

Quantifying averted burden from infectious disease control policies

Austin Carter

A dissertation

submitted in partial fulfillment of the
requirements for the degree of

Doctor of Philosophy

University of Washington

2025

Reading Committee:

David L Smith, Chair

Hmwe Hmwe Kyu

Nicholas Kassebaum

Thomas Richardson

Program Authorized to Offer Degree:

Department of Health Metrics Sciences

©Copyright 2025

Austin Carter

University of Washington

Abstract

Quantifying averted burden from infectious disease control policies

Austin Carter

Chair of the Supervisory Committee:

David L Smith

Department of Health Metrics Sciences

This dissertation explores methods for calculating averted burden from infectious disease control policies and applies those methods in the contexts of pediatric HIV in sub-Saharan Africa and malaria in Uganda. In my first aim, I describe the Shapley value estimate as the preferred approach to decomposition and then introduce an innovation to the application of Shapley value estimation in the context of interventions implemented at different times with overlapping effects. I call this innovation sequential Shapley value estimation and detail the algorithm for its application. In my second aim, I apply sequential Shapley value estimation to calibrated estimates of pediatric HIV burden in sub-Saharan Africa. I present intervention coverage level for three biomedical interventions and highlight variable impact made through these interventions in preventing new HIV infections and mortality among children under-15. I also estimate avertable burden in 2023 and discuss the implications for future policy. In my third aim, I propose a framework for evaluating outbreak detection approaches using averted burden and apply it to malaria outbreak detection in Uganda.

Aim 1: Estimating the impact of public health policies in terms of averted burden with Shapley value decomposition

Abstract

Attribution of impact to public health policies involving multiple interventions implemented simultaneously is non-trivial. This paper introduces notation for discussing averted and avertable burden, links them to the traditional epidemiological concepts of efficacy and population attributable fraction and then uses the notation to demonstrate the application of Shapley value decomposition to attribute averted burden to interventions. We highlight that another decomposition method, Das Gupta rate decomposition, is a special case of Shapley value decomposition. Finally, we introduce a novel algorithm for the application of Shapley value decomposition in the context of interventions whose impact extends beyond the time of implementation - we call this algorithm sequential Shapley value decomposition.

Introduction

Averted burden is a health metric that quantifies the difference in estimated health burden between two scenarios, often the observed level of burden and an estimated counterfactual that would have occurred under a different set of circumstances. The metric can be used to characterize the impact of an implemented policy or to compare hypothetical policies. In the case of retrospective evaluation of an implemented policy, the observed level of burden is contrasted with the counterfactual level of burden that would have occurred in the absence of the policy. When comparing two hypothetical policies, the difference in the estimated level of burden associated with each policy represents the marginal averted burden from implementing the better policy. Measures of burden include incidence, prevalence, mortality, etc.; common terms for averted burden include “lives saved” or “cases averted”.

Two key challenges arise when attributing averted burden to the separate components of a series of policies implemented over time. The first is attribution of impact in the case of simultaneously implemented interventions, a well understood challenge addressed through decomposition methods. Shapley value decomposition provides a game theoretic approach to attribution with several desirable properties (Shapley, 1952), which we highlight as a general case of the Das Gupta rate decomposition found in economics (Gupta, 1978). The fundamental reason

decomposition is necessary when multiple interventions are implemented simultaneously is that intervention impact varies according to the context within which an intervention is implemented. An example of this for malaria interventions would be the implementation of mass drug administration (MDA). If MDA is implemented in a location where indoor residual spraying (IRS) of insecticide as has driven down malaria transmission, the effect of administering treatment for malaria would be much lower than in a context where malaria transmission is high. If MDA and IRS are implemented simultaneously, it is not clear which intervention should be credited for driving down transmission. Decomposition methods provide approaches to attribution that assess the marginal contributions of an intervention across the multiple contexts within which the intervention could have been combined with other interventions.

The second challenge faced when attributing impact is quantifying the time-varying impact of an intervention whose effects are not isolated to the time of intervention implementation. Policies with overlapping challenge the traditional approach of decomposing each time-point separately without consideration for other policies. We introduce a novel algorithmic approach to decomposition called sequential Shapley value estimation, which captures policy impacts conditional on already implemented policies and independent of possible future policies. We argue that the impact associated with a given policy should be conditional on all previous policies and independent of future policies. A clear description of the algorithm and an example of its results enable implementation of this approach to generating sensible policy impact estimates that are attributed across interventions.

Impact

In order to clearly discuss decomposition and averted burden, we start by introducing notation and terminology for bridging between common epidemiological concepts like intervention coverage and efficacy to and the techniques necessary for producing and processing counterfactual estimates of burden into attributed impact. Starting with a single intervention, we introduce formulae for averted and avertable burden, and then progress to the case of multiple interventions, where decomposition methods become necessary.

Impact for a single intervention

The impact of a single intervention is estimated by comparing the counterfactual level of burden in the absence of the intervention to the observed level of burden. We call this difference in burden the averted burden attributable to the intervention. To construct the counterfactual, we need a model of how the level of burden varies according to levels of intervention coverage. A simple model of this relationship requires three inputs: the observed level of burden - $B(X = \hat{X})$, the observed level of intervention coverage - \hat{X} , and the efficacy of the intervention - ε_X . The efficacy of an intervention can be conceived as the ratio of burden in the presence of 100% coverage of intervention X to the burden in the absence of intervention X :

$$\varepsilon_X = \frac{B(1)}{B(0)}$$

Estimates of this ratio arise from research, where individuals or populations are randomly assigned to receive an intervention - the treatment group - or not receive the intervention - the control group. Comparison of the burden occurring in the treatment and control group provides estimates of efficacy, also referred to as the relative-risk associated with an intervention.

The observed level of burden in the context of some level of intervention coverage, $B(X = \hat{X})$, is estimated through observational data, for example deaths recorded in vital registration systems of or new cases reported by health facilities. The observed and the counterfactual burden can then be linked through an equation that scales the counterfactual burden according to those protected from the intervention. A simple static example is:

$$B(X) = B(0) \cdot (1 - \varepsilon_X \cdot X)$$

where the term $1 - \varepsilon_X \cdot X$ represents people not protected from disease by the intervention. From this equation we can solve for $B(0)$ and compare it to $B(\hat{X})$, to obtain an estimate of averted burden, $A(\hat{X})$:

$$A(\hat{X}) = B(0) - B(\hat{X})$$

Similarly, the avertable burden, $\bar{A}(\hat{X})$, can be estimated through comparing the observed burden to the burden that would have occurred in the presence of 100% coverage of intervention X :

$$B(1) = B(0) \cdot (1 - \varepsilon_X)$$

$$\bar{A}(\hat{X}) = B(\hat{X}) - B(1)$$

Avertable burden can be connected to the traditional epidemiological concept of the population attributable fraction (PAF) by taking the ratio of the avertable burden to the observed burden:

$$PAF(\hat{X}) = \frac{\bar{A}(\hat{X})}{B(\hat{X})}$$

The complement, for averted burden, is the burden averted fraction (BAF):

$$BAF(\hat{X}) = \frac{A(\hat{X})}{B(\hat{X})}$$

Single intervention example

To illustrate each of these concepts with an example (Figure 1), assume that we estimate 100 cases of a disease occurred in the presence of 40% coverage from an intervention X , ($B(\hat{X}) = B(X = 0.4) = 100$), and the intervention reduces cases by 50%, ($\epsilon_x = 0.5$). With these three inputs to Equation 1, we can solve for the counterfactual level of burden, $B(0)$:

$$100 = B(0) \cdot (1 - 0.4 \cdot 0.5) \Rightarrow B(0) = \frac{100}{1 - 0.2} = 125$$

