_{h \in \{0,1\}} \mathbb{E}(u_i(C) - P(h))$ subject to $C = w_i(\Omega_i - D)$.

Define $\bar{u}_i = u_i(w_i(\Omega_i - D^g))$ as utility under the good health outcome, and $\underline{u}_i = u_i(w_i(\Omega_i - D^b))$ as utility under the bad outcome. Expanding the expected value above using the definition of D from equation 1 and collecting terms, we can express the maximization problem as the following: individual i adopts in period t if and only if

$$q_t m(1 - p)(\bar{u}_i - \underline{u}_i) > P_1 - P_0. \quad (6)$$

Define $\Delta u_i = \bar{u}_i - \underline{u}_i$ and $\Delta P = P_1 - P_0$. The utility maximization problem can then be expressed as a simple cutoff rule: the acutely ill individual adopts if and only if the current-period prior on effectiveness exceeds a person-specific cutoff value:

$$h_{it} = \mathbf{1} \left(q_t > \frac{\Delta P}{m(1 - p)\Delta u_i} \right). \quad (7)$$

We will denote $\kappa_i = \frac{\Delta P}{m(1 - p)\Delta u_i}$. This cutoff responds in intuitive ways to changes in the model's parameters. An increase in the relative cost of adoption (ΔP) increases κ_i (i.e. makes adoption less likely). An increase in the rate of misdiagnosis ($1 - m$) increases κ_i . An increase in the effectiveness of the outside option p also increases the cutoff. Finally, an increase in the utility difference between quick and slow recovery from illness decreases κ_i .

2.4 Misdiagnosis and the adoption rate

In each period, let us denote the number of individuals who adopt as $n_{1t} = \sum_{i \in N_t} \mathbf{1}(q_t > \kappa_i)$ and those who do not as $n_{0t} = n_t - n_{1t}$. Define the period- t rate of adoption as $r_t = \frac{n_{1t}}{n_t}$, that is,

the fraction of sick individuals who adopt in a given period. Proposition 1 below states that as the rate of misdiagnosis ($1 - m$) increases, the rate of adoption decreases.

Proposition 1 r_t is weakly decreasing in $(1 - m)$.

Proof. Consider two levels of misdiagnosis, $1 - m' > 1 - m''$. From above, we know that κ_i is increasing in $1 - m$; thus $\kappa_i|_{1-m'} > \kappa_i|_{1-m''}$. But this implies that for a given q_t , $\sum_{i \in N_t} \mathbf{1}(q_t > \kappa_i|_{1-m'}) \leq \sum_{i \in N_t} \mathbf{1}(q_t > \kappa_i|_{1-m''})$, or $n_{1t}|_{1-m'} \leq n_{1t}|_{1-m''}$. Dividing by n_t on both sides, we obtain the desired result: $r_t|_{1-m'} \leq r_t|_{1-m''}$. ■

2.5 Misdiagnosis and the rate of learning

Next, we investigate how beliefs evolve over time through learning, and how misdiagnosis changes the learning process. Define the log-likelihood ratio of q_t as

$$\lambda_t = \log \left(\frac{q_t}{1 - q_t} \right). \quad (8)$$

From period to period, the log-likelihood ratio (equivalently, the belief q_t) evolves as individuals update the prior by incorporating new information contained in $\{h_{it-1}, D_{it-1} | i \in N_{t-1}\}$. Applying Bayes' rule, we can express the updating equation as:

$$\lambda_{t+1} = \lambda_t + \sum_{i \in N_t} h_{it} \log \left(\frac{P(D_{it} | \theta = 1)}{P(D_{it} | \theta = 0)} \right). \quad (9)$$

Using the expressions for the probabilities above from equations 2 through 5, we can write the updating equation as:

$$\lambda_{t+1} - \lambda_t = n_{1t}^g \log \left(\frac{m + (1 - m)\tilde{p}}{mp + (1 - m)\tilde{p}} \right) + n_{1t}^b \log \left(\frac{(1 - m)(1 - \tilde{p})}{m(1 - p) + (1 - m)(1 - \tilde{p})} \right), \quad (10)$$

where $n_{1t}^g = \sum_{i \in N_t} h_{it} \mathbf{1}(D_{it} = D^g)$ and $n_{1t}^b = \sum_{i \in N_t} h_{it} \mathbf{1}(D_{it} = D^b)$, so that $n_{1t}^g + n_{1t}^b = n_{1t}$.

Intuitively, if an individual adopts and realizes the good health outcome, the common prior on effectiveness should be revised upwards; if the adopter realizes the bad outcome, the opposite should happen. Finally, if the individual does not adopt, then no new information about effectiveness is revealed, and thus beliefs should not change.

To determine how misdiagnosis changes the rate of learning, following Christophe Chamley (2004), we examine the expected drift (taking the expected value over the distribution of M_{it}) in the log-likelihood ratio, conditional on $\theta = 1$: $\mathbb{E}(\lambda_{t+1} - \lambda_t | \theta = 1)$. Conditioning on $\theta = 1$ reflects the fact that the drift should be calculated for the true state, which in the case of effective therapy

is $\theta = 1$. Denoting $x_g = \log\left(\frac{m+(1-m)\tilde{p}}{mp+(1-m)\tilde{p}}\right)$ and $x_b = \log\left(\frac{(1-m)(1-\tilde{p})}{m(1-p)+(1-m)(1-\tilde{p})}\right)$, from equation 10 the expected drift can be expressed as:

$$\mathbb{E}(\lambda_{t+1} - \lambda_t | \theta = 1) = x_g \mathbb{E}(n_1^g | \theta = 1) + x_b \mathbb{E}(n_1^b | \theta = 1). \quad (11)$$

Now, we study how this expected drift varies with the rate of misdiagnosis, $1 - m$; this exercise enables us to understand how the rate of learning changes when misdiagnosis increases. The following proposition states that the way in which the expected drift varies with the misdiagnosis rate depends on the magnitudes of p and \tilde{p} , i.e., the extent to which the existing malarial treatment (the outside option) is effective, compared to the rate at which non-malarial fevers resolve without intervention. Intuitively, the proposition states that if the existing treatment is sufficiently ineffective, higher misdiagnosis will generate slower learning.

Proposition 2 For p, \tilde{p} such that $p < \tilde{p}$, $\mathbb{E}(\lambda_{t+1} - \lambda_t | \theta = 1)$ is weakly decreasing in $1 - m$.

Proof. Using equations 2 and 4, we can express $\mathbb{E}(n_1^g | \theta = 1)$ and $\mathbb{E}(n_1^b | \theta = 1)$:

$$\mathbb{E}(n_1^g | \theta = 1) = (m + (1 - m)\tilde{p}) n_{1t} \quad (12)$$

$$\mathbb{E}(n_1^b | \theta = 1) = (1 - m)(1 - \tilde{p}) n_{1t}. \quad (13)$$

Substituting the above expected value expressions into equation 11, we obtain

$$\mathbb{E}(\lambda_{t+1} - \lambda_t | \theta = 1) = \left(x_g (m + (1 - m)\tilde{p}) + x_b (1 - m)(1 - \tilde{p}) \right) n_{1t}. \quad (14)$$

Consider first the derivative of $B := x_g (m + (1 - m)\tilde{p}) + x_b (1 - m)(1 - \tilde{p})$ with respect to m in equation 14 above. This derivative can be expressed as

$$\frac{\partial B}{\partial m} = \frac{\partial x_g}{\partial m} (m + (1 - m)\tilde{p}) + \frac{\partial x_b}{\partial m} (1 - m)(1 - \tilde{p}) + (x_g - x_b)(1 - \tilde{p}). \quad (15)$$

Evaluating $\frac{\partial x_g}{\partial m}$ and $\frac{\partial x_b}{\partial m}$ and plugging the expressions into expression 15 above, we obtain

$$\frac{\partial B}{\partial m} = (\tilde{p} - p)(e^{x_g} - e^{x_b}) + (x_g - x_b)(1 - \tilde{p}). \quad (16)$$

Thus when $p < \tilde{p}$, $\frac{\partial B}{\partial m} > 0$, since $x_g > x_b$.

