

Malaria Control and Infant Mortality in Africa

Denis COGNEAU and Pauline ROSSI*

August 2016

Abstract

Have malaria control efforts contributed to the reduction in infant mortality in Sub-Saharan Africa over the past 15 years? Using large household surveys collected in 19 countries between 2000 and 2015, we estimate the correlation between the distribution of bednets and the progress in child survival. We find that the large increase in bednets ownership observed between 2000 and 2015 is associated with a decrease in infant mortality by 1.3pp, which amounts to one third of the total decrease in infant mortality over the period. We further discuss to which extent this correlation might be interpreted as a causal impact.

Keywords: Child mortality, Malaria, Africa, Foreign aid.

JEL Codes: I1, J1, O1, F35.

*Denis Cogneau (corresponding author): Paris School of Economics and IRD, 48 boulevard Jourdan 75014 Paris, denis.cogneau@psemail.eu. Pauline Rossi: Paris School of Economics and CREST, 48 boulevard Jourdan 75014 Paris, pauline.rossi@ensae.fr.

1 Introduction

One of the big success stories of the global fight against poverty is the substantial reduction in child mortality experienced worldwide in recent years. The under 5 mortality rate was divided by 2 between 1990 and 2015, reaching 4.3% today, and there was an acceleration of the progress in the last 10 years. In this global endeavor, Sub-Saharan Africa has played and will play a crucial part, because it is the world region with the largest room for improvement. It is the region where the decrease was strongest in absolute terms, from 18% in 1990 down to 8.3% in 2015. There has been substantial progress, but the mortality rate is still twice as large as the world average. As a result, half of child deaths in 2015 occurred in Africa ([UN Inter-agency Group for Child Mortality Estimation, 2015](#)). In the future, this share is predicted to increase because the population growth is much larger in Africa than anywhere else in the world. Indeed, an estimated 40% of births will take place in Africa in 2050. Child mortality will thus increasingly become an African issue ([UNICEF, 2014](#)). Together with immunization campaigns, one large-scale health intervention is usually cited as a key success factor: the fight against malaria, coordinated by the public-private partnership Roll Back Malaria ([United Nations, 2015](#)).

Malaria has been eradicated from various parts of the world during the twentieth century. Today, the burden weighs mostly on Sub-Saharan Africa, where we find the largest prevalence rates and the most dangerous form of malaria. At the beginning of the twenty-first century, there was a series of initiatives launched by the international community to start the fight against malaria in Africa. Very large-scale interventions have been implemented, among which the most emblematic is the distribution of Insecticide-Treated Nets (henceforth ITNs). 900 millions of bednets have been distributed since the early 2000s. Nowadays, an estimated 2/3 of children sleep under an ITN against virtually none before the distribution started. Another type of intervention is improving access to curative treatments called

Artemisinin-based Combination Therapy. The scope of this intervention is more modest with only an estimated 16% of children being treated when they are sick in 2015. The last type of intervention is to have public agents spray the inside of dwellings with an insecticide (Indoor Residual Spraying). This type of intervention historically eradicated malaria in many places. Nowadays, it is less promoted and covers less than 5% of the population at risk in Africa. Since the start of Roll Back Malaria, the evolution of the disease in terms of prevalence and mortality has been closely monitored by the WHO. According to their estimates, the prevalence among children decreased from 33% in 2000 down to 16% in 2015, and the number of deaths caused by malaria among children under 5 years old decreased from 700K per year down to 300K (WHO, 2015).¹ The Millenium Development Goals Report 2015 concludes that the "substantial expansion of anti-malaria efforts have helped [...] improve child survival rates" (United Nations, 2015).

The goal of this paper is to challenge these statistics and assertions. Have malaria control efforts really contributed to the reduction in infant mortality in Sub-Saharan Africa over the past 15 years? Or have they simply coincided with favorable time trends? If they did make a difference, by how much and for whom? We face two main empirical challenges to answer these questions. First, measuring properly the progress in malaria control in countries with weak public health systems and no vital statistics; second, identifying a causal relationship when the intervention was not randomly allocated.

We believe that it is worth trying to overcome these issues because the question is of primary relevance for policy-makers. The Roll-Back Malaria initiative is a very large-scale, costly program. 2.5 billion dollars were spent in 2015, of which 1.5 went to Sub-Saharan Africa. The cumulated costs since 2000 exceed 20 billion dollars. As a point of comparison, approximately 5.5 billion dollars are spent each year by the international community on

1. We will later discuss these estimates in more details and highlight their limitations.

food aid.² Anti-malaria efforts are mostly funded by international organizations, such as the Global Fund to fight AIDS, tuberculosis and malaria, as well as development aid agencies of developed countries, the main contributors being the US and the UK. In many African countries, the program is in fact entirely funded by international aid (WHO, 2015). This paper contributes to the lively debate on the efficiency of aid by analyzing if a prominent global health intervention yielded some returns, how much and for whom. The answer is not obvious ex-ante. In spite of the enthusiasm of international institutions, the efficiency of the program has been questioned. In particular, it has been argued that bednets were not used properly.³

Using the Demographic and Health Surveys collected in 19 African countries between 2000 and 2015, we find that infant mortality did decrease more where more bednets were distributed. We estimate that the surge in ITN ownership during the period is associated with a reduction in infant mortality by 1.3pp, accounting for one third of the total decrease in infant mortality. The association is stronger for more disadvantaged households. In terms of magnitude, the correlation is large, and much larger than the impact of ITNs on child mortality estimated in the medical literature. Indeed, a Cochrane review of randomized control trials (henceforth RCTs) concludes that, in the African setting, ITNs reduce child mortality by 0.4pp to 0.8pp (Lengeler, 2004). We argue that the discrepancy can be explained by two reasons: first, medical RCTs lack external validity; second, the correlation captures an endogenous placement of ITNs and a bunching of health interventions, in addition to the causal impact of malaria control on mortality. Overall, we make a fairly positive assessment of the distribution of ITNs.