The averted burden is then: $A(\hat{X}) = B(0) - B(\hat{X}) = 125 - 100 = 25$. Burden in the context of 100% coverage is: $B(1) = B(0) \cdot (1 - \epsilon_x) = 125 \cdot (1 - 0.5) = 62.5$. Avertable burden is then:

$$\bar{A}(\hat{X}) = B(\hat{X}) - B(1) = 100 - 62.5 = 37.5$$

We can now calculate the PAF using Levin's formula as an alternative approach. With protective interventions – in contrast to risk factors – we need to use the complement of intervention coverage as our prevalence of exposure, and we use the inverse of the efficacy to get the relative-risk of exposure. Continuing with the example above we get:

$$\begin{aligned} PAF(\hat{X}) &= \frac{P_e \cdot (RR - 1)}{P_e \cdot (RR - 1) + 1} = \frac{(1 - \hat{X}) \cdot (\epsilon_x^{-1} - 1)}{(1 - \hat{X}) \cdot (\epsilon_x^{-1} - 1) + 1} \\ &= \frac{0.6 \cdot (2 - 1)}{0.6 \cdot (2 - 1) + 1} = \frac{0.6}{1.6} = 0.375 \end{aligned}$$

Now, rearranging our previous definition of PAF, we can calculate the avertable burden: $\bar{A}(\hat{X}) = PAF_{\hat{X}} \cdot B(\hat{X}) = 0.375 \cdot 100 = 37.5$. This lines up with our counterfactual based approach to calculating avertable burden. The same calculation can be done to validate the relationship between the BAF and averted burden.




Fig. 1. An example illustrating averted burden, $A(\hat{X})$, and avertable burden, $\bar{A}(\hat{X})$, for a disease with an observed level of burden, $B(\hat{X})$. The x-axis is the level of coverage for intervention X and the y-axis is the burden of disease associated with different levels of coverage.

Impact for two interventions

Expanding from assessing the impact of one intervention to two reveals the need for decomposition methods. Two interventions implemented simultaneously can interact, modifying each others effects on burden. This leads to a challenge in evaluating impact, where the effect of removing an intervention depends on whether the other intervention is present or not. If we assess the impact of each intervention in the presence of the other, we estimate much smaller

effects than when we assess the impact of each intervention when the other intervention is absent. Further, we want estimates of impact that, when added together, total an amount equal to the difference between the observed burden and the level of burden in the absence of both interventions.

Two intervention example

In Figure 2, we present an example where we have added a second intervention, expanding our model of the relationship between burden and intervention coverage to be:

$$B(X) = B(0)(1 - \varepsilon_{X1} \cdot X1)(1 - \varepsilon_{X2} \cdot X2)$$

where $X1$ and $X2$ represent the coverages of each intervention and ε_{X1} and ε_{X2} represent their efficacies, respectively. For example, we set $B(0)$ again to 125, the observed level of coverages for $X1$ and $X2$ to 60% and 70%, and their efficacies to 0.5 and 0.25. In Figure 2A, we show a heatmap with the value of burden for all possible levels of intervention coverage for each intervention, and overlay points representing the observed level of coverage, the absence of each intervention, and the absence of both interventions. The arrows show the direction of moving from the observed coverage level to the absence of coverage through two paths that vary in which intervention is removed first. In Figures 2B and 2C, we look at the marginal level of burden for different levels of each intervention while holding the other intervention at a fixed level. In these, we can see that the difference in burden associated with an intervention's impact depends on the level of the other intervention. Finally, we show in Figure 2D that the two pathways lead to the same total level of burden of averted, with the different amounts of averted burden attributed to each intervention.

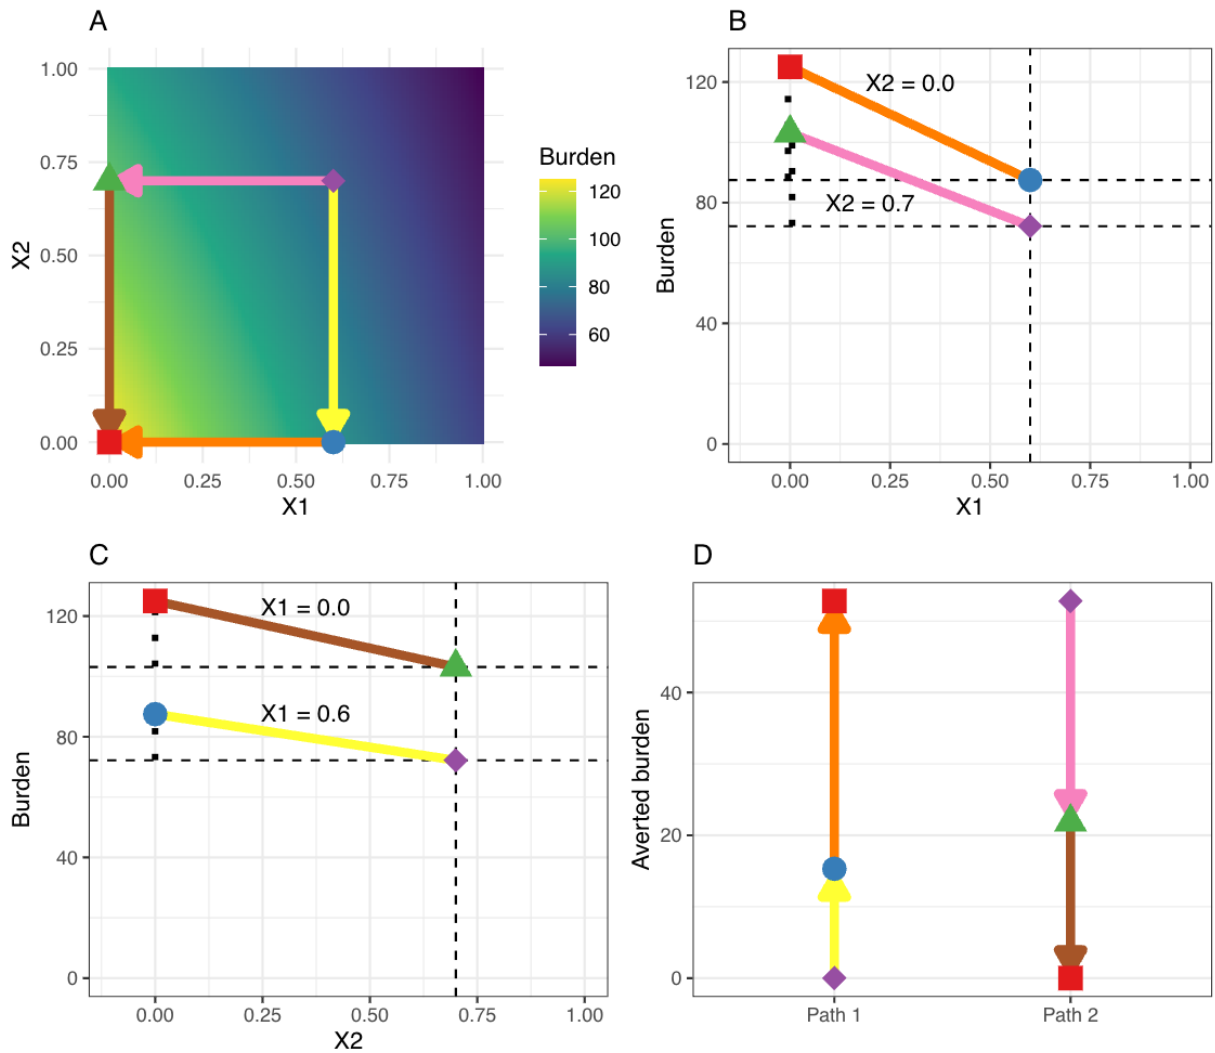


Fig. 2. (A) A heatmap showing the level of burden according to different levels of coverage for interventions X1 and X2. The colored symbols represent different scenarios and the arrows indicate the pathways of moving from the observed level of coverages to zero coverages (B) The marginal value of burden according to different levels of coverage for X1 at specific coverage levels for X2. (C) The marginal value of burden according to different levels of coverage for X2 at specific coverage levels for X1. (D) The averted burden for each step in the two pathways from observed coverages to zero coverages.