Now consider equation 14. Take two levels of misdiagnosis, $1 - m'' < 1 - m'$. The difference in $\mathbb{E}(\lambda_{t+1} - \lambda_t | \theta = 1)$ evaluated from $1 - m''$ (initial) to $1 - m'$ (final) is

$$\Delta \mathbb{E}(\lambda_{t+1} - \lambda_t | \theta = 1) = \Delta B \left(n_{1t} \Big|_{1-m''} \right) + \Delta n_{1t} \left(B \Big|_{1-m''} \right). \quad (17)$$

We know $\Delta B < 0$, since $\frac{\partial B}{\partial m} > 0$ for $p < \tilde{p}$. Proposition 1 demonstrates that $\Delta n_{1t} \leq 0$. Finally, $n_{1t}|_{1-m''} \geq 0$ and $B|_{1-m''} > 0$. Thus, $\Delta \mathbb{E}(\lambda_{t+1} - \lambda_t | \theta = 1) \leq 0$, as we set out to show. ■

2.6 Summary of predictions

In summary, this model makes two key predictions:

1. Greater misdiagnosis discourages adoption via the lower expected benefits of adoption.
2. Greater misdiagnosis decreases the rate of learning by introducing excess noise in the learning process.

In the following sections, I test these predictions in the context of the introduction of ACT in Tanzania.

3 Pilot program and data

3.1 Pilot program

The ACT pilot program, named the Interdisciplinary Monitoring Project for Antimalarial Combination Therapy in Tanzania, was implemented by The Centers for Disease Control and Prevention (CDC) and the Ifakara Health Institute in Rufiji, a rural district in southeast Tanzania, from February 2003 to the end of 2006.⁴ Under the auspices of the program, the piloted ACT therapy, artesunate plus sulphadoxine pyrimethamine, was prescribed to all individuals seeking care at government- or NGO-operated health facilities with fever or a recent history of fever. ACT was not available at any health care provider or store outside of government and NGO health facilities in the treatment district (Joseph Njau et al. 2008). Using this fact, I later define a proxy for ACT adoption based on health facility usage for individuals with self-reported acute illness.

3.2 Data

This study uses data from household surveys conducted before and after the introduction of ACT, in Rufiji, the treatment district, and in two comparison districts, Kilombero and Ulanga. Households in villages which were part of the Demographic Surveillance Systems in the treatment and comparison districts were sampled randomly to be surveyed. I use two years of data preceding the ACT intervention (2001 and 2002), and three years of post-intervention data (2004,

⁴For more details on the pilot and household survey, please refer to S. Patrick Kachur et al. (2001), Kachur et al. (2004) and Joseph Njau et al. (2008).

2005 and 2006). There was no survey round in 2003, the year that the new therapy was introduced. The treatment and comparison districts are geographically contiguous but separated by a large game reserve (the Selous Reserve).

I focus my analysis on the sample of individuals who reported being acutely ill with fever in the past two weeks. The individual module of the household survey asks questions about treatment-seeking following an episode of fever; I construct the dependent and independent variables of interest based on the answers of individuals who were acutely ill in the recent past to questions about their health care choices and health outcomes.

Table 1 reports the means and standard deviations of variables used in analysis for the sample of individuals who reported being acutely ill with fever starting in the two weeks prior to survey. I present summary statistics for the whole sample, and split by treatment versus comparison districts and high- versus low-misdiagnosis villages. High (low) misdiagnosis village is defined as a village in which the prevalence of *P. falciparum* was below (above) the median level at baseline (2001).

The total number of individuals sampled is just above 50,000, divided roughly evenly across the treatment and comparison districts. Of these individuals, 5826 reported being ill with fever that began in the two weeks preceding survey.

The unit for the age variable is years. It is constructed by subtracting the survey date from the date of birth of the individual, and dividing by 365.25. In all regression analyses, I include age decile fixed effects, as well as a quadratic in age, as controls. The average age in the sample is about 24.

The average years of completed education of household heads is approximately 4.5. The household survey only asked for the educational attainment of the household head. This attainment is reported in years, but for all analyses, I divide education into four categories, and include dummy variables for each category in the regressions (the omitted category is zero years of education):

1. *No formal education*: household head has zero years of formal education (28.9%)
2. *Less than primary education*: household head has greater than 0 and less than 7 years of formal education (23.3%)
3. *Primary education*: household head has exactly 7 years of education (39.3%)
4. *More than primary education*: household head has greater than 7 years of education (8.45%).

Pre-intervention (in 2001 and 2002), approximately 14 percent of the sample reported fever in the 2 weeks preceding survey. Among these individuals, about 30 percent sought care at a formal-sector health center, hospital or dispensary (public and NGO). In the treatment district, those who sought treatment at a health facility after acute illness received ACT. I thus use

formal-sector care usage as a proxy for ACT adoption in the treatment district. Njau et al. (2008) confirm, via surprise visits to health facilities and drug stores in the treatment district, that 1) there was no leakage of ACT into the informal sector, and 2) that all individuals presenting with fevers at government and NGO health facilities were prescribed ACT.

The rest of the individuals, those who did not visit the aforementioned options but who reported being acutely ill with fever, went to a medicine shop, street doctor, general store, kiosk, traditional healer, private laboratory, or could have used modern or traditional medicines from home or from a neighbor, or could have sought no treatment at all.

Among these individuals who sought health facility care, the length of their illness was just over 4 days. They reported about 1.5 additional symptoms besides fever. Among those who did not choose health facility care, the length of illness is shorter—just over 3 days—and the number of additional symptoms is approximately the same.

Columns 2 and 3 of Table 1 report summary statistics separately for the treatment and comparison districts. The average age among acutely ill individuals is slightly higher in the treatment district compared to the comparison, and the average educational attainment of household heads is lower. I address this imbalance in the empirical strategy section. On all the treatment-seeking and health outcome variables, the two samples are similar on average.

Columns 4 and 5 report summary statistics separately for high- and low-misdiagnosis villages. High (low) misdiagnosis village is defined as a village in which the malarial prevalence level is below (above) the median level at baseline. I use this distinction to test for the role of misdiagnosis in learning and adoption. A more complete definition is provided in section 4. Besides the mean malarial prevalence rate, which is by definition significantly larger in the low-misdiagnosis group, the two groups are similar across all the demographic and health variables on average.

4 Empirical strategy

The goal of this section is to develop an empirical strategy to test the implications of the learning model using data from household surveys before and after the introduction of the ACT pilot program. The main implications of the model are that where misdiagnosis of malaria is more prevalent, the adoption rate should be lower and learning—that is, the relationship between current-period adoption and previous-period health outcomes of adopters—should take place more slowly.