2. Data from the World Food Program <http://www.wfp.org/funding/year/2014>

3. See for example an article published on January, 25 2015 by the New York Times entitled "Meant to Keep Malaria Out, Mosquito Nets Are Used to Haul Fish In".

2 Literature

This paper relates both to the medical literature and the economics literature. On the one hand, the medical literature is primarily interested in assessing the burden of malaria on child deaths. This is a challenging task since there is no national vital registration systems categorizing all deaths by cause in African countries. The second-best source of data are verbal autopsy reports, in which the caregiver (sometimes a health worker, often a family member) is asked about the cause of the death. These reports provide information for some locations at some points in time. Next, researchers use statistical methods to extrapolate the results to larger units, like the continent. In this way, [Black et al. \(2014\)](#) estimate that 16% of under-5 deaths in Africa were caused by malaria in 2008.

Another strand of the medical literature assesses the causal impact of curative or preventive treatments against malaria on child mortality using RCTs. Researchers further perform some meta-analysis to come up with a global estimate of the treatment effect. [Lengeler \(2004\)](#) review the results of 5 RCTs estimating the impact of ITNs on under-5 mortality in Africa. Estimates range from -0.4pp to -0.8pp. Several points are important to note. First, 3 out of 5 RCTs were conducted in Kenya and the Gambia, where malaria is not endemic. If bednets save more lives where more children initially die from malaria, these coefficients underestimate the impact of distributing ITNs in endemic areas. Second, the under-5 mortality rates are extremely low in the 5 settings considered: between 1.3% and 5.2% in the control groups against 15% on average in Sub-Saharan Africa when the RCTs took place (in the 90s). Again, if the treatment effect is higher when the initial mortality rates are higher, RCTs estimates are biased downwards.

Overall, from the medical literature, we get fragmented, short-run and local information. Although internally valid, estimates derived from RCTs probably give us a lower bound of the impact of large-scale distribution campaigns on mortality.

As far as economists are concerned, they have been mostly interested in identifying the causal impact of malaria on socio-economic outcomes, like income or education. This strand of the economics literature fits into a global debate about the relationship between disease and development. Three papers set the golden rule for identification: [Bleakley \(2010\)](#) in the US, Brazil, Colombia and Mexico; [Cutler, Fung, Kremer, Singhal, et Vogl \(2010\)](#) in India; [Lucas \(2010\)](#) in Paraguay and Sri Lanka. They all use eradication campaigns as natural experiments, exploiting the variation in initial prevalence of malaria. The idea is to perform a difference-in-difference before and after the campaign, comparing areas with high and low initial prevalence. The assumption is that the eradication campaign caused a bigger change in areas more affected by malaria. The general message of these papers is that exposure to malaria in childhood deteriorates human capital. Note that the three papers abstract from child mortality, which might in principle raise an issue of selection. However, they study countries in America and South Asia, where malaria causes milder symptoms, rarely death, so the scope for such a concern is limited. Following these seminal contributions, several papers have exploited the same strategy to examine other outcomes and/or other contexts: [Venkataramani \(2010\)](#) on cognitive skills in Mexico, [Kuecken, Thuilliez, et Valfort \(2015\)](#) on education in Africa, and [Apouey, Picone, Wilde, Coleman, et Kibler \(2016\)](#) on anemia in Africa.

One study by [Kudamatsu, Persson, et Stromberg \(2012\)](#) is close to our paper in the sense that authors are interested in the relationship between malaria and infant mortality in Africa. But their perspective is quite different from ours. Their goal is to identify the causal impact of weather shocks on infant mortality. The motivation is to better understand the health consequences of climate change. For them, malaria is a channel through which unusually hot and humid weather conditions may cause extra deaths. They find that non-malarious weather conditions in utero reduce infant mortality by 0.8pp.

In this paper, we examine the relationship between malaria control efforts and infant

mortality. Contrary to the medical literature, we are able to make strong claims in terms of external validity because we exploit nationally representative surveys in 19 countries, covering more than two thirds of the population in Sub-Saharan Africa. Another advantage of our database is the large number of observations, which gives us the high statistical power needed to examine relatively rare events like child deaths.

However, we face a challenge in terms of internal validity because it is not possible to use the same identification strategy as the rest of the literature. Indeed, we cannot exploit the variation in initial prevalence of malaria for two reasons. First, the assumption of parallel trends does not hold (cf. Figure A.1 in Appendix). Second, malaria has not been eradicated yet, so the initial prevalence is not mechanically related to the intensity of the change. Instead, we need a measure of the treatment, i.e. observing the distribution of ITNs, which is not easy to construct as explained below.