Decomposition

In general, decomposition is the activity of breaking something down into parts. In the context of health metrics, decomposition refers to the attribution of changes in health outcomes over time or

between locations to changes in known drivers of the health outcome. Drivers of health outcomes can be broken up into biomedical factors, behavioral factors, and structural factors. Given a time series of health outcomes and drivers in one location, execution of a decomposition results in a set of proportions characterizing the amount of change in the health outcome attributable to variation in each driver. It is possible to have both positive and negative proportions, corresponding to drivers either decreasing or increasing disease burden, respectively. A reduction in burden associated with a driver is referred to as the averted burden attributable to the driver. The steps of producing a decomposition are (1) define a model relating drivers to specific health outcomes, (2) fit the model to data from a specific location and time period, and (3) apply decomposition methods to the calibrated model to attribute changes in outcomes to changes in drivers. Many valid approaches to decomposition exist, but the best approach to decomposing drivers of infectious disease outcomes is Shapley value decomposition coupled with a mathematical model relating drivers to disease dynamics.

Shapley value decomposition

$P = \{a, b, c\}$ is a policy consisting of interventions a , b , and c . In the context of our child HIV decomposition, we are investigating policies that include three interventions: a , b , and c . The total number of policy orderings for three interventions is $3! = 6$:

1st	2nd	3rd
a	b	c
a	c	b
b	a	c
b	c	a
c	a	b
c	b	a

Table 1: Policy orderings for three interventions

For a policy with three interventions, there are $2^3 = 8$ possible simulations arising from all combinations of each intervention being either implemented or not implemented. These are eight simulations then form four simulation pairs where a given intervention is added to a set of other interventions. Let C be the set of interventions implemented in a particular simulation. For intervention a , the four pairs of simulations that provide the marginal contribution of a to an existing set of interventions, C , are:

Pair	C	C ∪ a	Policy orderings
1	∅	{a}	2
2	{b}	{a, b}	1
3	{c}	{a, c}	1
4	{b, c}	{a, b, c}	2

Table 2: Counterfactual pairings and the number of policy orderings that include each pair

The Shapley value for intervention i , $\phi(i)$, is the weighted average of the marginal value of intervention i when added to policy C . The Shapley value equation is:

$$\phi(i) = \sum_{C \subseteq N \setminus \{i\}} \frac{|C|! (n - |C| - 1)!}{n!} (v(C \cup \{i\}) - v(C))$$

where n is the total number of interventions under consideration, C is a policy that excludes intervention i , and $v(S)$ is the output of a simulation with intervention set S . $\frac{|C|!(n-|C|-1)!}{n!}$ is the weight which accounts for the frequency with which intervention i is added to set C across all permutations of policy orderings. $v(C \cup \{i\}) - v(C)$ is the change in the value of the simulation when i is added to a policy (i.e. the marginal value of i).

Continuing our example with three interventions, we can see that the Shapley value for intervention a is:

$$\begin{aligned} \phi(a) &= \sum_{C \subseteq N \setminus \{a\}} \frac{|C|! (3 - |C| - 1)!}{3!} (v(C \cup \{a\}) - v(C)) \\ &= \frac{2}{6} (v(\{a\}) - v(\emptyset)) + \frac{1}{6} (v(\{a, b\}) - v\{a\}) + \\ &\quad \frac{1}{6} (v(\{a, c\}) - v\{a\}) + \frac{2}{6} (v(\{a, b, c\}) - v\{b, c\}) \end{aligned}$$

we can see that the numerator in the weights lines up with the frequency that a pairing appears in the permutations of policy orderings found in Table 1, and the denominator is the total number of permutations of policy ordering.

Das Gupta decomposition is a special case of Shapley value decomposition

Das Gupta decomposition is a method that was introduced in economics for decomposing changes in rates that are the result of the product of multiple factors. To fend off confusion about the differences between Das Gupta and Shapley value decomposition, we briefly highlight that Das Gupta decomposition, which is the decomposition method of choice for changes in rates over time in a number of public health papers, is a special case of Shapley value decomposition where the function translating factors into outputs is simply the product of all factors. The Das Gupta equation for decomposing the contributions of each factor to changes in rates between time periods a and b is:

$$\phi(i) = \sum_{j=1}^{n-1} [n \binom{n-1}{j-1}]^{-1} R(j-1, i) (x_i^a - x_i^b)$$

where n is the total number of factors under consideration, $R(j-1, i)$ is the product of all non- i factor combinations of size $j-1$, and x_i^t is the intervention value at time t . $[n \binom{n-1}{j-1}]^{-1}$ is the weight associated with each product difference, and $R(j-1, i) (x_i^a - x_i^b)$ is the marginal value of factor i when changing from the level at time a to the level at time b in the context of all other $j-1$ factor products.

To demonstrate that Das Gupta is a special of Shapley, we must show that the weights are equivalent and that the term for the marginal contribution of a factor to the changes in the product of factors aligns with the Shapley idea of changes in simulation output values in the presence of absence of a factor.

The equivalence of the weights:

$$\begin{aligned} [n \binom{n-1}{j-1}]^{-1} &= \left[\frac{n(n-1)!}{(j-1)! (n-j-2)!} \right]^{-1} \\ &= \frac{(j-1)! (n-j-2)!}{n!} \Big|_{j-1=|C|} \\ &= \frac{|C|! (n-|C|-1)!}{n!} \end{aligned}$$

Similarly, the marginal values are equivalent when we recognize that the difference in simulation outputs from Shapley is the same as the difference in the product of factors in Das Gupta:

$$(v(C \cup i) - v(C)) = R(j - 1, i)(x_i^a - x_i^b)$$

Sequential Shapley value decomposition

In its original form, Shapley value decomposition applies to a single policy whose impact occurs immediately. In reality, many interventions have effects that extend beyond the time period of their administration. To account for this, we propose sequential Shapley value decomposition, which involves step forward in time within a simulation and applying the policy for each time point conditional all preceding policies and independent of future policies. This is achieved through the following algorithm:

Algorithm 1 Sequential Shapley value estimation

Require: $1 \leq t \leq N$

$i \leftarrow 1$

while $i \leq N$ **do**

$P(t) = \hat{P}(t) \quad \forall t < i$

▷ Set all policies prior to time i to their observed values

$P(t) = 0 \quad \forall t > i$

▷ Set all policies after time i to zero

Execute Shapley value decomposition for $P(i)$

$i \leftarrow i + 1$

Sum Shapley values for each intervention across time points to get the total contribution of each intervention

In Figure 3, we present an example of the application of sequential Shapley value estimation. Again, we are working with same model of burden that is modified by two interventions, $X1$ and $X2$. We modify the coverage of these two interventions to have temporal impact profiles, where their efficacy decays over time, leading to effective coverages that vary of time. Intervention $X1$ is administered every 60 days and decays in efficacy much faster than intervention $X2$ which is administered only once (Fig 3A). The effective intervention coverage can be calculated by taking the convolution of the temporal impact profiles with a binary vector logging the time points where the intervention was administered with a one and all other time points with a zero. With a value for the effective intervention coverages at each time point, we can calculate the level of burden over time, which now exhibits variation according to the coverage levels at each time point (Fig 3B).