4.1 Defining high and low misdiagnosis villages

I construct areas of greater and lesser extents of misdiagnosis by exploiting baseline data on village-level malarial prevalence. Out of the 57 villages in the sample, 54 have baseline prevalence data. I divide these 54 villages into two groups—high and low prevalence—based on whether they fell above or below the village-level median of malarial prevalence in 2001. Since malaria treatment is prescribed presumptively at health facilities in both areas, higher malarial prevalence will imply a lower rate of misdiagnosis.

The main potential worry in making this classification is that high- and low-prevalence areas may be different on a variety of dimensions; these differences may generate variation in the rates of adoption and learning that are not due to misdiagnosis alone. At the end of this section, I present a set of analyses in support of the validity of this classification. In particular, I show that the villages in the high- and low-prevalence groups are not significantly different on average on observable dimensions, including the incidence of fever. We might be concerned that since misdiagnosis is defined as the proportion of fevers that are not malarial, if the fever rate highly correlates with malarial prevalence, then the effective rate of misdiagnosis will not be dissimilar across high- and low-prevalence areas. But in fact the correlation between malarial prevalence and fever at the village level is 0.0027; at the end of this section, I also show that there is no evidence that the *trends* in fever prevalence are differential across high and low misdiagnosis villages.

4.2 Misdiagnosis and the adoption rate

To test the prediction that misdiagnosis should decrease the ACT adoption rate, I estimate the extent of differential adoption in high versus low misdiagnosis villages. I use a difference in differences approach to measure the adoption rate, comparing health facility usage in the treatment and comparison groups before and after the intervention. I proxy for adoption with health facility usage for individuals reporting fever, since ACT was administered only at government and NGO health facilities, and was prescribed to all individuals presenting with a fever or history of fever. As noted above, ACT was not available in the informal sector and was indeed prescribed for fever or a history of fever at government and NGO health facilities (Njau et al. 2008).

I attribute the differential change in health facility usage over time in the treatment vis-a-vis the comparison group to the introduction of ACT. This differential change is thus a measure of the adoption rate in the sample of acutely ill individuals. The checks presented at the end of this section ensure that the differential trends across groups are not due to differential selection into acute illness.

To examine whether the adoption rate was different across high and low misdiagnosis vil-

lages, I first estimate the following difference in differences specification separately in the two samples:

$$h_{ijw} = \gamma T_j P_w + \alpha_w + \beta_j + \mathbf{X}'_{ijw} \delta + \epsilon_{ijw}. \quad (18)$$

The coefficient of interest is γ , the difference in differences estimate of the impact of ACT introduction on health facility usage. Here i denotes individual, j denotes village (and β_j are village fixed effects), w denotes round (wave) of survey (and α_w are round fixed effects), and \mathbf{X}_{ijw} is a vector of individual- and village-level controls, including the following variables: week of survey dummies to capture week-by-week seasonal variation; dummies for categories of educational attainment of the household head; and age decile dummies. We denote the dummy variable for health facility usage among acutely ill individuals as h_{ijw} , which equals 1 if the individual sought care at a health facility, and 0 if the individual sought care at an informal care option or did not seek care at all. T_j is a treatment district dummy and P_w is a post-intervention dummy, which equals 1 in post-intervention rounds.

I then interact the treatment x post-intervention term with a dummy for high misdiagnosis village (M_j equals 1 if the individual lives in a high misdiagnosis village, 0 otherwise), include the second-order interactions, and estimate a triple interaction specification in the pooled sample, measuring the differential adoption rate across high and low misdiagnosis villages:

$$h_{ijw} = \gamma_1 M_j T_j P_w + \gamma_3 M_j P_w + \gamma_4 T_j P_w + \alpha_w + \beta_j + \mathbf{X}'_{ijw} \delta + \epsilon_{ijw}. \quad (19)$$

Here the coefficient of interest is γ_1 , measuring the differential adoption across high- and low-misdiagnosis villages. Note that $M_j T_j$ and M_j will both be absorbed by the village fixed effects (β_j).

4.3 Identification of the learning effect

4.3.1 Definitions

We begin by defining the variables of interest. Suppose an individual i in village j falls sick with acute illness on date t . He makes an adoption choice, $h_{ijt} \in \{0, 1\}$ after falling ill, and his eventual health outcome, D_{ijt} , is measured as the length of illness in days. Note that if the individual is still ill when surveyed, the length of his illness will not be recorded (i.e., it will be coded as missing). I discuss the ramifications of this right-censoring of the distribution of the length of illness at the end of this section, and present evidence that it does not bias the estimate of the learning effect.

In line with the theory, the empirical model should reflect the intertemporal nature of the learning process: sick individuals should use the past health outcomes of adopters in their learn-

ing reference groups to update their priors on the quality of the new therapy. To construct the empirical analog, we must first define reference groups for learning.

My definition exploits the plausibly exogenous timing and location of survey enumeration to construct groups based on geographic and temporal proximity to the sick individual i . For geographic proximity I use the individual's village j , under the assumption that when individuals make healthcare choices, they learn from their fellow villagers who made similar choices in the past.⁵ For temporal proximity, I use information on the date the individual's acute illness began (t , as defined above). In particular, I assume that individuals falling sick on date t look back in time m days at the outcomes of adopters and non-adopters *in their village* who fell sick and made healthcare choices from date $t - m$ to date $t - 1$. At the end of this section, I present evidence that the order of survey enumeration was plausibly random.

Let N_{jt} be the set of individuals who fell ill in village j on date t . Let $N_{jt}^1 \subseteq N_{jt}$ denote the set of all individuals with fever in village j on date t who adopted the new therapy, and $N_{jt}^0 \subseteq N_{jt}$ denote those who did not adopt, such that $N_{jt}^1 \cup N_{jt}^0 = N_{jt}$. Then $\bar{D}_{j,(t-m,t)}^1$ and $\bar{D}_{j,(t-m,t)}^0$, the average length of illness for adopters and non-adopters, respectively, from dates $t - m$ to t , are defined as follows:

$$\bar{D}_{j,(t-m,t)}^1 = \frac{\sum_{a=1}^m \sum_{i \in N_{j,t-a}^1} D_{ij,t-a}}{\sum_{b=1}^m |N_{j,t-b}^1|} \quad (20)$$

$$\bar{D}_{j,(t-m,t)}^0 = \frac{\sum_{a=1}^m \sum_{i \in N_{j,t-a}^0} D_{ij,t-a}}{\sum_{b=1}^m |N_{j,t-b}^0|}. \quad (21)$$

I make four notes regarding these definitions. First, defining the reference group in this way implies that the group's choices and outcomes vary in general at the village x day level. Second, I only average over observations for whom the length of illness is recorded, i.e., for those whose acute illnesses are complete by the date of survey. At the end of this section, I check that right-censoring of the length of illness is non-differential across treatment and control districts and misdiagnosis categories.

Third, if, for some village j , choice $k \in \{0, 1\}$, and time span $(t - m, t)$, $\sum_{b=1}^m |N_{j,t-b}^k| = 0$, I replace the missing value of $\bar{D}_{j,(t-m,t)}^k$ with the average length of illness for all individuals who chose k between $t - m$ and t as calculated across the health facility catchment area (a group of spatially proximate villages). If this value is missing as well, then I replace it with the same average across the entire district (further broadening the definition of the individual's reference group for these observations). This process is necessary for less than 10 percent of acutely sick individuals.

⁵The survey instrument did not include data on social networks, which would have better reflected individuals' reference groups. On the other hand, the endogenous formation of social networks arguably introduces larger upward bias on the learning effect estimate than the endogenous formation of villages.