3 Data

Ideally, we would like to have time series to monitor the evolution of the outcomes, infant and child mortality rates, and the evolution of malaria control efforts, both in terms of outputs like the malaria prevalence among children, and in terms of inputs like the coverage of interventions. We would like to observe these quantities at the finest geographical level and for the period of interest between 2000 and 2015, as well as the decade before to look at pre-trends. If these time series existed, estimating the correlation between malaria control efforts and the change in infant mortality would be straightforward.

The Demographic and Health Surveys (DHS) conducted in African countries between 2000 and 2015 contain part of the information needed.⁴ In particular, we can reconstruct

4. DHS provide stratified samples of mothers aged 15 to 49. Individual survey weights ensure that the survey sample is representative of all mothers at the country level. Nonetheless, sample size of surveys is not proportional to population size. To obtain a representative sample of the countries studied, we will reweigh the whole sample.

the evolution of mortality from the retrospective information on birth histories of mothers, including date of birth, survival status and, if applicable, date of death. However, the information on malaria control is limited. We have no data on malaria prevalence, and no retrospective information on bednets. We know if the household owns and uses a bednet, and what type of bednet, at the time of the survey, but not since when. So we cannot infer which child has been protected by a bednet since he was born.⁵

Our strategy to reconstruct the evolution of the intervention coverage is to exploit the spatial dimension. As explained in the next section, we need at least two waves of DHS in the same country since the early 2000s, when the distribution of bednets started. We also need the GPS coordinates of the geographical cluster where households are located. In the end, there are 19 countries with geocoded data in at least 2 waves of DHS over the relevant period. The corresponding surveys are listed in Table A.1 in Appendix.

An alternative measure of malaria control efforts is provided by the Malaria Atlas ([Bhatt et al., 2015](#)). This is a map constructed by epidemiologists showing malaria prevalence and ITN use at a very fine geographical level for each year between 2000 and 2015. It is possible to merge the Malaria Atlas with the DHS to reconstruct the environment in which each child was born. We consider this alternative measure in a robustness test but not in our main specification because their statistical model relies on very strong assumptions. **[MORE DETAILS ON MALARIA ATLAS]**

5. To be more specific, we know (i) if the household owns a bednet, and if it does, (ii) if children under 5 slept under a bednet the night before and whether these bednets were treated or untreated. In our analysis, we rely on the ownership variable because we want to assess the impact on the whole population, not only on bednet-users, so we are interested in an intent-to-treat estimator. Moreover, bednet ownership seems more reliable: it can be observed by the enumerator and it is less prone to recall bias.

4 Measuring the correlation of interest

4.1 Preliminary evidence

Before describing our measure of malaria control efforts, we provide some descriptive evidence of the correlation between bednets and infant mortality.

4.1.1 Cross-sectional analysis

We start by estimating the correlation between ITN use and infant mortality in a cross-section of households using the latest wave of DHS in each country. In the first column of Table 1, we perform an analysis between clusters, regressing the infant mortality rate in a given geographical cluster on the proportion of households using an ITN in the same cluster. We find a positive relationship: children die more in areas with a high coverage of ITN. This is because bednets are not needed in malaria-free areas. To assess their impact on child survival, it is essential to control for the riskiness of the environment.

This is what we do in the second column of Table 1 by performing an analysis within cluster. We compare the mortality rates of children born the same year in the same cluster – so they are facing the same environment – some sleeping under an ITN while others do not. The correlation is now negative: those sleeping under an ITN are 0.9pp less likely to die. However, households using a bednet are probably different from the others, and these specific characteristics might explain the survival advantage. We control for mother’s age, education and wealth but there remains some unobserved heterogeneity such as mother’s care for the children that we cannot account for using a cross-section.

4.1.2 Panel analysis: country-level

We therefore turn to panel analysis. At the country level, we have data on mortality rates and ITN access at two points in time.⁶ Figure 1 plots, for each country in our sample, the variation in ITN access on the x axis and the variation in infant mortality on the y axis. The relationship is negative: countries with a strong take-up in ITN also have a fastest decline in mortality. When we fit a linear regression on the raw data, we get a coefficient equal to -0.036 significant at the 5% level.

To which extent is our sample of 19 countries representative of the whole continent? We can perform the same analysis on a larger sample of countries if we rely on the measures of ITN use provided by the Malaria Atlas, and on the mortality rates provided by the World Development Indicators for 39 African countries in 2000 and 2014. The estimate of the linear regression is very close to what we have found: -0.038 . This provides some level of reassurance that our sample is not too specific.

Another way to test if the mortality decreased more in countries where more bednets were distributed is to look at a difference-in-difference. Figure 2 shows the evolution of infant mortality rates by intensity of anti-malaria campaigns. In the "low intensity" group of countries, the proportion of households owning an ITN increased by less than the median (35pp) between the two waves of DHS. Over this period, the mortality rate decreased by 1.6 percentage points, from 7.7% down to 6.1%. In the "high intensity" group, ITN access increased by more than the median. The mortality rate was reduced by 3 percentage points, from 8.3% down to 5.3%.

