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

Achyuta R. Adhvaryu*

February 28, 2012

Abstract

This study examines how misallocation of new technology affects learning and adoption in the context of a new malaria therapy in Tanzania. Misdiagnosis of malaria reduces average therapeutic effectiveness, since not all adopters actually have malaria, and slows the rate of social learning due to increased noise. I exploit data on adoption choices, the timing and duration of acute illnesses, and biometric measures of malarial status to show that individuals whose reference groups experienced idiosyncratically fewer misdiagnoses exhibit stronger learning effects and are more likely to adopt. Improving initial targeting by subsidizing diagnostic technology may thus accelerate adoption through learning.

JEL codes: I15, O12, O33
Keywords: technology adoption, social 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 contexts[1]. 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; Alessandro Tarozzi et al. 2011) 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 that 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.

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

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.

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 allocated and used appropriately on a wide scale.

Second, misdiagnosis of malaria–and the resultant 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). For example, in one study in Tanzania, blood slide data showed that more than half the individuals receiving treatment for malaria at government hospitals were not actually infected (Hugh Reyburn et al. 2004). Studies examining the rate of misdiagnosis among individuals purchasing antimalarial therapies from informal private-sector providers (e.g., pharmacies or drug shops) find similar magnitudes (Jessica Cohen, Pascaline Dupas, and Simone Schaner 2011).

Misdiagnosis has grave consequences in both the short and long run for the health and economic well-being of affected populations. In the short run, inappropriate treatment induces losses in economic productivity, since patients are slower to recover when they are not treated

for the underlying cause of their symptoms (Mark Amexo et al. 2004). For young children, the effects can be much more severe: fever misdiagnoses account for an estimated three million child deaths each year that could be prevented if timely and appropriate treatment were administered. Misallocation of treatment also results in less effective usage of public and private healthcare resources. Estimates of the expenditure outlays per fever episode–including the prices of medicine, time and capital–imply that for each case of malaria that is appropriately treated, 3-5 dollars in public funds are wasted on an inappropriately treated case (not factoring in the additional costs of delayed recovery and subsequent healthcare for misdiagnosed patients) (Pascalina Chanda et al. 2011). In the long run, the over-prescription and overuse of proven therapies engendered by inadequate access to diagnostic technology leads to the rapid development of resistance (Kenneth Arrow, Claire Panosian and Hellen Gelband 2004). The costs of repeatedly engaging in research and development in hopes of finding the next effective antimalarial are staggering.

In this study, I identify a novel, behavioral channel through 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 nearly 20 percentage points when current adopters are sick for 1 less day than current non-adopters.

To test for the role of misdiagnosis, I exploit blood slide data on individual malarial diagnoses to calculate the rate of misdiagnosis faced by each individual's reference group. I define this rate as the proportion of fevers in the individual's reference group that were *not* malarial. Since households were randomly sampled and the timing of survey enumeration was plausibly exogenous (an assumption for which I provide some proof in the data), the inferred misdiagnosis rate should be an unbiased estimate of the actual rate faced by the reference group. I then compare the size of the learning effect for individuals whose reference groups by chance experienced different levels of misdiagnosis. I find, as the theory predicts, that the rate of adoption is lower and the learning effect smaller for individuals whose reference groups faced a higher misdiagnosis level.

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 clear 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. My results support the findings of Cohen et al.'s (2011) recent experimental study in calling for subsidies for rapid diagnostic tests (RDTs) for malaria. Not only would RDTs limit the negative clinical consequences of misallocating treatment, but they might also be effective in promoting the widespread acceptance and adoption of ACT.

Finally, this study emphasizes the optimal allocation of technology as a mechanism to promote adoption via learning, which has thus far been ignored. The basic insight that misallocation and subsequent 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, for example, on testing for soil type could thus improve the speed of learning as well as the rate of adoption.

The remainder of the paper is organized as follows. Section 2 develops the model. Section 3 describes 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 7 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. Each acutely ill individual draws at random a malarial status, which is unobserved to the individual himself, as well as to the other villagers. Each individual in N_t then makes an adoption decision, and realizes a health outcome. The adoption choice and the ensuing health outcome are observed by all villagers.[2] Information contained in the choices and outcomes of acutely ill individuals in period t is then used to form a posterior belief on the effectiveness of the new therapy; this posterior is then carried over into period $t + 1$.

2.2 Definitions

We begin by defining the following terms:

- Let $M \in \{0, 1\}$ be a random variable determining malarial status. $M = 1$ indicates the presence of malaria, and $M = 0$ indicates no malaria. Let $m = \Pr(M = 1)$. M_{it} is the realization of M for individual $i \in N_t$.

- $h_{it} \in \{0, 1\}$ denotes the adoption choice for $i \in N_t$, where $h_{it} = 1$ denotes adoption, and $h_{it} = 0$ denotes non-adoption.

- Let $D \in \{D^b, D^g\}$ be a random variable determining the length of illness, where $D^g < D^b$, i.e., the good health outcome is a speedier recovery from illness. D_{it} is the realization of D for $i \in N_t$.

[2] In addition, we assume that individual-specific characteristics, which play a role in the utility maximization problem described below, are common knowledge.

- The common period-t belief about the new therapy's effectiveness is:

  $q_t = \Pr(\theta = 1 | q_{t-1}, \{h_{it-1}, D_{it-1} | i \in N_{t-1}\})$.

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

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

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 with probability $\tilde{p}$ (despite having adopted the wrong therapy). This parameter captures the fact that some acutely ill individuals who do not have malaria 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, the probabilities of receiving good or bad outcomes when the therapy is of high or low are the following:

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

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

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

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

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

2.3 Timing

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. We assume that in $t = 0$, all individuals begin with a common initial belief, $q_0 = \Pr(\theta = 1)$. 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 a malarial status $M_{it} \in \{0, 1\}$, which is unobserved to the ill individual himself, as well as to the other villagers.

3. Each acutely ill individual makes an adoption choice $h_{it} \in \{0, 1\}$.

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.

2.4 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 $\max_{h\in\{0,1\}} \mathbb{E}_t \left(u_i(C) - P(h)\right)$ subject to $C = w_i(\Omega_i - D)$, where $\mathbb{E}_t$ is the expectation taken using all known information up to and including period t.