Fourth, in the estimates presented in section 5, I use a lag of 6 weeks ($m = 42$ days). I discuss the factors influencing this choice of lag length at the end of this section. As a check, I rerun the main analyses using lag lengths from 4 to 8 weeks, and the results are qualitatively similar (results available upon request).

4.3.2 Strategy

The theoretical model predicts that the learning process should generate a relationship between the past outcomes of adopters ($\bar{D}_{j,(t-m,t)}^1$) and the adoption choices of sick individuals in the current period (h_{ijt}). Of course, regressing the latter on the former would yield a biased learning effect estimate due to what Manski (1993) terms correlated effects. Individuals likely share common preferences for health with their reference group; have similar stocks of health as well as options for healthcare; and are exposed to the same local disease environment. Moreover, since this disease environment is often highly seasonal, outcomes and healthcare choices could be locally autocorrelated due to, for example, persistently heavy rainfall or high temperatures.

My empirical strategy aims to disentangle learning from the correlated effects described above. I take several steps to address bias arising from the fact that sick individuals and their reference groups have similar characteristics (health stocks, common shocks, preferences, availability of healthcare options, etc.).

First, I include village \times round-of-survey fixed effects (denoted η_{jw} , where w indicates survey round), which allow for flexible, village-specific trends in unobservables which may simultaneously affect current-period adoption and past-period health outcomes. Each survey round lasted approximately 3 to 4 months; thus there exists substantial day-to-day variation in h_{ijt} , $\bar{D}_{j,(t-m,t)}^1$ and $\bar{D}_{j,(t-m,t)}^0$ within fixed effect cells.

But even within village-by-survey-round cells, geographically and temporally local shocks—epidemics, weather fluctuations, drug stock-outs and the like—could potentially bias learning effect estimates. To deal with this possibility, the second aspect of the empirical strategy is to difference the past outcomes of adopters and non-adopters: $\Delta \bar{D}_{j,(t-m,t)} := \left(\bar{D}_{j,(t-m,t)}^1 - \bar{D}_{j,(t-m,t)}^0 \right)$. To the extent that sick adopters and non-adopters in the same reference group are affected equally by these common shocks, differencing their outcomes will remove the effect of the common shock from the health outcome measure.⁶

Finally, we must deal with the possibility that the health outcomes of adopters and non-adopters may indeed *not* react in the same way to shocks. Adoption is inherently driven by choice, and the unobserved characteristics of sick individuals—for example, the severity of their

⁶Thus we are positing that individuals learn from the *differential* outcomes of adopters as compared to non-adopters in their reference groups. Note that although this distinction cannot be made in the theoretical model, as for simplicity we assumed only one individual making an adoption choice per period, it is nevertheless crucial to make in our empirical setting, in which common autocorrelated shocks may generate significant bias in estimates of the learning effect.

illness or their preferences for health—likely drive their adoption choices, and will also be correlated with the way in which they react to a common shock. As a result, shocks which affect current-period choices may be correlated with $\Delta\bar{D}_{j,(t-m,t)}$ as well.

To account for this possibility, I exploit data on the choices and outcomes of 1) individuals in the comparison districts, and 2) individuals before the introduction of the new therapy. The intuition behind my strategy is that, for these individuals, the correlation between current-period healthcare choices (h_{ijt}) and previous-period (differential) health outcomes ($\Delta\bar{D}_{j,(t-m,t)}$) should represent only the spurious effects induced by common shocks, since the therapy was not introduced to these individuals, so no learning effect should be present for these groups. To purge the coefficient on $\Delta\bar{D}_{j,(t-m,t)}$ of these spurious effects, I interact the variable with a treatment x post-intervention dummy (denoted T_jP_t), as well as the main effects—treatment (T_j) and post-intervention (P_t). The resulting triple difference specification is:

$$h_{ijtw} = \left(\gamma_1 T_j P_w + \gamma_2 T_j + \gamma_3 P_w + \gamma_4 \right) \Delta\bar{D}_{j,(t-m,t)} + \eta_{jw} + \mathbf{X}'_{ijw} \delta + \epsilon_{ijtw}. \quad (22)$$

The learning effect is captured by the coefficient γ_1 . As mentioned before, η_{jw} are village x survey round fixed effects, and \mathbf{X}_{ijt} is a vector of individual and time-varying controls. \mathbf{X} includes the following variables: week of survey dummies to capture week-by-week seasonal variation; dummies for categories of educational attainment of the household head; and age decile dummies. Finally, \mathbf{X} includes lagged health average facility usage in the village, $\bar{h}_{j,(t-m,t)}$, and its interactions with treatment, post and treatment x post dummies.

4.4 Misdiagnosis and the rate of learning

To test the prediction that misdiagnosis decreases the rate of learning, I first estimate equation 22 separately in the high and low misdiagnosis samples. Then, I estimate a similar specification in the pooled sample to test for the difference in the learning effect across these samples, in which the coefficient of interest (α_1) is on the interaction of the high misdiagnosis village dummy with the learning effect estimate. The specification estimated in the pooled sample is:

$$h_{ijt} = \left(\alpha_1 T_j P_t + \alpha_2 T_j + \alpha_3 P_t + \alpha_4 \right) \Delta\bar{D}_{j,(t-m,t)} M_j + \left(\gamma_1 T_j P_t + \gamma_2 T_j + \gamma_3 P_t + \gamma_4 \right) \Delta\bar{D}_{j,(t-m,t)} + \eta_{jw} + \mathbf{X}'_{ijt} \delta + \epsilon_{ijt}. \quad (23)$$

Note that the main effect of the high misdiagnosis village dummy and its interactions with treatment, post-intervention and treatment x post-intervention are absorbed by the village x survey round fixed effects (η_{jw}).

4.5 Checks

4.5.1 Validity of classification of high- and low-misdiagnosis groups

Areas with different malarial prevalence rates may also be different in other ways associated with learning about and adopting new malaria therapy. I provide several pieces of evidence that the differences across these two groups (high and low misdiagnosis villages) along most dimensions are not large.

First, medical evidence suggests that the introduction of ACT does not affect a region's malarial ecology, including prevalence rates (Bousema et al. 2006). Second, as reported in Table 1, apart from malarial prevalence, which differs by definition across the two groups, along demographic and socioeconomic variables as well as healthcare choices and health outcomes (save for self-reported fever), the two groups are not different on average. Third, self-reported fever and village-level malarial prevalence are very weakly correlated: the correlation coefficient is 0.0027. The lack of correlation is necessary for the validity of the empirical strategy, since if fever and baseline malarial prevalence were very highly correlated in the sample, the rate of misdiagnosis ($1 - m$ over the prevalence of fever) would not vary enough across villages with differing malarial prevalence rates.

Fourth, I check that the *trends* in fever prevalence (which in my empirical context is equivalent to inclusion in the sample) are not differential across high and low misdiagnosis villages. If trends in sample selection were different across these two groups, we may worry that differential trends in ACT adoption may be due to differences in selection into acute illness rather than differences in the extent of misdiagnosis.

To check that this not indeed the case, I first estimate the following specification separately in the high and low misdiagnosis samples, and then in the pooled sample:

$$s_{ijw} = \gamma T_j P_w + \alpha_w + \beta_j + \mathbf{X}'_{ijw} \delta + \epsilon_{ijw}. \quad (24)$$

s_{ijw} is a dummy which equals 1 if individual i reported having fever in the 2 weeks preceding survey, and 0 otherwise. The results are reported in columns 1-3 of Table 2. The treatment x post-intervention interaction term is similar in magnitude across high- and low-misdiagnosis villages (columns 1 and 2), and is insignificantly different from 0 in both groups, as well as in the pooled sample regression (column 3).