One interpretation is the following. Before the Roll-Back Malaria initiative started, countries with high malaria prevalence lagged behind. Anti-malaria campaigns targeted these countries, enabling them to catch up with the others. The decrease in mortality was twice as

6. For countries with 3 waves of DHS during the period 2000-2015, we consider the first and third waves to have the longest time span.

large in the "high intensity" group as in the "low intensity" one. The difference-in-difference coefficient is equal to -1.4pp, significant at the 10% level. Again, if we perform the same exercise on a bigger sample of countries using statistics provided by the Malaria Atlas ⁷ and World Development Indicators, we find exactly the same coefficient: -1.4pp.

The main shortcoming of country-level analysis is that there are many variables driving both the evolution of mortality rates and the distribution of ITNs. They may either underestimate or overestimate the impact of bednets on mortality. For instance, in countries with a poor health infrastructure, the distribution of bednets might be slower as well as the progress in child survival. On the contrary, if international donors target specifically those countries in which public health is quickly deteriorating, a surge in bednets might coincide with a worsening of survival indicators. To get rid of variables determined at the country level, it is necessary to consider a finest geographical level. A good candidate is the geographical cluster of households defined in DHS, covering a village or a block in a city. The objective is to compare, within a country, the evolution of child health between clusters that received many or a few bednets.

4.2 A spatial analysis of the variation in ITN access

As explained in section 3, the challenge is that there is no indicator measuring the evolution of ITN access at the cluster level. For each DHS wave, we observe the proportion of households owning an ITN at the date of the survey in the clusters of the survey. However, sampling units are randomly drawn at each survey. So clusters drawn in the first wave are generally not surveyed in the second wave. Yet, other clusters in the neighborhood are. They provide information on the intensity of anti-malaria campaigns in the region. Our idea is to exploit the spatial correlation in the distribution of bednets.

7. Since we consider a longer time span (2000-2014 for all countries), the median increase in ITN use is larger than in our sample: +50pp.

Formally, we assume that mortality y of a child born to mother i in cluster c at time t is impacted by \overline{ITN} , the coverage of ITN in cluster c at time t :

$$y_{ict} = \alpha \overline{ITN}_{ct} + \gamma x_{ict} + \delta_t + \mu_i + \epsilon_{ict} \quad (1)$$

δ_t is a birthyear fixed effect, capturing all determinants of infant mortality specific to a birth cohort. μ_i is a mother fixed effect, accounting for all time-invariant characteristics related to the survival of children born to the same mother.⁸ x_{ict} is a vector of controls including birth-specific variables (rank, sex and mother's age at birth) and other time-varying characteristics described below. α is a kind of intent-to-treat coefficient: we estimate the impact on infant mortality of distributing bednets to households in a village, without caring for the actual use.

For a given cluster c surveyed in the last wave of DHS, we only observe \overline{ITN} at the date of the last wave. We also observe \overline{ITN} in previous waves for relatively close neighbors. We assume a spatial correlation of the form:

$$\overline{ITN}_{ct} = \rho \overline{ITN}_{n(c)t} + \theta_t + \eta_{ct} \quad (2)$$

With $E(\overline{ITN}_{n(c)t} \cdot \eta_{ct}) = 0$. $n(c)t$ is the nearest neighbor of cluster c surveyed at date t . We restrict our analysis to clusters with one neighbor in a 10 kilometers radius. The main drawback of our strategy is that we cannot exploit half of the clusters, because they have no close neighbor in previous waves. Our assumption is that the distribution of bednets is geographically clustered: campaigns start in some places, reach all villages around and then move to other regions. Empirically, when we estimate equation 2, we find $\hat{\rho} = 0.60$.

8. We have panel data at the mother level. For a given mother i , we observe several births occurring at different dates t .

Replacing equation 2 in equation 1, we get:

$$y_{ict} = \beta \overline{ITN}_{n(c)t} + \gamma x_{ict} + \kappa_t + \mu_i + u_{ict} \quad (3)$$

With $\beta = \alpha \times \rho$ in particular. For each mother i , we observe $\overline{ITN}_{n(c)}$ at each survey wave. We extrapolate $\overline{ITN}_{n(c)t}$ for each birth occurring five years before or after a DHS wave using a linear interpolation. Standard errors are clustered at the neighbor level to account for the correlation between the error terms of births occurring in the neighborhood of the same cluster. We include time-varying socio-economic characteristics of the neighbor in the vector of controls x : the share of mothers and fathers with some education, the average wealth index, the share of households owning a radio and the immunization coverage. The objective is to control as much as possible for the demand-side determinants of ITN ownership. We want to exploit the variation driven by distribution campaigns.

4.3 Results

4.3.1 Main correlation

We estimate equation 3 using a linear model with fixed effects. Results are reported in Table 2. We find $\hat{\beta} = -0.02$ significant at 5%. It corresponds to $\hat{\alpha} = \hat{\beta}/\hat{\rho} = -0.033$. Our estimates predict that infant mortality decreases by 3.3 percentage points when ITN coverage in a cluster moves from null to universal. Over the period of interest, the ITN coverage in our sample rose by 40 percentage points.⁹ It corresponds to a 1.3 percentage points reduction in infant mortality. In terms of magnitude, it amounts to 1/3 of the total decrease between 2000 and 2015, and to 13% of the baseline mortality.¹⁰

Table 3 investigates if estimates are heterogeneous. We find that the association between

9. Our variable of interest, bednet ownership, rose from 30% to 70%. The variation is driven by the variation in ITN use, which increased from 6% to 44%

10. The infant mortality rate in Sub-Saharan Africa decreased by 4pp from 10% in 2000 to 6% today.

ITN coverage and child mortality is stronger for non-educated mothers: $\hat{\beta}^{noedu} = -0.035$. The association for mothers with some education is not significantly different from zero. Results are qualitatively similar when we split the sample on the median wealth index. Disadvantaged households seem therefore to benefit more from anti-malaria campaigns.