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

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

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:

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

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.

2.5 Misdiagnosis and the adoption rate

In each period, let us denote the number of individuals who adopt as $n_{1t} = \sum_{i\in N_t} \mathbf{1}\left(q_t > \kappa_i\right)$ 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, $\sum_{i \in N_t} \mathbf{1}\left(q_t > \kappa_i|_{1-m'}\right) \leq \sum_{i \in N_t} \mathbf{1}\left(q_t > \kappa_i|_{1-m''}\right)$, or $n_{1t}|_{1-m'} \leq n_{1t}|_{1-m''}$. Dividing by n_t on both sides, we obtain the desired result: $r_t|_{1-m'} \leq r_t|_{1-m''}$. ■

2.6 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

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

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:

$$\lambda_{t+1} = \lambda_t + \sum_{k \in \{g,b\}} \sum_{i \in N_t} h_{it} \mathbf{1}(D_{it} = D^k) \log \left(\frac{\Pr(D_{it} = D^k | \theta = 1)}{\Pr(D_{it} = D^k | \theta = 0)} \right) . \tag{9}$$

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

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

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. Whether beliefs about effectiveness go up or down from t to $t+1$ depends on the proportion of adopters experiencing good and bad outcomes, and the magnitudes of the terms in logs, which reflect *how much* the belief

should be scaled up or down for each individual adopter who experiences, respectively, a good or bad outcome.

To determine how misdiagnosis changes the rate of learning, we examine the expected drift in the log-likelihood ratio conditional on $\theta = 1$, denoted as $\mathbb{E}(\lambda_{t+1} - \lambda_t|\theta = 1)$ (Christophe Chamley 2004).[4] Let $x_g := \log\left(\frac{m+(1-m)\tilde{p}}{mp+(1-m)\tilde{p}}\right) > 0$ and $x_b := \log\left(\frac{(1-m)(1-\tilde{p})}{m(1-p)+(1-m)(1-\tilde{p})}\right) < 0$. From equation 10 the expected drift can be expressed as:

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

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 *When $p \leq \tilde{p}$, $\mathbb{E}(\lambda_{t+1} - \lambda_t|\theta = 1)$ is weakly decreasing in $1 - m$.*

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

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

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

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

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

Consider first the derivative of $\Gamma := x_g\,(m + (1 - m)\tilde{p}) + x_b(1 - m)(1 - \tilde{p})$ with respect to m

[4]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 $\theta = 1$.

in equation 14 above. This derivative can be expressed as

$$\frac{\partial \Gamma}{\partial m} = \frac{\partial x_g}{\partial m}\left(m + (1-m)\tilde{p}\right) + \frac{\partial x_b}{\partial m}(1-m)(1-\tilde{p}) + (x_g - x_b)(1-\tilde{p}). \tag{15}$$

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

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

Thus when $p \leq \tilde{p}$, $\frac{\partial \Gamma}{\partial m} > 0$, 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

$$\Delta\mathbb{E}(\lambda_{t+1} - \lambda_t | \theta = 1) = \Delta\Gamma n_{1t}|_{1-m''} + \Delta n_{1t}\Gamma|_{1-m''}. \tag{17}$$

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

2.7 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.[5] 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 survey rounds in which blood slide data were collected, namely, the 2001 pre-intervention round and the 2004 post-intervention round. 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

[5]For 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).

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.

The total number of individuals in the selected rounds is just above 14,000, divided roughly evenly across the treatment and comparison districts. Of these individuals, 1891 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. I 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 23.5.

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 (2001), approximately 17 percent of the sample reported fever in the 2 weeks preceding survey. Among these individuals, about 25 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 average duration of illness was about 4.7 days. They reported about 1.5 symptoms apart from 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.

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 when 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 Definitions of key empirical concepts

4.1.1 ACT Adoption

I define ACT adoption as the choice of health facility care among individuals reporting recent fever in the treatment district (Rufiji) post-intervention. This is, perhaps, the most natural definition of adoption in the empirical context, for the following reasons. First, in the treatment district, ACT was only available through public health facilities; due to the CDC's continuous monitoring efforts, leakage of ACTs into the private sector was virtually nonexistent (Njau et al. 2008).

Second, this definition is appropriate because the key choice made by each acutely ill individual is whether and where to seek treatment–at a health facility vis-a-vis in the informal sector. While the patient may be able to exercise additional agency once at the health facility, the treatment she receives is, at that point, in large part decided by the health worker who administers her care. Thus, at the individual level, the primary decision of seeking out ACT as treatment or not–i.e., the adoption decision–is made through the choice of care.

Third, according to the CDC/IHI protocols regarding ACT distribution during the intervention, ACT was to be prescribed for every non-pregnant patient above the age of 2 months who presented at a public health facility with a fever or recent history of fever (Kachur et al. 2004). Thus, surveyed individuals in the treatment district who self-reported fever and visited a health facility should have, as per the distribution policy, received ACT. It is, however, possible that the policy may not have been enforced in all cases.

Two main possibilities are of concern. First, artemisinin stock-outs may have prevented some febrile patients from accessing combination therapy – in the event that artemisinin was not available, health workers were advised to prescribe SP monotherapy instead, which, as discussed earlier, is not as effective an antimalarial as ACT (Arrow et al. 2004). Second, in certain cases clinical judgement may have overruled the prescription policy. For example, if an individual presents with productive cough and shortness of breath along with fever, the health worker may conclude that pneumonia is more likely than malaria to be the underlying cause of

the symptoms, and may therefore prescribe antibiotics in lieu of ACT.

Julie Thwing et al. (2011) document ACT dispensing practices among health centers and dispensaries in the treatment district (Rufiji) before and during the intervention, from 2002 to 2005. The authors show that once the intervention began, 75 to 80 percent (depending on the quarter of survey) of patients with fever or recent history of fever were prescribed ACT, and of those prescribed, virtually all patients received the correct age-specific dosage (Thwing et al. 2011). Only 10 percent of febrile patients received SP monotherapy during the intervention, and most of this gap occurred during the second quarter of 2003, during which there was a brief artesunate stock-out.