Finally, I interact the treatment x post-intervention term with a dummy for high misdiagnosis village (M_j equals 1 if the individual lives in a high misdiagnosis village, 0 otherwise), and include the second-order interactions:

$$s_{ijw} = \gamma_1 M_j T_j P_w + \gamma_2 M_j T_j + \gamma_3 M_j P_w + \gamma_4 T_j P_w + \alpha_w + \beta_j + \mathbf{X}'_{ijw} \delta + \epsilon_{ijw}. \quad (25)$$

The results are reported in column 4 of Table 2. Based on the triple difference estimates, we find no evidence of differential selection into sickness across the two misdiagnosis groups.

4.5.2 Plausible randomness of survey enumeration

My strategy to estimate a learning effect relies on the fact that the ordering of survey enumeration was uncorrelated with observable household characteristics, the propensity for acute illness or treatment choice. We test this assumption directly, by regressing the date of survey (specifically the number of days since Jan. 1, 1960) on these variables. The econometric model we estimate is the following:

$$d_{ijt} = \delta' \mathbf{X}_{ijt} + \eta_{jw} + \epsilon_{ijt}. \quad (26)$$

d is the date of survey; \mathbf{X} are observable characteristics of individuals and households; η are village x survey round fixed effects, which, as described below, are a key part of the identification of the learning effect. The results of this estimation are reported in Table 3. In column 1, I estimate the above specification using the whole sample, and include a dummy variable for self-reported acute illness, in addition to observable characteristics (categories of educational attainment of the household head and a quadratic in age). The results in column 1 show that within village x survey round cells, there is no significant relationship between the date of survey and self-reported acute illness, as well as education and age.

Column 2 reports results of a similar specification estimated on the sample of acutely ill individuals, examining the relationship between date of survey and healthcare choice (a binary variable for seeking treatment at a formal-sector health facility), as well as the same education and age variables used in the previous specification. Again the results show no significant association between date of survey and these variables.

Columns 3 and 4 replicate columns 1 and 2, respectively, but use the actual integer ordering of survey enumeration (based on date of survey) within village x wave cells. That is, for each village in each wave, I order individuals according to when they were surveyed, with 1 being the earliest individual surveyed, 2 being the second earliest, and so on. This variable is then regressed on the same set of observable characteristics as before, as well as a village x wave fixed effect. The results, reported in columns 3 and 4, again show no systematic association between the order of survey enumeration and these variables.

Taken in sum, these results provide support for the plausible randomness of the timing of survey enumeration.

4.5.3 Non-differential trends in length of illness non-response

As mentioned above, since the length of illness is only recorded for those whose illnesses are complete, the right tail of the true distribution of this variable will be coded as a non-response. That is, conditional on the number of days prior to survey the illness began, the longest illnesses will be coded as missing and thus will not be used when calculating the average. This implies the average days of illness for the reference groups used to construct learning estimates will be underestimated.

Since my empirical strategy relies on a differencing approach, this underestimation will only be a problem if it occurs differentially across the treatment and comparison districts over time (and across high and low misdiagnosis villages when using a triple difference specification). To check that this is not the case, I construct a dummy variable c_{ijw} for individual i in village j surveyed in round w that equals 1 if the length of illness is not recorded due to right-censoring (i.e. due to the fact that the individual was still ill at the time of survey). I then estimate the following specification:

$$c_{ijw} = \gamma T_j P_w + \alpha_w + \beta_j + \mathbf{X}'_{ijw} \delta + \epsilon_{ijw}. \quad (27)$$

The coefficient of interest, γ , measures the extent to which trends in non-response in the length of illness were differential across treatment and comparison districts. The coefficient estimates are reported in column 1 of Table 4. The coefficient on treatment x post-intervention is small and not significantly different from 0.

In column 2, I report estimates of a triple difference specification, through which we can test whether the (difference in) trends in non-response in the length of illness was differential across high- and low-misdiagnosis villages:

$$c_{ijw} = \gamma_1 M_j T_j P_w + \gamma_2 M_j T_j + \gamma_3 M_j P_w + \gamma_4 T_j P_w + \alpha_w + \beta_j + \mathbf{X}'_{ijw} \delta + \epsilon_{ijw}. \quad (28)$$

The estimate of γ_1 reported in column 2 shows that this difference is not significantly different from 0. Thus, we find no evidence that the underestimation of average length of illness for reference groups is differential across treatment and comparison districts, and across high- and low-misdiagnosis groups.

4.5.4 Choice of lag length (m)

The choice of $m = 42$ reflects a balancing of two concerns. First, for a given lag length, there exist fewer observations of length of illness at the beginning of each survey round (which lasted 3-4 months for each year) than at the end. For example, if an individual is surveyed 1 week after the survey enumeration begins in a particular wave, the lagged variables for this individual is

calculated using only the outcomes of her fellow villagers who were surveyed in the previous week. For an individual surveyed near the end of a survey round, the lagged variables will use observations from the entire lag length. Therefore, the longer the lag length, the more relatively error-ridden the estimates at the beginning of the survey round will be compared to the end.

Second, while shorter lag length (i.e., a more temporally proximate reference group) may more accurately reflect reality, it also increases measurement error, for two reasons. First, with a smaller m , the number of observations included in the construction of the lagged variable decreases. Thus, when the number of observations over which I average to construct the lagged variable is low, the estimate of length of illness becomes more error-ridden. Second, if there are no observations in the given lag length, I fill in this missing number with the constructed average across a larger space than the village (first the health facility catchment area, and then the district on the whole). Thus, if the village-level average is missing and is filled in with a broader average, this value is less likely to reflect the individual's actual reference group outcomes in the previous period.

The 6-week ($m = 42$) lag length balances these two concerns. As mentioned above, I have rerun the main analyses using lag lengths of 4, 5, 7 and 8 weeks, and the results are qualitatively similar. These results are available upon request.

4.5.5 Controlling for lagged demographic and illness differences

Finally, I address the potential concern that the (differential) composition of adopters vis-a-vis non-adopters changed as a result of ACT introduction, rendering comparisons in the health outcomes of the two over time invalid. To the extent that these changes occurred on observable dimensions, we can control for the differential composition of adopters versus non-adopters, in terms of their demographic characteristics and the characteristics of their acute illnesses.

For a given characteristic x , define $\Delta\bar{x}_{j,(t-m,t)}$ as the difference in x across health facility users and non-health facility users in village j between dates $t - m$ and t . This difference is defined equivalently to the difference in the length of illness across adopters and non-adopters:

$$\bar{x}_{j,(t-m,t)}^1 = \frac{\sum_{a=1}^m \sum_{i \in N_{j,t-a}^1} x_{ij,t-a}}{\sum_{b=1}^m |N_{j,t-b}^1|} \quad (29)$$

$$\bar{x}_{j,(t-m,t)}^0 = \frac{\sum_{a=1}^m \sum_{i \in N_{j,t-a}^0} x_{ij,t-a}}{\sum_{b=1}^m |N_{j,t-b}^0|} \quad (30)$$

$$\Delta\bar{x}_{j,(t-m,t)} = \bar{x}_{j,(t-m,t)}^1 - \bar{x}_{j,(t-m,t)}^0. \quad (31)$$

I use age, education of the household head and a wealth index (generated via principal components analysis), and the number of additional self-reported symptoms (a measure of the severity of illness) as x variables.