4.3.2 Robustness

The estimate of β remains stable when: (i) we use cluster fixed effects instead of mother fixed effects (-0.022 , significant at 1%); (ii) we average all quantities at the cluster level (-0.025 , significant at 5%); (iii) we cluster the standard errors at the mother level (-0.02 , significant at 10%); (iv) we exclude countries one by one.

Next we consider alternative specifications. We first estimate a duration model of survival time with baseline hazards specific to each mother. We find a hazard ratio on $\overline{ITN}_{n(c)t}$ equal to 0.59 and significantly different from one at the 1% level. It means that, after controlling for birth-specific characteristics, a child survives longer than his older sibling when ITN have just been distributed in the neighborhood. Second, we use under-five mortality instead of infant mortality as the dependent variable in a linear model with fixed effects. We find an estimate of similar magnitude, $\hat{\beta}^{under5} = -0.025$.

Last, we directly estimate equation 1 using the measure of ITN use provided by the Malaria Atlas (cf. Table A.2 in Appendix). We find $\hat{\alpha}^{Atlas} = -0.038$ significant at 1%. The order of magnitude is very close to our estimate of α relying on the neighbors. It is important to assess the representativeness of our sample. Remember that we keep only countries with at least two DHS waves and clusters with at least one close neighbor, whereas the Malaria Atlas provides yearly estimates for each cluster. Consequently, the number of observations is 2.5 times higher with the Malaria Atlas than with our baseline strategy. And yet, estimates are very similar, suggesting that our measure of the correlation is credible.

5 Interpretation as a causal effect

So far, we have found that the distribution of ITN in Africa over the period 2000-2015 is associated with a reduction in infant mortality by 1.3 percentage points. The magnitude of our estimate is large compared to experimental results in the medical literature. In this section, we examine if our correlation captures, in addition to the causal impact of bednets on mortality, other variables likely to create an overestimating bias.

Our strategy allows child survival and ITN coverage to be jointly determined by time-invariant characteristics of the mother and by annual shocks affecting all births. The correlation can be interpreted as a causal impact under a strict exogeneity condition: the idiosyncratic error term u_{ict} should be uncorrelated with $\overline{ITN}_{n(c)t}$ of all past, current and future births of the same mother. It means that, in the absence of bednet distributions, the evolution of infant mortality should have been the same for all mothers, whatever the area of residence.

5.1 Endogenous placement of health interventions

The first threat to this assumption is an endogenous placement of bednets. Anti-malaria campaigns are not randomly allocated and the variation in ITN coverage in a region may well capture specific trends in infant mortality. For instance, one can imagine that campaigns targeted some areas depending on the prevalence rate, the population density, the road accessibility or the economic vulnerability. Describing the rules of allocation is beyond the scope of this paper because they may be different for different countries and different periods. A natural extension of our work would be to investigate the case of a specific country and to exploit the determinants of anti-malaria campaigns to design the identification strategy.

In this paper, we assess the scope for such a concern by conducting a placebo test on older birth cohorts. We estimate equation 3 on children born between 1985 and 2003, before the Roll-Back Malaria initiative started. We test if the coverage of ITN ten years after a

birth might explain the probability to die before age one. If there is a correlation, it cannot be driven by the causal impact of bednets on survival. It has to be the case that areas where more bednets have been distributed also have specific trends in mortality. Table 4 shows the results of the placebo test. The coefficient of interest is not significantly different from zero. The magnitude is close to -1 percentage point, twice as low as our estimate of β , yet substantial and of same sign. So the placebo test on older cohorts provides no conclusive evidence that anti-malaria campaigns have been more intense in regions with better trends in infant mortality, although we cannot definitely rule out this possibility.

5.2 Bunching of health interventions

Another threat to the exogeneity assumption is that bednets might be just one component of a global health intervention. For instance in Ethiopia, a massive rural health program started in the mid-2000s. Health workers were in charge of distributing ITNs among other tasks, which makes it difficult to separate out the causal impact of each component. Again, different countries might have coupled anti-malaria campaigns with different interventions at different periods, which complicates the analysis at the continent level.

In this section, we test if the variation in ITN coverage captures other improvements in maternal and child health. The idea is to look at the evolution of neonatal mortality. Children dying in the first month of life do not die from malaria. Indeed, newborns are resistant to the disease during a couple of months, after which they lose their immunity (Hviid & Staalsoe, 2004). Neonatal deaths are primarily caused by premature births as well as asphyxia and trauma during delivery. The evolution of neonatal mortality is therefore a good proxy for health interventions related to maternal health, prenatal care, skilled care at childbirth and access to emergency obstetric care. As a placebo test, we check if our baseline specification predicts the evolution of neonatal mortality. Table 5 shows that the variation in ITN is not related to neonatal deaths. So our measure of malaria control does not seem to

capture other improvements in maternal and neonatal health. However, we cannot rule out that some interventions affecting child survival after the first month of life drive the negative correlation between ITN ownership and mortality.