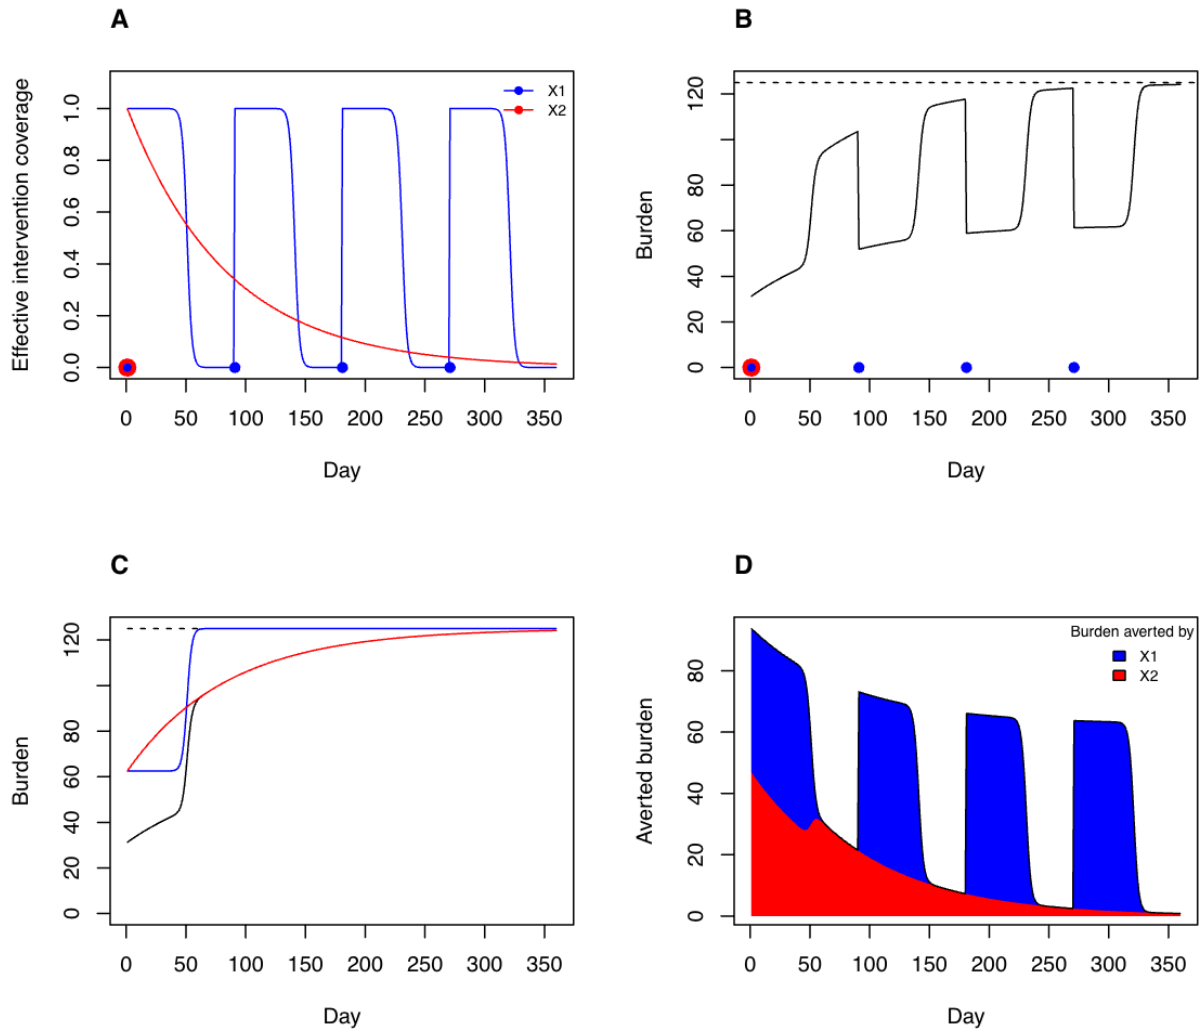


Fig. 3. (A) The effective intervention coverage levels over time for two interventions. Intervention X1 is implemented four times and has a faster decay in effectiveness than intervention X2, which is only implemented once. (B) Burden levels over time, varying as intervention coverage levels vary. (C) An example of burden levels for four scenarios (1) observed in black (2) only X1 is present in blue (3) only X2 is present in red and (4) no interventions are present in the dashed black line. (D) The averted burden attributable to each intervention estimated through application of sequential Shapley value decomposition

Now, with a series of policies implemented over time with overlapping impact, we must use the sequential Shapley value approach to accurately attribute impact to each intervention time point. This involves stepping forward through the time points in which policies were implemented at executing Shapley value decomposition with a set of scenarios that include the observed policies that were implemented in the past and none of the policies that are implemented in the future, while simulating all permutations of presence or absence of each intervention in the present. The first policy, implemented on Day 1, involves the administration of both interventions simultaneously, leading to the simulation of four combinations of intervention coverages and

four burden time series (Fig 3C). We then process these counterfactuals to calculate the weighted average marginal contribution in each time point throughout the period under assessment. This process is repeated for each time point in which a policy is administered, leading to a set of intervention and time of administration specific time series of averted burden. These can then be summed across time points of administration for each intervention to get the total impact of each intervention at every point in time (Fig 3D).

Of note in our example, the only time point of administration in which both interventions were implemented was the first day, and the impacts of each intervention are clearly modified by the presence of the other. In contrast, the future implementation of $X1$ does not modify the impact of $X2$ which was implemented in the past, but $X2$ does modify the impact of $X1$, as the impact of future implementations of $X1$ are conditional on the presence of $X2$ coverage. This reinforces the logic of evaluating policies conditional on the past but independent of the future.

Conclusion

This manuscript provides a guide to researchers looking to generate estimates of averted burden and attribute impact to individual interventions implemented as part of a series of policies. Estimates of averted burden are critical inputs to cost-effectiveness estimates and provide better summaries of impact than simply assessing variation in rates of burden. In addition, averted and avertable burden can serve as key inputs to the formulation of future policies and optimal resource allocation.

We motivate the use of decomposition to assess intervention impact while clarifying the methods available to estimate and attribute averted burden in a variety of policy settings. Through introduction of notation, comparison to common epidemiological metrics, and a series of examples, we prepare researchers to apply decomposition methods to their analyses and propose further innovations that address the many challenges that arise when estimating how action modified the events of the past. Future work is needed to highlight how to apply these methods across various models of disease modified by health policy. Two particular areas that invite focus are (1) decomposing contributions when the factors under assessment have a hierarchical relationship and (2) addressing the propagation of uncertainty through the averted burden estimation process. Increasingly, there is demand for estimates of averted burden, and clarification of these methods in our manuscript should improve the clarity of their application in quantifying public health intervention impact.

Citations

- Gupta, P. D. (1978). A General Method of Decomposing a Difference Between Two Rates into Several Components. *Demography*, 15(1), 99–112. JSTOR. <https://doi.org/10.2307/2060493>
- Shapley, L. S. (1952). *A Value for n-Person Games*. <https://www.rand.org/pubs/papers/P0295.html>

Aim 2: Averted and avertable burden from biomedical interventions for pediatric HIV in sub-Saharan Africa from 2000 to 2023

Abstract

Research in context

Evidence before this study

Estimates of children living with HIV indicate great progress in the reduction of new infections and HIV mortality in individuals ages 0 to 14. In addition, reported coverage of interventions focused on preventing mother-to-child-transmission and reducing the burden of pediatric HIV has increased dramatically in the last two decades. Despite this progress, further improvements are required to eliminate HIV burden among children and adolescents, and lessons should be gleaned from countries that have effectively reduced pediatric HIV burden. To our knowledge, no comprehensive estimates of averted burden attributable to biomedical pediatric HIV interventions in sub-Saharan Africa has been published.

Added value of this study

This study provides the first ever decomposition of averted and avertable pediatric HIV in sub-Saharan Africa. We apply novel decomposition methods that account for impact occurring beyond the time period of intervention delivery to recent pediatric HIV burden estimates, enabling attribution of impact to biomedical interventions delivered since 2000. Our country-level analysis allows for direct comparisons of averted burden achieved over time, and it supports the highlighting of suites of interventions that have been most impactful. Finally, estimates of avertable burden provide insight into areas to focus efforts to increase intervention coverage moving forward.