In sum, defining adoption via healthcare choice is appropriate in this context because the choice of care is the patient's primary decision, and for individuals with fever in the treatment district, choosing care at a public or mission health facility resulted in access to ACT in the vast majority of cases.

4.1.2 Lagged adoption and outcomes

We begin by introducing some notation. Suppose 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 learning 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.[6] 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 up to m days at the outcomes of 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 $N^1_{jt} \subseteq N_{jt}$ denote the set of all individuals with fever in village j on date t who adopted the new therapy, and $N^0_{jt} \subseteq N_{jt}$ denote those who did not adopt, such that $N^1_{jt} \cup N^0_{jt} = N_{jt}$. Then $\bar{D}^1_{j,(t-m,t)}$ and $\bar{D}^0_{j,(t-m,t)}$, the average length of illness for adopters and non-adopters, respectively, from dates $t-m$ to t, are defined as follows:

$$\bar{D}^1_{j,(t-m,t)} = \frac{\sum_{a=1}^{m} \sum_{i \in N^1_{j,t-a}} D_{ij,t-a}}{\sum_{b=1}^{m} |N^1_{j,t-b}|} \tag{18}$$

$$\bar{D}^0_{j,(t-m,t)} = \frac{\sum_{a=1}^{m} \sum_{i \in N^0_{j,t-a}} D_{ij,t-a}}{\sum_{b=1}^{m} |N^0_{j,t-b}|}. \tag{19}$$

Equivalently, we define lagged average adoption as the following:

$$\bar{h}_{j,(t-m,t)} = \frac{\sum_{a=1}^{m} \sum_{i \in N_{j,t-a}} h_{ij,t-a}}{\sum_{b=1}^{m} |N_{j,t-b}|}. \tag{20}$$

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 level of village x day on which illness began. Thus it is possible that two individuals surveyed on the same day in the same village may have different lagged outcome and lagged adoption realizations if their illnesses began on different dates.

Second, I only average over observations for whom the length of illness is recorded, i.e.,

[6] The 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.

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 $(t-m, t)$, $\sum_{b=1}^{m} |N^k_{j,t-b}| = 0$, I replace the missing value of $\bar{D}^k_{j,(t-m,t)}$ 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.

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 4, 5, 7 and 8 week lag lengths, and the results are qualitatively similar.

4.1.3 Misdiagnosis of malaria

I construct a measure of the extent of malaria misdiagnosis in each acutely ill individual's reference group, by exploiting individual-level blood slide data on malaria positivity collected at the time of survey enumeration. Define $M_{ijt} \in \{0, 1\}$ as the result of the blood slide test for individual i in village j surveyed on date t, where $M_{ijt} = 1$ denotes that the individual was found positive for malaria, and $M_{ijt} = 0$ denotes negative. The average level of reference group misdiagnosis is then calculated as the following:

$$\bar{M}_{j,(t-m,t)} = \frac{\sum_{a=1}^{m} \sum_{i \in N_{j,t-a}} (1 - M_{ij,t-a})}{\sum_{b=1}^{m} |N_{j,t-b}|}. \tag{21}$$

Intuitively, misdiagnosis is defined as the proportion of acutely ill individuals in the reference group who were confirmed as negative for malaria via blood slide test. I use this definition in the majority of the analysis presented in the next section.

The main difficulty with defining misdiagnosis in this way is that the blood slide data were

collected at the time of survey enumeration, not when the individual's illness actually began. The time lag between the start of illness and the time of survey could generate discrepancies in *prior* and *measured* malarial status. Specifically, two cases are problematic. First, some individuals who had malaria at the start of their illness may have received effective enough treatment that their observed malarial status (at the time of survey) is negative. Second, some individuals who did not have malaria at the start of their illness may have acquired it sometime between the start of illness and date of survey.

Let us address these two cases in turn. First, I argue that the proportion of acutely ill individuals whose prior malarial status was positive and measured status was negative is not large. Of 1891 acutely ill individuals with non-missing blood slide data, 1429 (about 75 percent) were measured as negative for malaria. Of these individuals, 1075 (about 57 percent of the total) had potential access to some antimalarial, via self-medication or health facility usage; the rest either sought/received no treatment or received traditional (non-antimalarial) treatment. Of these, some had malaria at the start of their illnesses and some did not. Supposing that this proportion is equal to the baseline rate of malarial prevalence among acutely ill individuals in the sample (about 22 percent), we are left with approximately 237 individuals (about 13 percent of the total 1891).

For these individuals, the only way to transition from positive to negative malarial status in the time between the start of sickness and time of survey is by taking effective treatment. The proportion of individuals receiving effective treatment is a function of the (local) resistance to antimalarial therapies other than ACT. For example, in many areas of Tanzania, more than half of *P. falciparum* samples had acquired resistance to sulphadoxine pyrimethamine (SP) by 2004 (Bloland 2005). In the extreme case, one might argue that only ACT was maximally effective, meaning that only those individuals in the treatment district post-intervention would plausibly change from positive to negative malarial status in the above defined time span. Taking 50 percent as a more conservative estimate, we are left with about 119 individuals, which is approximately 6 percent of the total 1891.

The second scenario, in which individuals acquire malaria sometime between the start of

illness and the time of survey, is likely very small, given that the average time elapsed (i.e. the average number of days after the start of illness individuals were surveyed) is just over 6 days. A very small proportion of individuals are likely to have acquired malaria during this short a time span on average.

Given these two facts, it is likely that for the great majority of acutely ill individuals (who comprise the reference groups for learning I construct here), measured malaria status (at the time of survey) equals prior malaria status (at the start of illness).

In a set of robustness checks, covered in detail in section 6, I define two alternative measures of misdiagnosis that address the potential issues still outstanding from the discussion above. In particular, I first restrict the calculation of the lagged average misdiagnosis level to only those individuals in the reference group whose illnesses started within one week of survey. This restriction further shrinks the length of time between the prior malarial status and measured status that is the source of error. Second, for each individual in the sample, I predict malaria positivity using a nonparametric function of age, and use only the predicted component of malaria status to calculate the misdiagnosis rate in the reference group. As shown by the results described in section 6, changing the definition of misdiagnosis in these ways does not substantially quantitatively alter my main findings.