6 Conclusion

Using an original measure of the variation in bednet coverage, this paper shows that anti-malaria campaigns in Africa over the past 15 year are strongly correlated to a reduction in infant mortality. We find a larger correlation for disadvantaged households, suggesting that the poor have benefited more. The magnitude of the correlation is larger than experimental estimates of the protective effect of bednets. On the one hand, medical RCTs lack external validity. They take place in specific settings with relatively low mortality and low malaria prevalence, providing a lower bound of the causal impact of large-scale campaigns. On the other hand, the surge in bednet ownership over the period 2000-2015 is not a natural experiment: the distribution of bednets was not random, and it could have been coupled with other health interventions. Placebo tests suggest that the correlation might indeed overestimate the causal impact, although not substantially.

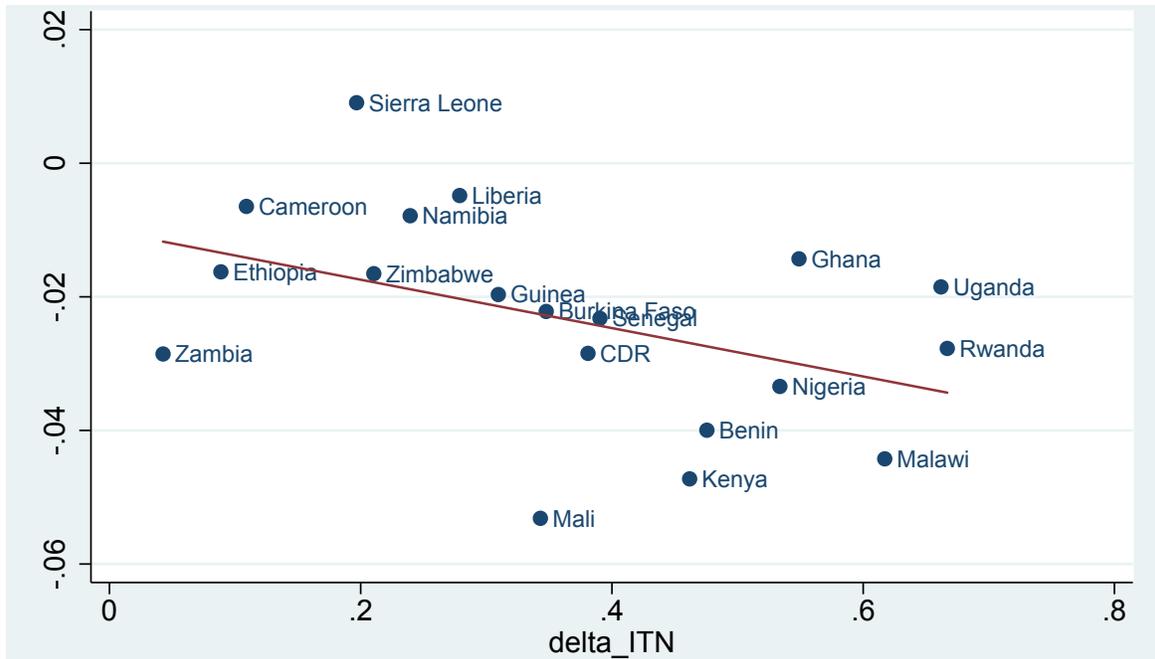
This paper highlights the trade-off between internal and external validity when assessing the impact of malaria control efforts on infant mortality. One way ahead would be to focus on one big country and document the implementation of Roll-Back Malaria: which areas were targeted first and why? Who was in charge of distributing the bednets and what else did they provide? The identification strategy would exploit the roll-out of the campaign and control for other interventions. Moreover, in future research, it will be important to keep in mind that surviving children are selected. By improving survival prospects, anti-malaria campaigns also change the population of survivors. Any analysis of the impact of malaria on economic outcomes in Sub-Saharan Africa should account for this selection.