Implications of all the available evidence

Our results indicate that while most of the previous reductions in pediatric HIV burden were achieved through increased coverage of ART among adults, most of the avertable burden that remains requires improved PMTCT coverage. In addition, this work highlights the dramatic impact achieved through all pediatric HIV biomedical interventions, including PMTCT, ART,

and cotrimoxazole, and these results can be used to support future investment in reducing pediatric HIV through improved coverage of biomedical intervention

Introduction

Over 1 million children under the age of fifteen were living with HIV in 2023, most of whom were infected through mother-to-child transmission (MTCT) during pregnancy, childbirth, or breastfeeding (Carter et al., 2024). While this high level of pediatric HIV burden is tragic, it also reflects a dramatic global improvement in the prevention of mother to child transmission and extension of life through treatment, as new infections in children have decreased by 65.6% between 2000 and 2023, and HIV-related mortality in children has declined by 79.0% over the same period. Prevention of mother-to-child transmission (PMTCT) requires identifying pregnant women living with HIV and enrolling them in treatment protocols which reduce their viral load and therefore their probability of MTCT. Increased HIV testing of pregnant women and enrollment in PMTCT can be credited with bringing about the dramatic reduction in global HIV burden among children over the past quarter century.

In 2023, the largest number children living with HIV (CLHIV) resided in sub-Saharan Africa, with 91.9% of all children living with HIV residing in SSA. Further, incidence rates in SSA are more than eight times higher in 2023 than any other region. For this reason, we focus on the region for our assessment of the impact of biomedical pediatric HIV interventions. Biomedical interventions, including treatment regimens for prevention of mother to child transmission (PMTCT), anti-retroviral therapy (ART) for children, and cotrimoxazole (cotrim) – an antibiotic which reduces mortality in HIV positive children – all contributed to observed reductions in child HIV burden. WHO recommendations on the use of these biomedical interventions for the prevention of HIV burden in children first came about in the early 2000s, driving increased coverage in resource constrained settings (World Health Organization, 2006, 2009, 2013). Increased access and uptake of these interventions, along with growing evidence for expanded eligibility and safety, led to dramatic reductions in MTCT and child HIV mortality, inspiring hope for elimination of child HIV. Repeatedly, goals targeting elimination have not been achieved, but nonetheless progress was made. Variation in policies and uptake of interventions across locations resulted in variable progress in fighting child HIV across locations. In this study we intend to assess that variation and attribute impact to the interventions that were delivered.

Several prior studies have assessed HIV program impacts, particularly those of the United States' President's Emergency Plan for AIDS Relief (PEPFAR), which funds a wide range of interventions including PMTCT and ART in SSA. Stover et al. retrospectively assessed the impact of PEPFAR on children, finding the 1.5 million child HIV-related deaths and 2.8 million new child HIV infections were prevented by the program between 2004 and 2022 (Stover et al.,

2023). This type of all-inclusive analysis provides helpful numbers for advocacy but does not assess the impact from each intervention. Although detailed insights into the specific contributions of biomedical interventions could significantly aid policymakers in understanding their relative importance and the context needed to achieve HIV prevention and control targets, such an analysis has not yet been conducted. This study represents the first-ever comprehensive decomposition of both averted and avertable pediatric HIV burden in sub-Saharan Africa.

To accurately attribute impact to each intervention, our study applies innovative decomposition methods that consider intervention impacts extending beyond the time point of intervention administration, this study not only enables the attribution of impacts to biomedical interventions since 2000 but also facilitates a nuanced country-level comparison of the averted burden over time. The output of the decomposition is averted burden attributable to each intervention for years 2000 to 2023 in every country in SSA. These results allow for assessment of which countries achieved the greatest success in preventing HIV burden, and for comparison of differences in policy combinations across locations. Finally, estimates of avertable deaths by intervention in 2023 point to where the greatest future opportunities lie. These results are critical to highlighting exemplar countries while offering insights into the path forward for pursuing the elimination of HIV/AIDS in children. Accordingly, the objective of this study is to summarize the impact of biomedical interventions in averting pediatric HIV burden across countries and over time in sub-Saharan Africa, allowing for assessment of progress made and gaps that remain in the effort to eliminate pediatric HIV.

Methods

Approach

The decomposition of averted and avertable HIV burden in children starts with location-specific estimates of the historical level of burden that occurred in the past – a modeling exercise that we carry out for each publication of the Global Burden of Disease. Our estimates result from calibrating a mechanistic model of HIV dynamics to available data on health outcomes and intervention coverages. With a calibrated model of what occurred in each location, we are then able to simulate counterfactuals by modifying intervention coverages and re-simulating HIV burden. One complication that arises in attributing averted burden to interventions is that intervention impact varies according to the context that an intervention is implemented. For example, the impact of ART in preventing child HIV-related deaths is much larger in a context where there is no PMTCT than when PMTCT has prevented many child HIV infections.

To address this, we simulate all permutations of combinations of interventions and assess the marginal impact of an intervention in all contexts in which it could have been implemented. This method of decomposition is called Shapley value estimation. Further, we employ a novel algorithmic approach to Shapley value decomposition that accounts for the fact that interventions implemented at different time points can have overlapping impacts when their effects extend

beyond the time-period of their implementation. We call this method sequential Shapley value estimation, and it enables us to assess the effects of interventions conditional on policies already implemented in the past, but independent of policies implemented in the future. The results of our decomposition are location- and year-specific averted HIV infections and HIV-related deaths for each intervention, and the sum of each of the individual intervention impacts is equal to the total impact of all biomedical intervention.

Pediatric HIV simulation

Our approach to decomposition relies on a mechanistic model of the effects of biomedical interventions on HIV transmission and mortality. Mathematical models representing the dynamics of mother-to-child-transmission, the natural history of HIV in children, and the effects of biomedical interventions on these dynamics (Stover et al., 2021) were developed for UNAIDS' estimation of HIV burden through the Spectrum software. The pediatric component of the model simulates the transmission of HIV from mother to child, sensitive to the prevalence of HIV in pregnant women and the level of PMTCT coverage among HIV positive pregnant women. Transmission can occur perinatally or via breastfeeding after birth, so transmission rates are also sensitive to the frequency of breastfeeding among HIV positive mothers. After calculation of transmission, disease progression is tracked by CD4 categories. Mortality rates vary by CD4 level, ART status (split by time since initiation and CD4 count at initiation), and TMP-SMX prophylaxis status. In summary, incidence rates are modified by PMTCT rates, and mortality rates are modified by ART and TMP-SMX. These dynamics are embedded within a population projection that includes birth, death, migration, and ageing, enabling investigation of the temporal evolution of disease modified by biomedical factors.

Sequential Shapley value decomposition

Sequential Shapley value estimation involves the repeated application of Shapley value decomposition at different time points, enabling calculation of the impact of a policy executed at a given time point conditional on interventions implemented in the past and independent of future interventions (Shapley, 1952). Practically, this means we start with intervention coverage levels at zero for all years except the first time point and execute Shapley value decomposition on the interventions delivered as part of the first year policy. We then move to performing Shapley decomposition on the second time point, keeping the first time point coverage levels at their observed levels and all future time points set to zero. We continue moving forward sequentially until a unique Shapley value decomposition is conducted for each year of the analysis.

Avertable HIV-related deaths

To assess the effects of gaps in intervention coverage, we calculated the avertable burden associated with the coverage gaps present in 2023. This was achieved by comparing the observed level of pediatric HIV-related deaths in 2023 to a counterfactual where all interventions have 100% intervention coverage. To attribute the change in the number of deaths to each intervention, we executed Shapley value decomposition for year 2023, looking at the difference between observed coverage and full coverage for each intervention combination. This contrasts with the decomposition used to calculate averted burden which compared observed coverage to the absence of coverage for each intervention combination. Aside from the specification of the counterfactual coverage level, none of the other details of the decomposition differed.