4.2 Misdiagnosis and the adoption rate

In this section, I develop an empirical test of the prediction that individuals whose reference groups for learning faced higher misdiagnosis rates should be less likely to adopt. 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 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 cat-

egories, I first estimate the following difference in differences specification separately for two types of acutely ill individuals: those whose learning reference groups faced above-median and below-median misdiagnosis rates:

$$h_{ijr} = \gamma T_j P_r + \alpha_r + \beta_j + \mathbf{X}'_{ijr}\delta + \epsilon_{ijr}. \tag{22}$$

Here i denotes individual, j denotes village (and β_j are village fixed effects), r denotes round of survey (and α_r are round fixed effects), and $\mathbf{X}_{ijr}$ 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 a quadratic term in age. We denote the dummy variable for health facility usage among acutely ill individuals as h_{ijr}, 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_r is a post-intervention dummy, which equals 1 in post-intervention rounds. The coefficient of interest is γ, the difference in differences estimate of the impact of ACT introduction on health facility usage.

I then interact the treatment x post-intervention term with the misdiagnosis rate faced by the individual's reference group (denoted $\bar{M}_{j,(t-m,t)}$), include the second-order interactions, and estimate a triple interaction specification in the pooled sample, measuring the differential adoption rate across reference-group misdiagnosis rates:

$$\begin{aligned} h_{ijtr} = \quad & \gamma_1 \bar{M}_{j,(t-m,t)} T_j P_r + \gamma_2 \bar{M}_{j,(t-m,t)} T_j + \gamma_3 \bar{M}_{j,(t-m,t)} P_r + \gamma_4 T_j P_r + \gamma_5 \bar{M}_{j,(t-m,t)} \\ & + \alpha_r + \beta_j + \mathbf{X}'_{ijr}\delta + \epsilon_{ijr}. \end{aligned} \tag{23}$$

The coefficient of interest is γ_1, measuring differential adoption across high- and low-misdiagnosis villages.

4.3 Learning effect estimation

4.3.1 Estimating equation

We begin with equation 7, which describes the cutoff rule governing the period-t adoption decision: $h_{it} = \mathbf{1} \quad q_t > \frac{\Delta P}{m(1-p)\Delta u_i}$. The decision rule specifies that an individual will adopt if and only if q_t is greater than an individual- and illness episode-specific cutoff. Expressing this cutoff as a function of individual, village, and time-varying observables ($\mathbf{X}'_{ijt}$) plus unobserved characteristics (ε_{ijt}), we can rewrite equation 7 as

$$h_{ijt} = \mathbf{1}\left(\varepsilon_{ijt} + \mathbf{X}'_{ijt}\beta < q_t\right). \tag{24}$$

q_t, the period-t prior on the therapy's effectiveness, is unobserved. From equation 10, we know that the current-period belief is an additive function of the previous period's belief (q_{t-1}) and the outcomes of previous-period adopters ($\bar{D}^1_{j,(t-m,t)}$, in the empirical notation introduced above). Since q_{t-1} is unobserved as well, we use average previous-period adoption ($\bar{h}_{j,(t-m,t)}$) as a proxy. Substituting these into equation 24 above, we obtain the following baseline estimating equation:

$$h_{ijt} = \mathbf{1} \quad \varepsilon_{ijt} < \gamma_1 \bar{D}^1_{j,(t-m,t)} + \gamma_2 \bar{h}_{j,(t-m,t)} - \mathbf{X}'_{ijt}\beta \quad . \tag{25}$$

The coefficient of interest is γ_1, which reflects the size of the learning effect – i.e., the extent to which information gleaned from the previous-period outcomes of adopters influences the current-period probability of adoption.

4.3.2 Contagion-based autocorrelation and selection into adoption

Of course, it is unlikely in reality that the only way in which past adopters' outcomes and current adoption probabilities are linked is through the learning process. We might expect that the learning effect estimate, γ_1 above, would be biased 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 these correlated effects. 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 flexibly control for local, time-varying unobserved factors by introducing the full set of village x week-of-survey fixed effects, denoted η_{jw}. These dummies stringently account for fluctuations in the local environment that simultaneously drive current-period adoption and past-period health outcomes. I am thus restricting attention to variation in lagged outcomes (as defined above) within village x week cells. Since the lagged adoption and outcomes variables are defined at the level of village x day the illness began, there exists substantial variation in these variables across *days and individuals* within fixed effect cells.

But even within fixed effect 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 my empirical strategy is to difference the past outcomes of adopters and non-adopters: $\Delta\bar{D}_{j,(t-m,t)} := \bar{D}^1_{j,(t-m,t)} - \bar{D}^0_{j,(t-m,t)}$. 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.[7]

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

[7]Thus 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 that the only intertemporal link between outcomes and adoption is through learning, it is nevertheless crucial to make in our empirical setting, in which common autocorrelated shocks may bias 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_r), as well as the main effects–treatment (T_j) and post-intervention (P_r), where r denotes round of survey. Representing equation 25 as a linear probability model, the resulting triple difference specification is:

$$h_{ijtrw} = \gamma_1 T_j P_r + \gamma_2 T_j + \gamma_3 P_r + \gamma_4 \Delta\bar{D}_{j,(t-m,t)} + \eta_{jw} + \gamma_2 \bar{h}_{j,(t-m,t)} + \mathbf{X}'_{ijt}\beta + \varepsilon_{ijtrw}. \quad (26)$$

The learning effect is captured by the coefficient γ_1. $\mathbf{X}$ includes the following variables: dummies for categories of educational attainment of the household head; a quadratic term in age; and the date on which the individual fell acutely ill.