References

- Apouey, B., Picone, G., Wilde, J., Coleman, J., & Kibler, R. (2016). Paludisme et anÃmie des enfants en afrique subsaharienne : effet de la distribution de moustiquaires. *Revue Economique, Forthcoming*.
- Bhatt, S., Weiss, D. J., Cameron, E., Bisanzio, D., Mappin, B., Dalrymple, U., et al. (2015). The effect of malaria control on plasmodium falciparum in africa between 2000 and 2015. *Nature*, 526, 207-211.
- Black, R., Cousens, S., Johnson, H., Lawn, J., Rudan, I., Bassani, D., et al. (2014). Global, regional, and national causes of child mortality in 2008: a systematic analysis. *The Lancet*, 375 No 9730, 1969-1987.
- Bleakley, H. (2010). Malaria eradication in the americas: A retrospective analysis of childhood exposure. *American Economic Journal: Applied Economics*, 2, 1-45.
- Cutler, D., Fung, W., Kremer, M., Singhal, M., & Vogl, T. (2010). Early-life malaria exposure and adult outcomes: Evidence from malaria eradication in india. *American Economic Journal: Applied Economics*, 2, 72-94.
- Hviid, L., & Staalsoe, T. (2004). Malaria immunity in infants: a special case of a general phenomenon? *Trends in Parasitology*, 20(2), 66-72.
- Kudamatsu, M., Persson, T., & Stromberg, D. (2012). Weather and infant mortality in africa. *Working Paper*.
- Kuecken, M., Thuilliez, J., & Valfort, M. (2015). Large-scale health interventions and education: Evidence from roll back malaria in africa. *Working Paper*.
- Lengeler, C. (2004). Insecticide-treated bed nets and curtains for preventing malaria. *Cochrane Database of Systematic Reviews*.
- Lucas, A. (2010). Malaria eradication and educational attainment: Evidence from paraguay and sri lanka. *American Economic Journal: Applied Economics*, 2, 46-71.
- UN Inter-agency Group for Child Mortality Estimation. (2015). Levels and trends in child mortality. *United Nations Children's Fund*.
- UNICEF. (2014). Generation 2030. *UNICEF publication, Division of Data, Research and Policy*.
- United Nations. (2015). The millenium development goals reports. *United Nations publication, ESA*.
- Venkataramani, A. (2010). Early life exposure to malaria and cognition in adulthood: Evidence from mexico. *Journal of Health Economics*, 31(5), 691-780.
- WHO. (2015). World malaria report. *WHO publication*.

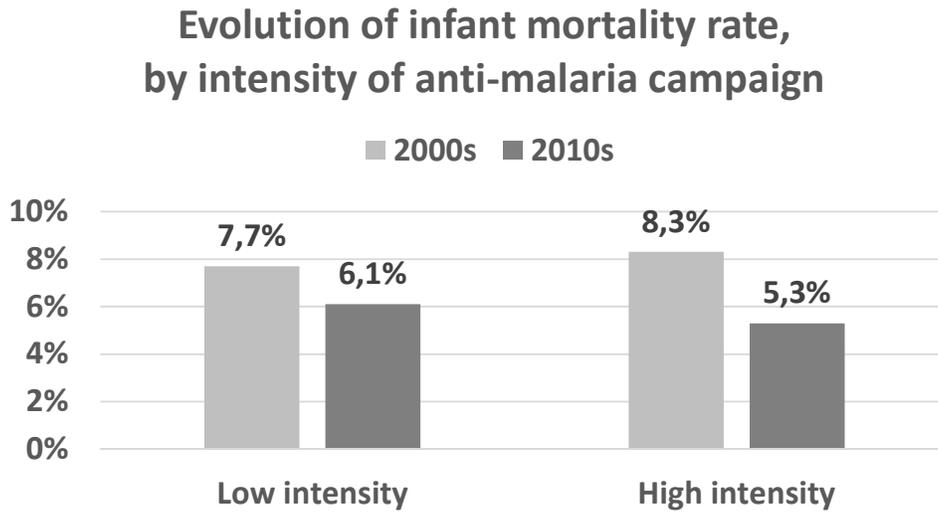
Figures and Tables

FIGURE 1: Panel analysis: country-level



The figure plots the variation in ITN access on the x axis and the variation in infant mortality on the y axis between 2 waves of DHS. Infant mortality rates are computed on children born in the 2 years preceding the survey. Each dot corresponds to a country. The red line corresponds to fitted values predicted by a linear regression. The best fit has a slope of -0.036 .

FIGURE 2: Difference-in-difference: country-level



The figure plots the evolution of infant mortality rates between 2 waves of DHS. Infant mortality rates are computed on children born in the 2 years preceding the survey. The first waves took place in the early 2000s and the second waves in the early 2010s (cf. Table A.1 in Appendix to know the exact years for each country). The evolution is computed separately for two groups of countries: those where the proportion of households owning a bednet increased by more than the median (high intensity of anti-malaria campaign) and those where the increase was below the median (low intensity). Figure 1 above shows which country belongs to which category. The median increase in ITN access is +35pp. The difference in difference estimate is equal to -1.4pp, with a p-value of 0.056.

TABLE 1: Cross-sectional analysis: correlation between ITN use and infant mortality

Dep. var.	Infant mortality	
	Between cluster	Within cluster
ITN use	0.025*** (0.005)	-0.009*** (0.002)
Controls	Mother's age, education and wealth	
Birth year FE	No	Yes
Cluster FE	No	Yes
Observations	10282	166491

Last DHS wave in each country. Children born in last 5 years. Standard errors clustered at the DHS cluster level in the within regression. Weights.

TABLE 2: Panel analysis: retrospective birth histories

Dep. var.	Probability of dying before age 1
% households owning ITN	-0.020** (0.010)
Controls	Birth-specific (mother's age at birth, sex, rank) and socio-economic variables in neighboring cluster
Birthyear FE	Yes
Mother FE	Yes
Observations	206,804

Children born between 1995 and 2013. The explanatory variable measures the proportion of households owning an ITN in the closest neighboring cluster at the date of birth. Standard errors clustered at neighbor level. Weights.

TABLE 3: Heterogenous estimates

Dep. var. Sample	Probability of dying before age 1	
	Non-educated	Educated
% households owning ITN	-0.035** (0.017)	-0.013 (0.012)
Controls	Birth-specific (mother's age at birth, sex, rank) and socio-economic variables in neighboring cluster	
Birthyear FE	Yes	Yes
Mother FE	Yes	Yes
Observations	73,409	133,389

Children born between 1995 and 2013. The explanatory variable measures the proportion of households owning an ITN in the closest neighboring cluster at the date of birth. Standard errors clustered at neighbor level. Weights.