Results

Between 2000 and 2023, incident HIV (fig 1a) and HIV deaths (fig 1b) in children <15 years of age declined in all countries in SSA. The total change in new infections among children across SSA was a reduction from 627,307.7 in 2000 to 208,667.5 in 2023, a 66.7% decline. For mortality, the number of HIV-related deaths among children dropped from 231,123.1 in 2000 to 50,530.6 in 2023, a 78.1% decline. The largest declines in incidence were observed in Zimbabwe, Burkina Faso, and Malawi, who saw declines of 95.7%, 95.1%, and 94.9%, respectively. For mortality, the largest declines were observed in Burkina Faso, Botswana, and Malawi, who saw declines of 97.6%, 96.9%, and 95.3%, respectively. Between the four regions in SSA – Central, Eastern, Southern, and Western – the largest overall declines in new infections in children occurred in Southern SSA, where new HIV infections dropped from 131,834 in 2000 to 13,003 in 2023, a 90.1% decline. The region with the largest declines in child HIV-related deaths was also Southern SSA, where child HIV related deaths dropped from 37,460 in 2000 to 4228 in 2023, an 89% decline.

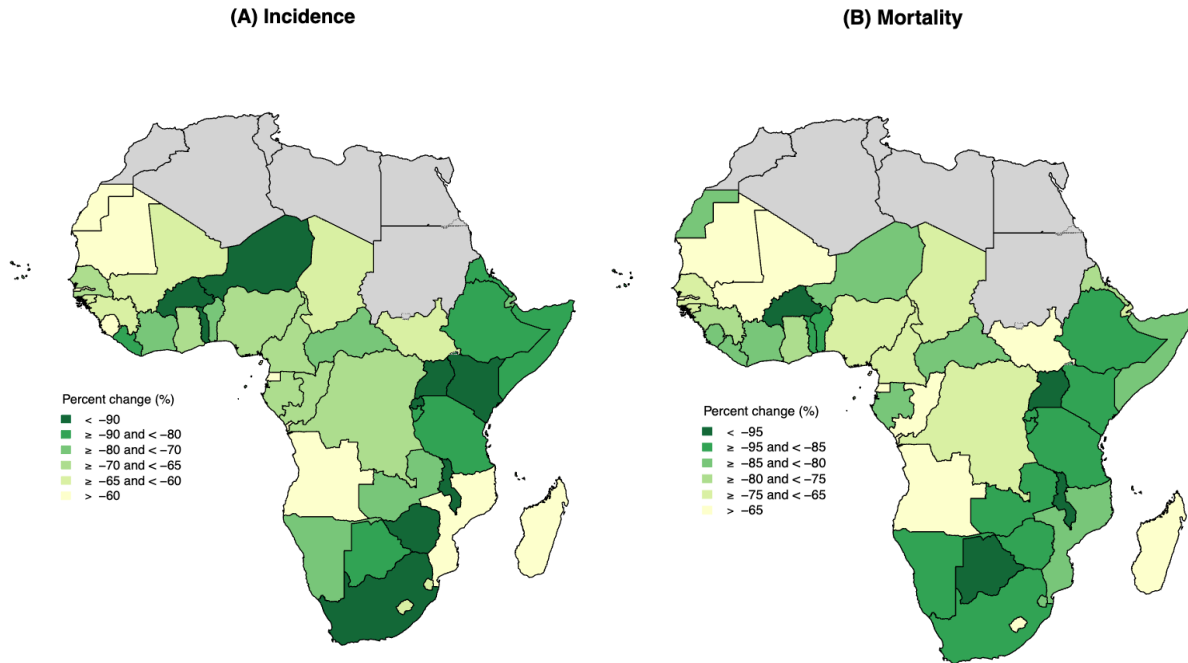


Figure 1. Map of the percentage change in incidence of pediatric HIV between 2000 and 2023

Over the same period, intervention coverage for PMTCT, TMP-SMX, and ART rose across all of SSA. The total number of women receiving PMTCT in SSA grew from less than five thousand in 2000 to 2.6 million in 2023, a jump from 0.2% coverage in 2000 to 96.7% coverage in 2023. The total number of CLHIV receiving ART and cotrim grew from 2105.2 and 6648.6 in 2000 to 478,525.6 and 165,507.0 in 2023, which resulted in an increase in ART coverage from 0.1% in 2000 to 61.9% in 2023 and an increase TMP-SMX coverage from 0.6% in 2000 to 74.9% in 2023. The region with the highest level of child ART coverage in 2023 was Eastern SSA, with coverage of 73.1% (Figure 2). For TMP-SMX, the region with the highest levels of coverage was Western SSA, with coverage of 83.8%. Finally, the highest level of PMTCT coverage in 2023 was found in Southern SSA, with coverage of 99.8%.

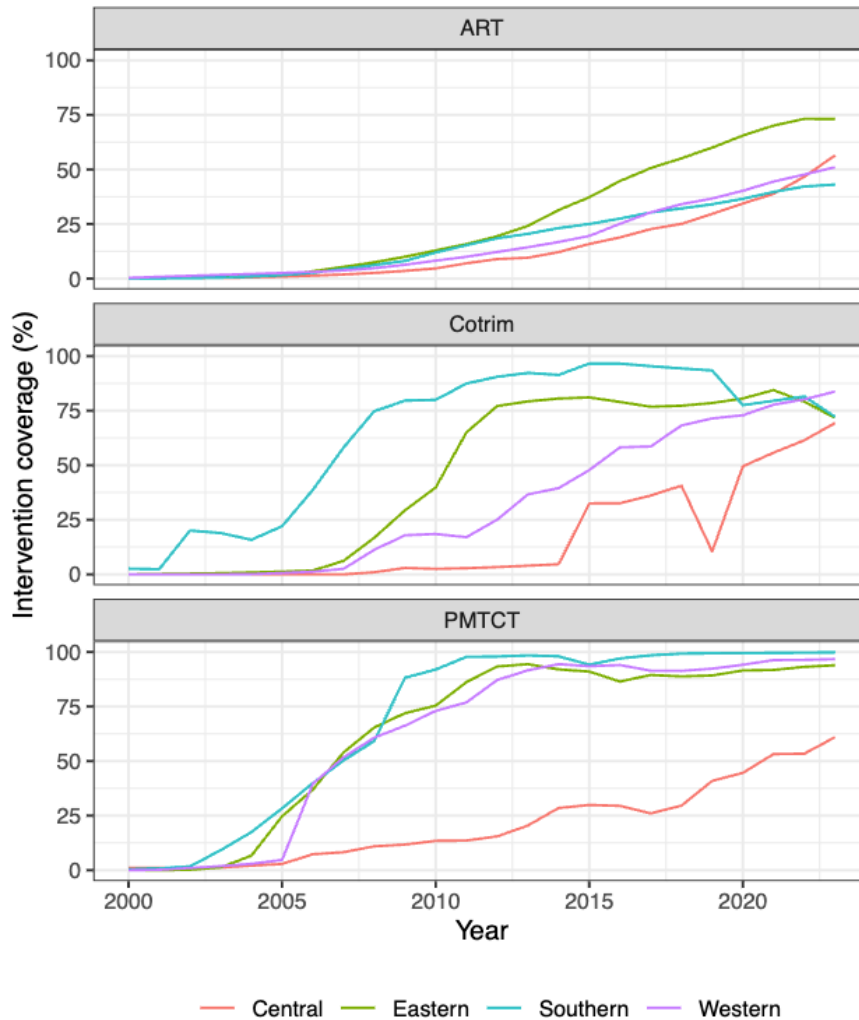


Fig 2. Pediatric biomedical intervention coverage by SSA region between 2000 and 2023

The decomposition estimates indicate that a total of 4.8 million new HIV infections and 3.6 million HIV-related deaths were averted due to all evaluated interventions combined delivered in SSA between 2000 and 2023. The country with the largest averted pediatric HIV burden was South Africa, with a total of 1.3 million new HIV infections averted and 920,000 HIV-related deaths averted. All cases averted were due to PMTCT, but the proportion of deaths averted attributable to each intervention varied by country and year. When aggregating intervention impact across all years by country and intervention, the country with the largest proportion of deaths averted due to PMTCT was Sierra Leone, with 89.8% of all HIV-related deaths averted, and the smallest proportion averted due to PMTCT was Democratic Republic of the Congo with 10.4%. For ART, the largest proportion was in Democratic Republic of the Congo with 64.6% of all deaths averted, and the smallest proportion averted due to ART was

Sierra Leone with 6.4%. For TMP-SMX, the largest proportion was in Somalia with 31.9% of all deaths averted, and the smallest proportion was in Angola with 1.6%.