4.4 Misdiagnosis and the size of the learning effect

To test the prediction that misdiagnosis decreases the rate of learning, I first estimate equation 26 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 (γ_1) is on the interaction of the learning effect estimate and the reference group misdiagnosis rate. The specification estimated in the pooled sample is:

$$h_{ijtrw} = \gamma_1 T_j P_r + \gamma_2 T_j + \gamma_3 P_r + \gamma_4 \Delta\bar{D}_{j,(t-m,t)}\bar{M}_{j,(t-m,t)}$$

$$
\begin{aligned}
&+ \ \gamma_5 T_j P_r + \gamma_6 T_j + \gamma_7 P_r + \gamma_8 \ \ \Delta \bar{D}_{j,(t-m,t)} \\
&+ \ \gamma_9 T_j P_r + \gamma_{10} T_j + \gamma_{11} P_r + \gamma_{12} \ \ \bar{M}_{j,(t-m,t)} \\
&+ \gamma_{13} \Delta \bar{D}_{j,(t-m,t)} \bar{M}_{j,(t-m,t)} + \gamma_{14} \Delta \bar{D}_{j,(t-m,t)} + \gamma_{15} \Delta \bar{D}_{j,(t-m,t)} \\
&+ \eta_{jw} + \gamma_2 \bar{h}_{j,(t-m,t)} + \mathbf{X}'_{ijt} \beta + \varepsilon_{ijtrw}. \qquad (27)
\end{aligned}
$$

4.4.1 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:

$$
\begin{aligned}
\bar{x}^1_{j,(t-m,t)} &= \frac{\sum_{a=1}^{m} \sum_{i \in N^1_{j,t-a}} x_{ij,t-a}}{\sum_{b=1}^{m} |N^1_{j,t-b}|} \qquad (28) \\
\bar{x}^0_{j,(t-m,t)} &= \frac{\sum_{a=1}^{m} \sum_{i \in N^0_{j,t-a}} x_{ij,t-a}}{\sum_{b=1}^{m} |N^0_{j,t-b}|} \qquad (29) \\
\Delta \bar{x}_{j,(t-m,t)} &= \bar{x}^1_{j,(t-m,t)} - \bar{x}^0_{j,(t-m,t)}. \qquad (30)
\end{aligned}
$$

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:

$$
h_{ijt} = \gamma_1 T_j P_t + \gamma_2 T_j + \gamma_3 P_t + \gamma_4 \ \ \Delta \bar{D}_{j,(t-m,t)}
$$

$$+ \quad \alpha_1 T_j P_t + \alpha_2 T_j + \alpha_3 P_t + \alpha_4 \quad \Delta\bar{x}_{j,(t-m,t)} + \eta_{jw} + \mathbf{X}'_{ijt}\delta + \quad _{ijt}. \tag{31}$$

5 Results

5.1 Misdiagnosis and ACT adoption

Table 2 reports estimates of ACT adoption in the whole sample, split by reference-group misdiagnosis classification, and in a triple difference specification which formally tests the prediction that for individuals whose reference groups experienced more misdiagnoses, the rate of adoption was lower. Columns 1-3 report the results of estimations of equation 22. In column 1, I estimate adoption in the pooled sample. We find that, overall, ACT introduction increased health facility usage by individuals with fever by more than 18 percent in the treatment district as compared to the control.

I then split the sample into individuals whose reference groups experienced greater than or less than the median misdiagnosis level, and estimate equation 22 separately in these two groups. The results are reported in column 2 and 3, respectively. In both columns, we find a positive (but imprecisely estimated) adoption rate, but the magnitude of the coefficient in column 3, for individuals whose reference groups experienced relatively fewer misdiagnoses, is nearly double the estimated coefficient in column 2.

In column 4, I test formally for a difference in the adoption rate across misdiagnosis levels by estimating the interaction specification in equation 23. As indicated by the results in columns 2 and 3, we find that ACT adoption was significantly greater when individuals' reference groups experienced less misdiagnoses. A one standard deviation increase in reference group misdiagnosis level (an increase of approximately 0.2) decreases adoption by about 13 percentage points.

5.2 Estimates of the learning effect

In Table 3, I report estimates of the learning effect. Column 1 reports estimates of the baseline specification, equation 26. 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 past adopters and non-adopters by 1 day *increases* the future adoption probability by just over 20 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. These columns report estimations of the augmented learning specification in equation 31.

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. Again, the results reported in column 3 show that the magnitude of the learning effect estimate and the precision with which it is estimated remains quite stable.

5.3 Misdiagnosis and the learning effect

Table 4 reports estimates of equation 26 separately for individuals whose reference groups experienced above- and below-median levels of misdiagnoses, and then reports estimates of equation 27 in the pooled sample. As the results reported columns 1 and 2 show, the learning effect is slightly smaller for individuals with reference groups having relatively high misdiagnosis levels as compared to low, but is not precisely estimated in the two samples. The pooled sample estimates of equation 27, reported in column 3, confirm that the learning effect is significantly different in high versus low misdiagnosis villages. Specifically, a one standard deviation increase in reference group misdiagnosis level diminishes the learning effect (i.e., makes the coefficient more positive) by nearly 12 percentage points.

6 Robustness and Other Checks

6.1 Selection into acute illness

I first check that trends in self-reporting of acute illness (equivalent to inclusion in the sample) are not differential by reference group misdiagnosis level. If trends in sample selection were different across these two groups, we may worry that differential ACT adoption and learning may be due to differences in selection into acute illness rather than differences in the extent of misdiagnosis.

To check whether this is indeed the case, I first estimate the following specification separately in the high and low misdiagnosis samples:

$$s_{ijr} = \gamma T_j P_r + \alpha_r + \beta_j + \mathbf{X}'_{ijr}\delta + \ _{ijr}. \tag{32}$$

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 and 2 of Table 5. The treatment x post-intervention interaction term is small and similar in magnitude across individuals whose reference groups experienced relatively high and low misdiagnosis, and is insignificantly different from 0 in both groups.

Finally, to estimate the differential trends in selection into acute illness across reference group misdiagnosis levels, I interact the treatment district x post-intervention term with the reference group's misdiagnosis level and include the second-order interactions:

$$\begin{aligned} s_{ijtr} = \quad & \gamma_1 \bar{M}_{j,(t-m,t)} T_j P_r + \gamma_2 \bar{M}_{j,(t-m,t)} T_j + \gamma_3 \bar{M}_{j,(t-m,t)} P_r + \gamma_4 T_j P_r + \gamma_5 \bar{M}_{j,(t-m,t)} \\ & + \alpha_r + \beta_j + \mathbf{X}'_{ijr}\delta + \ _{ijr}. \end{aligned} \tag{33}$$

The results are reported in column 3. Based on the triple difference estimates, we find no evidence of differential selection into sickness across reference group misdiagnosis levels.