I augment the baseline specification by adding these new variables ($\Delta\bar{x}_{j,(t-m,t)}$) and their interactions with treatment, post-intervention and treatment x post-intervention dummies. The resulting specification is:

$$h_{ijt} = \left(\gamma_1 T_j P_t + \gamma_2 T_j + \gamma_3 P_t + \gamma_4 \right) \Delta \bar{D}_{j,(t-m,t)} + \left(\alpha_1 T_j P_t + \alpha_2 T_j + \alpha_3 P_t + \alpha_4 \right) \Delta \bar{x}_{j,(t-m,t)} + \eta_{jw} + \mathbf{X}'_{ijt} \delta + \epsilon_{ijt}. \quad (32)$$

5 Results

5.1 Misdiagnosis and ACT adoption

Table 5 reports estimates of equations and , the difference in differences estimates of ACT adoption. Columns 1 and 2 report estimates of equation in the high and low misdiagnosis samples, respectively. The results in these columns show that the differential trend in health facility usage in high misdiagnosis villages had close to 0 slope, while the differential trend in low misdiagnosis villages was positive (though imprecisely estimated). The pooled sample estimates of equation , reported in column 3, confirm that these differential trends were significantly *different* across the two groups. In particular, the adoption rate was approximately 16 percentage points lower in high misdiagnosis villages.

5.2 Estimates of the learning effect

The estimates of equation 22 are reported in Table 6. Column 1 reports the estimates of this baseline specification. The learning effect estimate is the coefficient on the interaction of the differential illness length across adopters and non-adopters x treatment district x post-intervention. The learning effect estimate is large relative to the mean health facility usage rate and is precisely estimated. The interpretation of this estimate is that narrowing the difference in the length of illness across adopters and non-adopters by 1 day *increases* the future adoption probability by about 4.8 points.

Columns 2 and 3 report the results of robustness checks, in which lagged differences in demographic characteristics and symptoms across adopters and non-adopters, constructed in the same way as the lagged differences in the length of illness, are interacted with treatment and post-intervention dummies.

In column 2, I estimate the above specification using age, education of the household head and a wealth index (generated via principal components analysis) as x variables. The results make clear that the addition of these lagged differences and their interactions do no affect the magnitude or significance of the learning effect estimate.

In column 3, I again estimate the above specification using the demographic characteristics mentioned above, as well as the number of additional self-reported symptoms, a measure of the severity of illness. Note that these data were only collected for in 3 of the 5 survey rounds, so the sample in this estimation is smaller. Nevertheless, the results reported in column 3 show that the magnitude of the learning effect estimate remains fairly stable (though it is estimated with less precision than the baseline estimate).

5.3 Misdiagnosis and the learning effect

Table 7 reports estimates of equation 22 separately for the high and low misdiagnosis samples, and the estimates of equation 23 for the pooled sample. As the results reported columns 1 and 2 show, the learning effect is small (-0.0248) and insignificantly different from 0 in high misdiagnosis villages, but is larger and precisely estimated (-0.0735) in villages with a lower rate of misdiagnosis. The pooled sample estimates of equation 23, reported in column 3, confirm that the learning effect is significantly different in high versus low misdiagnosis villages.

6 Conclusion

In this study, I demonstrate how the acceptance and adoption of effective technologies can hinge on the way in which they are allocated. In the case of malaria therapy, I show that the misdiagnosis of malaria affects individuals' beliefs and subsequent adoption patterns through learning. When individuals are uncertain about the effectiveness of new therapy, misdiagnosis of malaria makes it more difficult to extract a signal about quality from the health outcomes of adopters. It also scales down the expected benefit of the therapy, since individuals who are unsure that they have malaria know that even if the therapy is effective, they will only realize its benefits if they actually have the disease. In both these ways, poor diagnostic policy can discourage the adoption of new malaria therapy even if the therapy is clinically effective.

I develop a strategy to test these hypotheses empirically, using household survey data from a pilot program through which ACT was prescribed at health facilities in Tanzania. I find evidence that 1) places in which misdiagnosis is more common experienced lower adoption rates over time, and 2) individuals learn from the health outcomes of past adopters, but misdiagnosis decreases the extent of this learning.

Most new technologies take time to gain acceptance, due to uncertainty about their inherent effectiveness (Kremer and Miguel 2007) or about their optimal usage (Foster and Rosenzweig 1995, Conley and Udry 2010). Social learning has been shown to be a key mechanism by which individuals overcome this uncertainty, and come to adopt the technologies which are most productive and profitable for them. Yet there remain many examples of beneficial technologies and

behaviors that are not adopted despite their proven effectiveness. In the developing country context, these barriers to adoption can stifle technology-driven economic growth in the aggregate. I demonstrate that the inappropriate allocation of new technology is one such barrier, which can be removed by changing policy to target only those individuals for whom the technology is intended.

Leaders in the fight against malaria are recognizing the need for the proliferation of better diagnostic technology. For example, the World Health Organization has called for private manufacturers to produce effective, cheap rapid diagnostic tests (RDTs) for malaria, for use in rural settings in Africa and southeast Asia (WHO 2008). New evidence shows that subsidizing RDTs along with ACTs for distribution in the private sector can yield high uptake of ACT even while limiting its inappropriate use by patients who do not have malaria (Cohen, Dupas and Simone Schaner 2011). These results, together with this paper's findings, suggest that investments in subsidies for diagnostic technology may reap high returns, not only by limiting the inappropriate allocation of ACT to non-malarial patients, but also by encouraging the sustained adoption of ACT via learning.

References

- [1] Amexo, Mark, Rachel Tolhurst, Guy Barnish, and Imedda Bates. 2004. "Malaria misdiagnosis: effects on the poor and vulnerable." *Lancet*, vol. 364: pp. 1896-1898.
- [2] Arrow, Kenneth, Claire Panosian and Hellen Gelband. 2004. *Saving Lives, Buying Time: Economics of Malaria Drugs in an Age of Resistance*. Published by National Academies Press.
- [3] Bandiera, Oriana and Imran Rasul. 2006. "Social Networks and Technology Adoption in Northern Mozambique." *The Economic Journal*, vol. 116 (October): pp. 869890.
- [4] Baird, J. Kevin. 2005. "Effectiveness of Antimalarial Drugs." *The New England Journal of Medicine*, 352: pp. 1565-1577.
- [5] Bleakley, Hoyt. 2010. "Malaria Eradication in the Americas: A Retrospective Analysis of Childhood Exposure." *American Economic Journal: Applied Economics*, vol. 2, no. 2: pp. 1-45.
- [6] Bousema, J. Teun, Petra Schneider, Louis Gouagna, Chris Drakeley, Alma Tostmann, Rein Houben, John Githure, Rosalynn Ord, Colin Sutherland, Sabah Omar, and Robert Sauerwein. 2006. "Moderate Effect of Artemisinin-Based Combination Therapy on Transmission of *Plasmodium falciparum*." *Journal of Infectious Diseases*, vol. 193, no. 8: pp. 1151-1159.
- [7] Chamley, Christophe. 2004. *Rational herds: economic models of social learning*. Cambridge University Press, Cambridge, United Kingdom.
- [8] Cohen, Jessica and Pascaline Dupas. 2010. "Free Distribution or Cost-Sharing? Evidence from a Randomized Malaria Prevention Experiment." *Quarterly Journal of Economics*, vol. 125, no. 1: pp. 1-45.
- [9] Cohen, Jessica, Pascaline Dupas, and Simone Schaner. 2011. "Price Subsidies, Diagnostic Tests, and Targeting of Malaria Treatment." *working paper*.
- [10] Conley, Timothy and Christopher Udry. 2010. "Learning About a New Technology: Pineapple in Ghana." *American Economic Review*, vol. 100, no. 1: pp. 3569.
- [11] Cutler, David, Winnie Fung, Michael Kremer, Monica Singhal and Tom Vogl. "Early-Life Malaria Exposure and Adult Outcomes: Evidence from Malaria Eradication in India." *American Economic Journal: Applied Economics*, vol. 2, no. 2: pp. 72-94.
- [12] Duflo, Esther, Michael Greenstone and Rema Hanna. 2008. "Cooking Stoves, Indoor Air Pollution and Respiratory Health in Rural Orissa, India." *Economic and Political Weekly*, vol. 43, no. 32: pp. 71-76.