TABLE 4: Placebo 1: impact on older cohorts

Dep. var.	Probability of dying before age 1	
	% households owning ITN <i>10 years after birth</i>	-0.009 (0.016)
Controls	Birth-specific (mother's age at birth, sex, rank) and socio-economic variables in neighboring cluster	
Birthyear FE	Yes	
Mother FE	Yes	
Observations	163,986	

Children born between 1985 and 2003. The explanatory variable measures the proportion of households owning an ITN in the closest neighboring cluster 10 years after birth. Standard errors clustered at neighbor level. Weights.

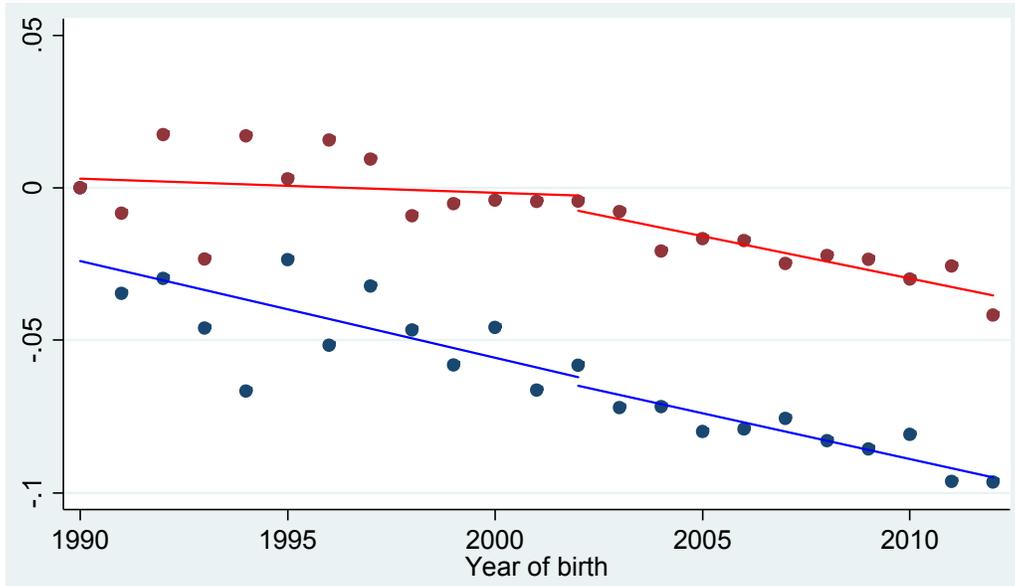
TABLE 5: Placebo 2: impact on neonatal mortality

Dep. var.	Probability of dying before <i>month 1</i>	
	% households owning ITN	-0.006 (0.008)
Controls	Birth-specific (mother's age at birth, sex, rank) and socio-economic variables in neighboring cluster	
Birthyear FE	Yes	
Mother FE	Yes	
Observations	223,842	

Children born between 1995 and 2013. The explanatory variable measures the proportion of households owning an ITN in the closest neighboring cluster at the date of birth. Standard errors clustered at neighbor level. Weights.

Appendix

FIGURE A.1: Trends in infant mortality, by initial prevalence of malaria



The figure plots the coefficients from a regression of infant mortality on year of birth fixed effects. The reference category is 1990. Regions with a high initial level of malaria ($\geq 50\%$) are in red and regions with a low initial level of malaria prevalence ($\leq 50\%$) are in blue. The solid lines represent the linear fit before (1990-2002) and after (2002-2012) the start of anti-malaria campaigns.

The figure shows that pre-trends in infant mortality are not parallel. Over the period 1990-2002, mortality decreased gradually in regions with low initial prevalence, whereas there was no progress in regions with high initial prevalence. Consequently, we cannot perform a difference-in-difference using the variation in initial malaria prevalence to identify the impact of malaria control efforts.

TABLE A.1: Survey waves and years

Country	Survey years
Benin	2001, 2011
Burkina Faso	2003, 2010
Cameroon	2004, 2011
DRC	2007, 2013
Ethiopia	2000, 2005, 2011
Ghana	2003, 2008, 2014
Guinea	2005, 2012
Kenya	2003, 2008, 2014
Liberia	2007, 2013
Malawi	2000, 2004, 2010
Mali	2001, 2006, 2012
Namibia	2000, 2006, 2013
Nigeria	2003, 2008, 2013
Rwanda	2005, 2010
Senegal	2005, 2010, 2012
Sierra Leone	2008, 2013
Uganda	2000, 2006, 2011
Zambia	2007, 2013
Zimbabwe	2005, 2010

TABLE A.2: Using data from the Malaria Atlas

Dep. var.	Probability of dying before age 1
% households using ITN	-0.038*** (0.010)
Controls	Birth-specific (mother's age at birth, sex, rank)
Birth year FE	Yes
Mother FE	Yes
Observations	525,919

Children born between 1995 and 2013. The explanatory variable measures the proportion of households using an ITN in the cluster at the date of birth (measure provided by the Malaria Atlas). Standard errors clustered at the DHS cluster level.