6.2 Plausible randomness of survey enumeration

My empirical strategy relies heavily 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 integer ordering of survey enumeration within villages (based on the date of survey) on these variables. That is, for each village we order individuals by date of survey (when the data set is collapsed at the date-of-survey x village level), and then assign numbers such that the earliest individual surveyed would be assigned 1, the next earliest 2, and so forth. We denote the integer order of survey by o_{ijt}. The specification we then estimate is the following:

$$o_{ijt} = \delta' \mathbf{X}_{ijt} + \eta_{jw} + \ _{ijt}. \tag{34}$$

The results of this estimation are 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 as well as the individual's blood slide result, 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 villages cells, there is no significant relationship between the ordering of survey enumeration and self-reported acute illness, and the same holds for 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 and malaria positivity variable used in the previous specification. Again the results show no significant association between date of survey and any of these variables.

Columns 3 and 4 replicate columns 1 and 2, respectively, but in these columns I add additional variables related to the individual's reference group. Specifically, I add reference group misdiagnosis level, lagged average health facility usage in the reference group, and the difference in lagged sickness length for adopters and non-adopters in the reference group. 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.

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

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 reference group misdiagnosis levels. To check whether this is the case, I construct a dummy variable c_{ijr} for individual i in village j surveyed in round r 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:

$$c_{ijr} = \gamma T_j P_r + \alpha_r + \beta_j + \mathbf{X}'_{ijr}\delta + \ _{ijr}. \tag{35}$$

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. I first estimate this equation separately for individuals whose reference groups experienced greater than or less than the median misdiagnosis level. The coefficient estimates are reported in columns 1 and 2 of Table 7. The coefficients on treatment x post-intervention are not significantly different from 0.

In column 3, I report estimates of a triple difference specification, through which it is possible to test whether there exist differential trends in non-response in the length of illness across

reference group misdiagnosis levels:

$$\begin{aligned} c_{ijtr} = \quad & \gamma_1 \bar{M}_{j,(t-m,t)} T_j P_r + \gamma_2 \bar{M}_{j,(t-m,t)} T_j + \gamma_3 \bar{M}_{j,(t-m,t)} P_r + \gamma_4 T_j P_r + \gamma_5 \bar{M}_{j,(t-m,t)} \\ & + \alpha_r + \beta_j + \mathbf{X}'_{ijr}\delta + \ _{ijr}. \end{aligned} \tag{36}$$

The estimate of γ_1 reported in column 3 shows that this difference is not significantly different from 0. Thus, we find no evidence that trends in the censoring of average length of illness is differential across treatment and comparison districts, and across individuals whose reference groups experienced differing misdiagnosis levels.

6.4 Choice of lag length

The choice of a lag length of six weeks for defining reference groups 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. In this section, I present results in which the main analyses, related to the effects of misdiagnosis on adoption and learning, are rerun using lag lengths of 4, 5, 7 and 8 weeks. The results are reported in Tables 8 and 9. Table 8 shows the overall adoption rate estimate, as well as the differential adoption across reference group misdiagnosis levels, for the four aforementioned alternative lag lengths. Across all lag lengths shown, we see that the adoption results are quite consistent in magnitude, as is the magnitude of the differential adoption rate across reference group misdiagnosis levels.

In Table 9, I rerun the learning effects results, and find the same pattern: across all lag lengths (perhaps with the exception of 4 weeks, for which the learning effect is imprecisely estimated), the magnitude of the learning effect is very similar, and the extent to which reference group misdiagnosis levels affect the learning effect is also consistent. Overall, these results suggest that my main findings are robust to a wide range of choices of reference group lag length.

6.5 Definition of reference group misdiagnosis

As mentioned in section 4, there may be some discrepancy between malarial status at the start of illness (the negative of which is the most natural definition of misdiagnosis) and measured malarial status at the time of survey. While the discussion in section 4 makes clear that the discrepancy should, in fact, occur in only a small percentage of cases, it is important to provide evidence that the results are robust to changing the definition of reference group misdiagnosis.

Along these lines, I use two variants of the previously proposed definition. First, I compute misdiagnosis level using only individuals whose illnesses began within a week preceding survey. Shrinking the time gap between actual misdiagnosis and measured misdiagnosis should further reduce the discrepancy between the two. Second, I predict malaria positivity by estimating malaria as a nonparametric function of age, and then using only the predicted malaria probability (utilizing the projected values from the nonparametric estimation) instead of actual malaria positivity when computing reference group misdiagnosis level.

I then rerun the main analyses related to adoption and learning using these alternate definitions of misdiagnosis. The results are reported in Tables 10 and 11. Table 10 shows that both alternate definitions yield adoption results similar to the main findings, and Table 11 similarly shows that the learning effect estimates and learning effects by misdiagnosis are again qualitatively much like the main findings. Overall, the results from the alternate definitions of misdiagnosis show that my main results are indeed robust to these alternate definitions.

7 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) individuals whose reference groups experienced idiosyncratically more frequent misdiagnoses have lower adoption rates, 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 Mwakasungula, 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)	
Number of individuals	14107		6898		7209	
Number of individuals reporting fever in 2 weeks preceding survey	1891		986		905	
	Mean	SD	Mean	SD	Mean	SD
Among acutely ill individuals (pre and post):						
Age (years)	23.453	22.074	24.318	23.418	22.888	20.792
Educational attainment of household head (years)	4.279	3.367	3.224	3.405	5.353	2.965
Pre-intervention (2001):						
Proportion reporting fever in 2 weeks preceding survey	0.172	0.377	0.190	0.393	0.153	0.360
Among individuals reporting fever:						
Proportion who sought care at health facility	0.248	0.432	0.263	0.441	0.230	0.421
Among individuals who sought health facility care:						
Length of illness (days)	4.731	2.667	4.992	2.912	4.333	2.203
Number of additional symptoms	1.500	1.115	1.319	1.059	1.750	1.153
Among individuals who did not seek health facility care:						
Length of illness (days)	3.299	2.348	3.325	2.468	3.267	2.199
Number of additional symptoms	1.420	1.054	1.278	1.017	1.588	1.074

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.