- [13] Dupas, Pascaline. 2010. "Short-Run Subsidies and Long-Term Adoption of New Health Products: Experimental Evidence from Kenya." *working paper*.
- [14] Foster, Andrew and Mark Rosenzweig. 1995. "Learning by Doing and Learning from Others: Human Capital and Technological Change in Agriculture." *Journal of Political Economy*, 103, pp. 1176-1209.
- [15] Hong, Sok Chul. 2007. "The Burden of Early Exposure to Malaria in the United States, 1850-1860: Malnutrition and Immune Disorders." *Journal of Economic History*, vol. 67, no. 4.
- [16] Kachur, S. Patrick, Rashid Khatib, Ellen Kaizer, Susan Fox, Salim Abdulla, and Peter Bloland. 2004. "Adherence to Antimalarial Combination Therapy with Sulphadoxine-Pyrimethamine and Artesunate in Rural Tanzania." *American Journal of Tropical Medicine and Hygiene*, 71(6): pp. 715-722.
- [17] Kachur, S. Patrick, Salim Abdulla, Karen Barnes, Hassan Mshinda, David Durrheim, Andrew Kitua, and Peter Bloland. 2001. "RE: complex and large trials of pragmatic malaria interventions." *Tropical Medicine and International Health*, 6, 324-325.
- [18] Kremer, Michael and Edward Miguel. 2007. "The Illusion of Sustainability." *Quarterly Journal of Economics*, 112(3): pp. 1007-1065.
- [19] Kremer, Michael, Jessica Leino and Alix Zwane. 2010. "Spring Cleaning: Rural Water Impacts, Valuation, and Institutions." *Quarterly Journal of Economics*, forthcoming.
- [20] Lucas, Adrienne. 2010. "Malaria Eradication and Educational Attainment: Evidence from Paraguay and Sri Lanka." *American Economic Journal: Applied Economics*, vol. 2, no. 2: pp. 46-71.
- [21] Manski, Charles. 1993. "Identification of Endogenous Social Effects: The Reflection Problem." *The Review of Economic Studies*, volume 60, no. 3, pp. 531-542.
- [22] Munshi, Kaivan. 2004. "Social Learning in a Heterogeneous Population: Technology Diffusion in the Indian Green Revolution." *Journal of Development Economics*, 73(1): pp. 185-215.
- [23] Njau, Joseph, Catherine Goodman, S. Patrick Kachur, Jo Mulligan, John Munkondya, Naiman Mchomvu, Salim Abdulla, Peter Bloland, and Anne Mills. 2008. "The costs of introducing artemisinin-based combination therapy: evidence from district-wide implementation in rural Tanzania." *Malaria Journal*, 7:4.
- [24] Pitt, Mark, Mark Rosenzweig and Md. Nazmul Hassan. 2006. "Sharing the Burden of Disease: Gender, the Household Division of Labor and the Health Effects of Indoor Air Pollution." *working paper*.

- [25] Reyburn, Hugh, Hilda Mbakilwa, Rose Mwangi, Ombeni Mwerinde, Raimos Olomi, Chris Drakeley and Christopher Whitty. 2007. "Rapid diagnostic tests compared with malaria microscopy for guiding outpatient treatment of febrile illness in Tanzania: randomized trial." *British Medical Journal*, 10: 1136.
- [26] Reyburn, Hugh, Redempta Mbatia, Chris Drakeley, Ilona Carneiro, Emmanuel Mwaka-sungula, Ombeni Mwerinde, Kapalala Saganda, John Shao, Andrew Kitua, Raimos Olomi, Brian Greenwood, and Christopher Whitty. 2004. "Overdiagnosis of malaria in patients with severe febrile illness in Tanzania: a prospective study." *British Medical Journal*, 329: 1212.
- [27] Sachs, Jeffrey and Pia Malaney. 2002. "The economic and social burden of malaria." *Nature*, vol. 415.
- [28] World Health Organization. 2008. "WHO Malaria Diagnostics Product Testing Programme - Call for expressions of interest (EOI) from manufacturers for product submission to Round Two." URL accessed August 20, 2008:
<http://www.who.int/tdr/svc/grants/calls/malaria-rdts>

Table 1

Summary statistics

	Whole sample		Treatment district (Rufiji)		Comparison districts (Kilombero & Ulanga)		High misdiagnosis villages		Low misdiagnosis villages	
	Mean	SD	Mean	SD	Mean	SD	Mean	SD	Mean	SE
Number of individuals	50474		24920		25554		23514		26960	
Number of individuals reporting fever in 2 weeks preceding survey	5826		3022		2804		2630		3196	
<i>All years:</i>										
Age in years	23.882	21.273	24.749	22.738	23.036	19.703	24.163	21.515	23.637	21.057
Educational attainment of household head (years)	4.402	3.314	3.453	3.481	5.303	2.869	4.468	3.346	4.346	3.285
<i>Pre-intervention years (2001 & 2002):</i>										
Proportion reporting fever in 2 weeks preceding survey	0.142	0.349	0.141	0.348	0.143	0.350	0.136	0.343	0.147	0.354
<i>Among individuals reporting fever:</i>										
Proportion who sought care at health facility	0.291	0.454	0.290	0.454	0.292	0.455	0.325	0.469	0.261	0.439
<i>Among individuals who sought health facility care:</i>										
Length of illness (days)	4.260	2.389	4.473	2.613	4.085	2.175	4.462	2.548	4.024	2.169
Number of additional symptoms	1.471	1.115	1.333	1.077	1.645	1.147	1.403	1.134	1.544	1.099
<i>Among individuals who did not seek health facility care:</i>										
Length of illness (days)	3.277	2.466	3.205	2.355	3.340	2.557	3.475	2.655	3.118	2.292
Number of additional symptoms	1.426	1.066	1.250	1.012	1.622	1.092	1.404	1.039	1.445	1.091

Notes: age in years calculated as (date of survey - date of birth)/365.25; additional symptoms besides fever include body pains, headache, diarrhoea, chills, cough, convulsions, vomiting, fainting, fast breathing, and dizziness; high (low) misdiagnosis village is defined as a village in which the prevalence of *P. falciparum* was below (above) the median level at baseline (2001).