Fig 3. Deaths and infections averted by country and intervention

Estimates of avertable burden by intervention and region in 2023 indicate that the largest opportunity for the reduction of pediatric HIV-related deaths is in scaling up child ART coverage in the Eastern region from 73.1% to 100%. For the Central region, the total avertable deaths were 4674.8, with the largest avertable burden from child ART. For the Eastern region, the total

avertable deaths were 17,112.2, with the largest avertable burden from child ART. For the Western region, the total avertable deaths were 21,061.1, with the largest avertable burden from child ART. Across all of SSA, 88.4% of all HIV-related deaths could have been averted if there was complete coverage of biomedical interventions.

	Avertable 0-14 HIV-related deaths			Total	% of estimated deaths
	PMTCT	ART	TMP-SMX		
Central	89.5	4371.1	214.2	4674.8	84.5%
Eastern	2312.5	14453.1	346.6	17112.2	90.5%
Southern	369	4349.3	160	4878.2	91.2%
Western	9252.3	11247	561.8	21061.1	87.3%
Total	26014.5	34417.9	1282.9	61715.3	88.4%

Table 1. Avertable HIV-related deaths by SSA region and intervention between 2000 and 2023

Discussion

The results of this analysis highlight the critical role that targeted medical interventions played in contributing to reduced pediatric HIV burden over the period from 2000 to 2023. An estimated 4.8 million cases and 3.6 million deaths averted in SSA over this period represents an enormous achievement in implementing interventions to better the lives of children. Our estimate of 61,715.3 avertable deaths in 2023 alone points to the need to continue efforts to scale access and coverage of biomedical interventions. From a regional standpoint, Central and Southern SSA exhibit similar opportunities to avert child HIV burden, with the majority of the potential in child ART. In contrast, the Western region has a much larger opportunity to avert deaths through increased PMTCT coverage. This degree of variation indicates that policy makers have the opportunity to tune country-specific plans to address gaps in coverage that are contributing to the greatest level of burden.

In 2023, representatives from twelve high HIV burden African countries signed the Dar-es-Salaam Declaration, laying out their plans for ending AIDS in children globally by 2030 (Lancet HIV Editorial Board, 2023). Achievement of this ambitious goal will require escalated efforts to identify HIV positive pregnant women and ensure that they are enrolled in PMTCT. In the meantime, children living with HIV will need lifelong access to treatment to reduce HIV-related deaths. A crucial aspect of tackling pediatric HIV that this analysis is unable to assess is the role of increasing HIV testing among pregnant women and at-risk children. The pathway to increased coverage relies on improvements in testing, as an individual must first know their status before biomedical interventions become available to them. In addition, retention in care is

crucial, as dropout is a key source of coverage gaps. Further, behavioral and structural factors must be taken into consideration to obtain a holistic view of the drivers of pediatric HIV burden.

A number of limitations affect our interpretation of results and point to opportunities for future analyses. First, the quality of our decomposition estimates relies on the accuracy of the model structure and parameters governing how rates of transmission, disease progression, and mortality are modified by each intervention. Both the model structure and parameters have been continually refined as evidence accumulated and priority model updates were implemented. This process of continuous model improvement indicates that the accuracy of decomposition estimates improves as the pediatric model improves. Similarly, the density of HIV burden data available to calibrate the model for each location impacts the uncertainty intervals surrounding estimates. While this uncertainty is propagated through the decomposition, we can produce more precise decomposition methods when there is more data available to fit to during model calibration. In addition, these estimates rely on country-reported levels of biomedical intervention coverage. The accuracy of these reports is difficult to assess, as there are limited opportunities for validation, and data sources are not always transparent. Future work on addressing data inconsistencies and validating would increase confidence in decomposition estimates.

This work highlights the enormous impact that investment in pediatric HIV biomedical interventions has made in preventing HIV burden in children. Efforts from throughout the global health ecosystem, including scientists, donors, international organizations, local ministries of health, doctors, and community members all contributed to the staggering number of HIV infections and deaths that did not occur due to the administration of effective treatments. Averted burden is a powerful metric that concisely captures the impact of interventions beyond simply reporting the level of disease. Particularly with infectious diseases, intervention impact can be underestimated due to the difficulty in calculating the counterfactual. Only through simulating and decomposing impact is it possible to quantify how pediatric HIV policies have reduced HIV burden in children. Further, simply assessing coverage gaps does not communicate the varying potential impact of closing gaps. Our avertable burden results offer a valuable contribution to future policy formulation, and, hopefully, support the push to end HIV among children in the near future.

Citations

Carter, A., Zhang, M., Tram, K. H., Walters, M. K., Jahagirdar, D., Brewer, E. D., Novotney, A., Lasher, D., Mpolya, E. A., Vongpradith, A., Ma, J., Verma, M., Frank, T. D., He, J., Byrne, S., Lin, C., Dominguez, R.-M. V., Pease, S. A., Comfort, H., ... Kyu, H. (2024). Global, regional, and national burden of HIV/AIDS, 1990–2021, and forecasts to 2050, for 204 countries and territories: The Global Burden of Disease Study 2021. *The Lancet HIV*, 11(12), e807–e822. [https://doi.org/10.1016/S2352-3018\(24\)00212-1](https://doi.org/10.1016/S2352-3018(24)00212-1)

- Lancet HIV Editorial Board. (2023). Declaration commits to ending AIDS in children. *The Lancet HIV*, 10(4), e209. [https://doi.org/10.1016/S2352-3018\(23\)00057-7](https://doi.org/10.1016/S2352-3018(23)00057-7)
- Mahy, M., Penazzato, M., Ciaranello, A., Mofenson, L., Yianoutsos, C. T., Davies, M.-A., & Stover, J. (2017). Improving estimates of children living with HIV from the Spectrum AIDS Impact Model. *AIDS (London, England)*, 31 Suppl 1(Suppl 1), S13–S22. <https://doi.org/10.1097/QAD.0000000000001306>
- Shapley, L. S. (1952). *A Value for n-Person Games* [Product Page]. <https://www.rand.org/pubs/papers/P0295.html>
- Spectrum Manual—Spectrum System of Policy Models*. (n.d.). <https://avenirhealth.org/Download/Spectrum/Manuals/SpectrumManualE.pdf>
- Stover, J., Flanagan, C. F., & Teng, Y. (2023). *The Impact of the President's Emergency Program for AIDS Relief (PEPFAR) on Children* (p. 2023.09.18.23295712). medRxiv. <https://doi.org/10.1101/2023.09.18.23295712>
- Stover, J., Glaubius, R., Kassanje, R., & Dugdale, C. M. (2021). Updates to the Spectrum/AIM model for the UNAIDS 2020 HIV estimates. *Journal of the International AIDS Society*, 24(Suppl 5), e25778. <https://doi.org/10.1002/jia2.25778>
- Stover, J., Walker, N., Grassly, N. C., & Marston, M. (2006). Projecting the demographic impact of AIDS and the number of people in need of treatment: Updates to the Spectrum projection package. *Sexually Transmitted Infections*, 82(suppl 3), iii45–iii50. <https://doi.org/10.1136/sti.2006.020172>
- Updates to the Spectrum/AIM model for the UNAIDS 2020 HIV estimates—Stover—2021—Journal of the International AIDS Society—Wiley Online Library*. (n.d.). Retrieved May 18, 2024, from <https://onlinelibrary-wiley-com.offcampus.lib.washington.edu/doi/full/10.1002/jia2.25778>
- World Health Organization. (2006). Guidelines on co-trimoxazole prophylaxis for HIV-related infections among children, adolescents and adults in resource-limited settings: Recommendations for a public health approach. *Directives Sur l'utilisation Du Cotrimoxazole Pour La Prophylaxie Des Infections Liées Au VIH Chez l'enfant, l'adolescent et l'adulte : Recommandations Pour Une Approche de Santé Publique*, 64.
- World Health Organization. (2009). Rapid advice: Use of antiretroviral drugs for treating pregnant women and preventing HIV infection in infants, version 2. *Recommandations Rapides : Médicaments Antirétroviraux Pour Traiter La Femme Enceinte et Prévenir l'infection à VIH Chez l'enfant, Version 2*, 23.
- World Health Organization. (2013). *Consolidated guidelines on the use of antiretroviral drugs for treating and preventing HIV infection: Recommendations for a public health approach*. World Health Organization. <https://iris.who.int/handle/10665/85321>