Table 2

Trends in ACT Adoption by Reference Group's Level of Misdiagnosis

Dependent var: 1 if individual sought care at a government or NGO health facility, 0 otherwise				
		ACT Adoption		*Adoption by misdiagnosis level*
	Whole sample	1(Reference group misdiagnosis > median)	1(Reference group misdiagnosis < median)	Whole sample
Treatment x post	0.183**	0.106	0.205	0.694**
	(0.0732)	(0.0913)	(0.133)	(0.286)
Reference group misdiagnosis level x				
Treatment x post				-0.635*
				(0.325)
Treatment				0.449*
				(0.266)
Post				0.537**
				(0.229)
Reference group misdiagnosis level	-0.0300	0.0615	-0.212*	-0.416*
	(0.0808)	(0.256)	(0.107)	(0.222)
Fixed effects		*Village & survey round*		
Number of observations	1,749	864	885	1,749

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, and a quadratic term in age.

Table 3

Learning Effect Estimates

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
Difference in lagged sickness length in days (health facility users - non-health facility users) x			
Treatment x post	-0.219***	-0.254***	-0.220**
	(0.0765)	(0.0866)	(0.0915)
Treatment	0.0948**	0.129**	0.103*
	(0.0465)	(0.0548)	(0.0584)
Post	0.0885	0.120*	0.0914
	(0.0560)	(0.0676)	(0.0738)
Difference in lagged sickness length in days (health facility users - non-health facility users)	-0.0252	-0.0499	-0.0240
	(0.0349)	(0.0428)	(0.0468)
Reference group misdiagnosis level	0.0377	0.00228	0.0520
	(0.140)	(0.147)	(0.164)
Fixed effects		*Village x year x week*	
Number of observations	1,691	1,593	1,561

Notes: *** p<0.01, ** p<0.05, * p<0.1; all specifications are estimated using OLS; reference group for learning is all sick individuals in village within 6 weeks of date illness began; proportion of acutely ill individuals in reference group visiting health facility is included as a control; robust standard errors are reported in parentheses below coefficient estimates, and allow for correlation in the error term within villages; additional controls are dummies for categories of educational attainment of the household head and a quadratic term 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.

Table 4

Learning Effect Estimates by Reference Group's Level of Misdiagnosis

Dependent var: 1 if individual sought care at a government or NGO health facility, 0 otherwise			
	Reference group misdiagnosis > median	Reference group misdiagnosis < median	Whole sample
Learning effect x Reference group misdiagnosis level			0.574***
			(0.205)
Learning effect estimate	-0.115	-0.146	-0.676***
	(0.0929)	(0.117)	(0.208)
Reference group misdiagnosis level	1.478***	-0.399	0.360
	(0.468)	(0.239)	(0.558)
Fixed effects		*Village x year x week*	
Number of observations	834	857	1,691

Notes: *** p<0.01, ** p<0.05, * p<0.1; all specifications are estimated using OLS; 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) x treatment x post; main effect of this difference and its interactions with treatment and post-intervention dummies are included in all specifications; proportion visiting health facility in reference group is also included in all specifications; column 3 reports results of an interaction specification, in which the learning effect is interacted with a dummy for individuals for whom reference group misdiagnosis level was below median; robust standard errors are reported in parentheses below coefficient estimates, and allow for correlation in the error term within villages; additional controls are dummies for categories of educational attainment of the household head and a quadratic term in age.

Table 5

Sample Selection

Dependent var: 1 if individual sought care at a government or NGO health facility, 0 otherwise			
	1(Reference group misdiagnosis > median)	1(Reference group misdiagnosis < median)	Whole sample
Treatment x post	-0.0228	-0.0492	-0.0919
	(0.0144)	(0.0306)	(0.120)
Reference group misdiagnosis level x			
Treatment x post			0.0558
			(0.146)
Treatment			-0.0426
			(0.132)
Post			0.00748
			(0.116)
Reference group misdiagnosis level	0.0833	-0.00310	-0.00598
	(0.0819)	(0.0434)	(0.110)
Fixed effects		*Village & survey round*	
Number of observations	8,503	8,436	16,939

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, and a quadratic term in age.

Table 6

Timing of Survey Enumeration

Dependent var: Integer order of survey				
	Whole sample	Acutely ill individuals	Whole sample	Acutely ill individuals
1(Blood slide test for malaria was positive)	0.0221	0.000966	0.0280	0.000602
	(0.0255)	(0.0833)	(0.0238)	(0.0953)
Educational attainment of household head:				
Less than primary (< 7 years)	0.0691	0.132	0.0594	0.114
	(0.0646)	(0.137)	(0.0643)	(0.137)
Completed primary school (= 7 years)	0.0294	0.0998	0.0385	0.117
	(0.0617)	(0.122)	(0.0621)	(0.126)
More than primary school (> 7 years)	-0.0318	-0.0174	-0.0335	0.0244
	(0.122)	(0.201)	(0.126)	(0.234)
Age in years	0.000589	-0.00257	0.000685	-0.00374
	(0.000871)	(0.00555)	(0.000890)	(0.00593)
Age squared	-2.66e-07	7.00e-05	-1.12e-06	8.70e-05
	(1.18e-05)	(6.99e-05)	(1.16e-05)	(7.55e-05)
Reported being acutely ill in 2 weeks preceding survey	-0.0340		-0.0390*	
	(0.0205)		(0.0220)	
Sought treatment at health facility		-0.0121		-0.0172
		(0.0916)		(0.0882)
Reference group misdiagnosis level			-0.486	0.936
			(1.112)	(1.335)
Lagged average health facility usage in reference group			-2.324	0.489
			(2.604)	(5.998)
Difference in lagged sickness length in days (health facility users - non-health facility users)			0.0456	0.125
			(0.0890)	(0.124)
Fixed effects		*Village x year x week*		
Number of observations	13,589	1,506	13,094	1,433

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.