Table 2

Sample selection

<i>Dependent var: 1 if individual reported having fever in two weeks preceding survey, 0 otherwise</i>				
	High-misdiagnosis villages	Low-misdiagnosis villages	Whole sample	Whole sample
Treatment x post	0.0122 (0.0159)	0.0202 (0.0153)	0.0161 (0.0106)	0.0203 (0.0141)
<i>High-misdiagnosis village x</i> Treatment x post				-0.00940 (0.0223)
Post				0.00673 (0.0173)
Fixed effects	<i>Village & survey round</i>			
Number of observations	25729	22232	47961	47961

Notes: *** p<0.01, ** p<0.05, * p<0.1; all specifications are estimated using OLS; robust standard errors are reported in parentheses below coefficient estimates, and allow for correlation in the error term within villages; specifications control for week-of-survey dummies to capture seasonal variation, dummies for categories of educational attainment of the household head, age decile dummies and a quadratic function in age; note that village fixed effects subsume the high-misdiagnosis village dummy and its interaction with treatment district.

Table 3

Timing of survey enumeration

<i>Dependent variable:</i>	<i>Date of survey</i> <i>(number of days since January 1, 1960)</i>		<i>Integer order of survey</i> <i>(within village x survey round cell)</i>	
	Whole sample	Acutely ill individuals	Whole sample	Acutely ill individuals
Reported fever in two weeks preceding survey	0.326 (0.463)		-1.019 (1.122)	
Sought treatment at health facility		-0.255 (0.202)		1.055 (2.691)
<i>Educational attainment of household head:</i>				
Less than primary (< 7 years)	-1.397 (1.672)	0.0249 (0.0856)	-0.572 (1.691)	-2.021 (2.177)
Completed primary school (= 7 years)	0.530 (0.338)	-0.122 (0.167)	-0.714 (2.055)	-0.337 (2.048)
More than primary school (> 7 years)	0.557 (0.358)	0.00558 (0.165)	-1.318 (2.362)	-3.441 (3.880)
Age in years	0.114 (0.114)	0.0127 (0.0106)	-0.00139 (0.0255)	0.0587 (0.121)
Age squared	-0.00123 (0.00126)	-0.000154 (0.000139)	-0.000163 (0.000481)	-0.00186 (0.00165)
Fixed effects		<i>Village x survey round</i>		
Number of observations	47961	5505	47961	5505

Notes: *** p<0.01, ** p<0.05, * p<0.1; all specifications are estimated using OLS; robust standard errors are reported in parentheses below coefficient estimates, and allow for correlation in the error term within village x survey round; specifications control for week-of-survey dummies to capture seasonal variation, dummies for categories of educational attainment of the household head, age decile dummies and a quadratic function in age.

Table 4

Trends in length of illness non-response

<i>Dependent var: 1 if length of illness right-censored for individual (i.e. individual's acute illness was not complete on the day of survey), 0 otherwise</i>		
	Whole sample	Acutely ill individuals
Treatment x post	0.0359 (0.0379)	-0.0187 (0.0540)
<i>High-misdiagnosis village x</i>		
Treatment x post		0.112 (0.0685)
Post		-0.00200 (0.0483)
Fixed effects	<i>Village & survey round</i>	
Number of observations	5,500	5,500

Notes: *** p<0.01, ** p<0.05, * p<0.1; all specifications are estimated using OLS; robust standard errors are reported in parentheses below coefficient estimates, and allow for correlation in the error term within villages; specifications control for week-of-survey dummies to capture seasonal variation, dummies for categories of educational attainment of the household head, age decile dummies and a quadratic function in age; note that village fixed effects subsume the high-misdiagnosis village dummy and its interaction with treatment district.

Table 5

Trends in ACT adoption over time for high and low misdiagnosis villages

<i>Dependent var: 1 if individual sought care at a government or NGO health facility, 0 otherwise</i>			
	<i>Difference in differences estimates</i>		<i>Triple difference estimates</i>
	High misdiagnosis villages	Low misdiagnosis villages	Whole sample
Treatment x post	-0.0254 (0.0578)	0.117 (0.0687)	0.120* (0.0672)
<i>High misdiagnosis village dummy x</i>			
Treatment x post	-	-	-0.161* (0.0834)
Post	-	-	0.0428 (0.0705)
Fixed effects	<i>Village & survey round</i>		
Number of observations	2471	3034	5505

Notes: *** p<0.01, ** p<0.05, * p<0.1; all specifications are estimated using OLS; robust standard errors are reported in parentheses below coefficient estimates, and allow for correlation in the error term within villages; specifications control for week-of-survey dummies to capture seasonal variation, dummies for categories of educational attainment of the household head, age decile dummies and a quadratic function in age; note that village fixed effects subsume the high-misdiagnosis village dummy and its interaction with treatment district.

Table 6

Estimates of the learning effect

Dependent var: 1 if individual sought care at a government or NGO health facility, 0 otherwise

	Baseline specification	Baseline + lagged differences in demographics	Baseline + lagged differences in demographics and symptoms
<i>Difference in lagged sickness length in days (health facility users - non-health facility users) x</i>			
Treatment*Post dummy	-0.0475*** (0.0164)	-0.0452*** (0.0168)	-0.0428* (0.0259)
Treatment dummy	0.0303*** (0.0101)	0.0335*** (0.0101)	0.0268 (0.0182)
Post dummy	0.0287** (0.0114)	0.0272** (0.0116)	0.00946 (0.0185)
Difference in lagged sickness length in days (health facility users - non-health facility users)	-0.0178*** (0.00652)	-0.0221*** (0.00649)	-0.0101 (0.0136)
Fixed effects		<i>Village x survey round</i>	
Number of observations	5294	5166	2958

Notes: *** p<0.01, ** p<0.05, * p<0.1; all specifications are estimated using OLS; reference group for learning is the individual's village; 6-week lags are used, beginning on the day the individual reported falling ill; proportion of ref group who visited health facility, and its interactions with post, treatment, and treatment x post dummies, are included in the specification; robust standard errors are reported in parentheses below coefficient estimates, and allow for correlation in the error term within village x date of survey; additional controls are week-of-survey dummies to capture seasonal variation, dummies for categories of educational attainment of the household head, age decile dummies and a quadratic function in age; column 1 reports results from the baseline learning effect specification; columns 2 and 3 test robustness to addition of lagged differences (across facility users and non-users) in demographics (age, education and asset index) and number of symptoms, respectively; columns 1 and 2 are estimated on full sample; column 3 is run on the sample of individuals from 2001, 2004 and 2006, for whom symptoms data were collected.

Table 7

Estimates of the learning effect for high- and low-misdiagnosis villages

<i>Dependent var: 1 if individual sought care at a government or NGO health facility, 0 otherwise</i>			
	High misdiagnosis villages	Low misdiagnosis villages	Whole sample
Learning effect estimate	-0.0248 (0.0232)	-0.0735*** (0.0229)	-0.0789*** (0.0226)
Learning effect x high misdiagnosis village dummy	-	-	0.0634** (0.0316)
Fixed effects	<i>Village x survey round</i>		
Number of observations	2346	2948	5294

Notes: *** p<0.01, ** p<0.05, * p<0.1; all specifications are estimated using OLS; reference group for learning is the individual's village; columns 1 and 2 report learning effect estimates, which are the coefficients on the difference in lagged sickness length in days (health facility users - non-health facility users) interacted with treatment x post dummy in high and low misdiagnosis villages, respectively; main effect of this difference and its interactions with treatment and post-intervention dummies are included in all specifications; prop visiting health facility in ref group and its interactions with post, treatment, and treatment x post are included in all specifications; column 3 reports results of an interaction specification, in which the learning effect is interacted with a high misdiagnosis village dummy; robust standard errors are reported in parentheses below coefficient estimates, and allow for correlation in the error term within village x date of survey; additional controls are week-of-survey dummies to capture seasonal variation, dummies for categories of educational attainment of the household head, age decile dummies and a quadratic function in age.