Table 7

Trends in Truncation of Length of Acute Illness

Dependent var: 1 if individual sought care at a government or NGO health facility, 0 otherwise			
	1(Reference group misdiagnosis > median)	1(Reference group misdiagnosis < median)	Whole sample
Treatment x post	-0.103	-0.0484	0.0962
	(0.0790)	(0.0839)	(0.160)
Reference group misdiagnosis level x			
Treatment x post			-0.168
			(0.185)
Treatment			0.502***
			(0.139)
Post			0.118
			(0.106)
Reference group misdiagnosis level	-0.265	0.107	-0.319***
	(0.211)	(0.0841)	(0.113)
Fixed effects		*Village & survey round*	
Number of observations	864	885	1,749

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, and a quadratic term in age.

Table 8

Robustness of ACT Adoption Results to Lag Length Used to Define Reference Group

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

Lag length for reference group:	Four weeks		Five weeks		Seven weeks		Eight weeks	
	Overall adoption	*Adoption by misdiagnosis level*	*Overall adoption*	*Adoption by misdiagnosis level*	*Overall adoption*	*Adoption by misdiagnosis level*	*Overall adoption*	*Adoption by misdiagnosis level*
Treatment x post	0.179**	0.455*	0.178**	0.581**	0.180**	0.765***	0.182**	0.725**
	(0.0740)	(0.256)	(0.0736)	(0.263)	(0.0746)	(0.274)	(0.0743)	(0.275)
Reference group misdiagnosis level x								
Treatment x post		-0.339		-0.492		-0.730**		-0.672**
		(0.279)		(0.298)		(0.314)		(0.318)
Treatment		0.259		0.407		0.447		0.466*
		(0.234)		(0.251)		(0.271)		(0.269)
Post		0.351*		0.487**		0.615***		0.610***
		(0.197)		(0.208)		(0.213)		(0.221)
Reference group misdiagnosis level	-0.0140	-0.269	-0.00368	-0.374*	-0.0143	-0.417*	-0.0266	-0.449*
	(0.0739)	(0.202)	(0.0786)	(0.213)	(0.0860)	(0.228)	(0.0820)	(0.226)
Fixed effects	*Village & survey round*		*Village & survey round*		*Village & survey round*		*Village & survey round*	
Number of observations	1,748	1,748	1,749	1,749	1,749	1,749	1,749	1,749

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, and a quadratic term in age.

Table 9

Robustness of Learning Results to Lag Length Used to Define Reference Group

Dependent var: 1 if individual sought care at a government or NGO health facility, 0 otherwise								
Lag length for reference group:	Four weeks		Five weeks		Seven weeks		Eight weeks	
	Overall learning effect	*Learning by misdiagnosis level*	*Overall learning effect*	*Learning by misdiagnosis level*	*Overall learning effect*	*Learning by misdiagnosis level*	*Overall learning effect*	*Learning by misdiagnosis level*
Learning effect x Reference group misdiagnosis level		0.490**		0.480**		0.573***		0.429*
		(0.189)		(0.199)		(0.212)		(0.224)
Learning effect estimate	-0.0927	-0.503**	-0.160**	-0.543**	-0.223***	-0.679***	-0.219***	-0.552**
	(0.0751)	(0.194)	(0.0711)	(0.204)	(0.0768)	(0.216)	(0.0782)	(0.224)
Reference group misdiagnosis level	-0.0155	0.415	0.0215	0.292	0.105	0.400	0.0234	0.408
	(0.155)	(0.446)	(0.139)	(0.579)	(0.171)	(0.544)	(0.161)	(0.601)
Fixed effects	*Village x year x week*		*Village x year x week*		*Village x year x week*		*Village x year x week*	
Number of observations	1,688	1,688	1,687	1,687	1,685	1,685	1,685	1,685

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 dummies for categories of educational attainment of the household head and a quadratic term in age.

Table 10

Robustness of ACT Adoption Results to Definition of Misdiagnosis Level in Reference Group

Dependent var: 1 if individual sought care at a government or NGO health facility, 0 otherwise				
Reference group misdiagnosis definition:	Only recent illnesses used in reference group misdiagnosis		Predicted malariosity based on non-parametric f. of age used	
	Overall adoption	*Adoption by misdiagnosis level*	*Overall adoption*	*Adoption by misdiagnosis level*
Treatment x post	0.108	0.531**	0.184***	0.439
	(0.0681)	(0.236)	(0.0676)	(0.775)
Reference group misdiagnosis level x				
Treatment x post		-0.567**		-0.307
		(0.280)		(0.980)
Treatment		0.278		0.713
		(0.219)		(0.725)
Post		0.413*		0.536
		(0.243)		(0.827)
Reference group misdiagnosis level	0.136**	-0.111	-0.319	-0.929
	(0.0520)	(0.204)	(0.265)	(0.627)
Fixed effects	*Village & survey round*		*Village & survey round*	
Number of observations	1,680	1,680	1,749	1,749

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, and a quadratic term in age.

Table 11

Robustness of Learning Results to Definition of Misdiagnosis Level in Reference Group

Dependent var: 1 if individual sought care at a government or NGO health facility, 0 otherwise				
Reference group misdiagnosis definition:	Only recent illnesses used in reference group misdiagnosis		Predicted malariosity based on non-parametric f. of age used	
	Overall learning effect	*Learning by misdiagnosis level*	*Overall learning effect*	*Learning by misdiagnosis level*
Learning effect x Reference group misdiagnosis level		0.534***		0.457
		(0.175)		(0.718)
Learning effect estimate	-0.155*	-0.634***	-0.206***	-0.560
	(0.0859)	(0.136)	(0.0770)	(0.600)
Reference group misdiagnosis level	0.0662	-0.185	-0.649	0.0402
	(0.119)	(0.308)	(0.581)	(0.778)
Fixed effects	*Village x year x week*		*Village x year x week*	
Number of observations	1,627	1,627	1,691	1,691

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 dummies for categories of educational attainment of the household head and a quadratic term in age.