Aim 3: Uganda Malaria Outbreaks: Using Averted Burden to Evaluate Algorithms for Outbreak Detection and Response

Introduction

While outbreaks are a core concept in the study of infectious diseases, there is not a canonical definition for what constitutes an outbreak that applies to all pathogens or contexts (O'Neil & Naumova, 2007). For some pathogens and contexts, a single case warrants declaration of an outbreak, like the recent Ebola outbreak in Uganda (Nakkazi, 2025). In contrast, endemic infectious diseases - which have persistent local transmission - pose a particular challenge for outbreak identification. According to the World Health Organization, a disease outbreak is an above-normal number of cases (World Health Organization, 2018). While this definition aligns with our intuitive understanding of the concept of an outbreak, it is not rigorous enough to operationalize. What is the normally expected number of cases? What constitutes an excess number? How can it be computed by disease control programs to trigger a response?

Many outbreak detection algorithms have been proposed and implemented over time, across locations, and for a range of pathogens (Brady et al., 2015). The WHO has proposed a range of algorithms that could be useful for a given context, accounting for challenges like seasonality, poor data quality, and natural variation in transmission (World Health Organization, 2018). Previous efforts to evaluate the outbreak detection algorithms look across these multiple options and [insert conclusions from paper comparing approaches]. In Uganda, a recent comparison of approaches to detecting malaria outbreaks across districts in the country reported on the frequency with which outbreaks occurred according to different outbreak definitions, but it stopped short of determining which approach is best for each location in Uganda (Zalwango et al., 2024). Frequency of outbreak detection falls short of providing sufficient information to determine the strength of an outbreak detection approach, as a high frequency of outbreaks detected could be either bad or good, depending on the capacity to respond and the impact of a response. We seek to build upon that work by proposing a framework for differentiating the performances of various approaches to outbreak detection that are intended to trigger an outbreak response when thresholds are met.

To identify a metric for comparing algorithms, we turn to incremental cost effectiveness ratios (ICERs), which are used in health economics to compare hypothetical interventions that could be added to a policy (Weinstein & Stason, 1977). An ICER is the ratio of the marginal cost of an intervention to the averted burden resulting from the intervention. This metric is useful for pursuing optimal resource allocation during policy formulation. Here, we propose calculating the hypothetical ICERs associated with different outbreak detection algorithms applied to historical data as a means of differentiating which algorithms would have performed better and therefore may be the best option to use to detect outbreaks in the future.

To demonstrate the value of this framework, we retrospectively apply two outbreak detection algorithms to health facility reported malaria case data from three districts in Uganda. Through simulation of the hypothetical impact of a response delivered in response to detection of an outbreak, we estimated the averted burden resulting from each response and compared the cost-effectiveness of the responses triggered by each algorithm. This manuscript provides a clear framework through which countries can tune their choice outbreak detection approach to maximize impact given variable capacity to respond and temporal and geographic variation in transmission.

Methods

Data

Monthly reports from health facilities in Uganda include counts of individuals testing positive for malaria. We maintain a database of all historical report values, dating back to 2015, and update the database as new reports are submitted. Our team developed a set of data transformations that are applied to clean the extracted data, including excluding outliers and imputing missingness at the facility level. Outliers were identified through a two-step screening that first removed large outliers failing more than five standard deviations outside of the distribution of the facility data and then used a seasonal time series outlier function available in the *forecast* R package. Imputation was also carried out using the *tsclean* function from the *forecast* package. This analysis utilizes both the original and the cleaned data, aggregated to a total number of monthly malaria cases at the district level.

Outbreak detection algorithms

The WHO recommends an outbreak detection algorithm that sets a threshold for outbreak at the 75th quantile across past data from the same month across all years. This accounts for seasonal variation by ignoring data from other months in the past, dramatically reducing the

number of observations included in quantile calculations. The WHO algorithm does not provide explicit data cleaning instructions, so we apply the algorithm to the original data.

We developed an outbreak detection approach that establishes a baseline by estimating a district-specific canonical seasonal pattern and applying it to the median historical value. We then apply a kernel smooth with a Gaussian kernel and a bandwidth of two months to the cleaned case data and then take the ratio of the resulting smoothed time-series to baseline to estimate an outbreak index, against which an outbreak threshold can be established. For this analysis, we use an outbreak index threshold of two, which equates to cases being at least twice as large as the baseline, to determine when a location has transitioned into outbreak.

Model implementation and calibration

To assess the impact of outbreak response in terms of variation in malaria cases, we implemented a simple model of malaria transmission modified by treatment and transmission-suppressing interventions. The model is a discrete-time state space model with three states: Susceptible, Infected, and Treated. Our time step is ten days, and to account for the thirty days of protection provided by treatment, we split the treated class into two and have all newly treated individuals progress from the first treated state to the second treated state, and then all individuals in the second treated state progress to the susceptible state, resulting in a minimum of thirty days between infections for treated individuals. Transition from the susceptible to the infected state is governed by a force of infection parameter which is calibrated to location specific data. The recovery rate in the absence treated is uniform across locations and aligns with an expected length of infection of 200 days. Treatment rates are set to 70% across all locations, aligning with estimates from the Malaria Atlas Project that are based on survey data. We use a fourth order B-spline as the functional form for our force of infection time series, optimizing the knot values or control points during model calibration. Spline knots are placed every two months spanning the whole period of observation.

We implement our malaria model using `torch`, an R package which serves as an interface to PyTorch, an open-source machine learning framework that executes automatic differentiation to accelerate model fitting. For our loss function, we use the root mean squared error between the simulated incidence and the observed incidence – facility reported cases divided by population – divided by 0.4, the proportion of individuals who seek care at a public facility. We use gradient descent to minimize the error of our calibrated model, stopping training when our loss value changes by less than $1e-6$ across iterations.

Simulating outbreak detection and response

To emulate the experience of a program applying an outbreak detection algorithm in real time, we simulated the application of outbreak detection while stepping forward in time, month by month, starting in the beginning of 2018. At each time point, we used only data available up to that time point to establish a baseline and assess outbreak. We then took the periods identified as a transition into outbreak by each algorithm and simulated a response in the next time-step using our calibrated model. Our simulated outbreak response consisted of mass drug administration reaching 95% of the population, and transmission suppressing interventions that decreased the transmission rate by 50% for six months. We waited at least six months between outbreak responses, even if a location transitioned into and out of an outbreak multiple times in a six-month time interval. For construction of ICERs, we assume that the cost to respond is proportional to the population of a district and the total case load in the district. These costs are hypothetical and should be replaced with estimates of the real costs when calibrating outbreak detection

Results

Data outliering and imputation were carried out at the facility-level and both imputed data and reported data were aggregated to the district level, with the imputed data differing from the reported data at multiple time points (Figure 1). Outbreak thresholds were computed for both outbreak detection approaches, and data were compared to the thresholds to determine outbreaks detected by each approach. We analyzed all 146 districts, but we present and discuss results from three districts in some detail that illustrate the variable performance of each outbreak detection approach.

For the Bundibugyo District, the application of the RAMP algorithm resulted in zero outbreaks detected, while the 75th percentile approach detected 10 distinct outbreaks (Figure 1, Table 1). In Rukiga District, there was one RAMP outbreak detected and two 75th percentile outbreaks detected. For Tororo District, the RAMP outbreak approach detected one outbreak while the 75th percentile approach detected only one outbreak.

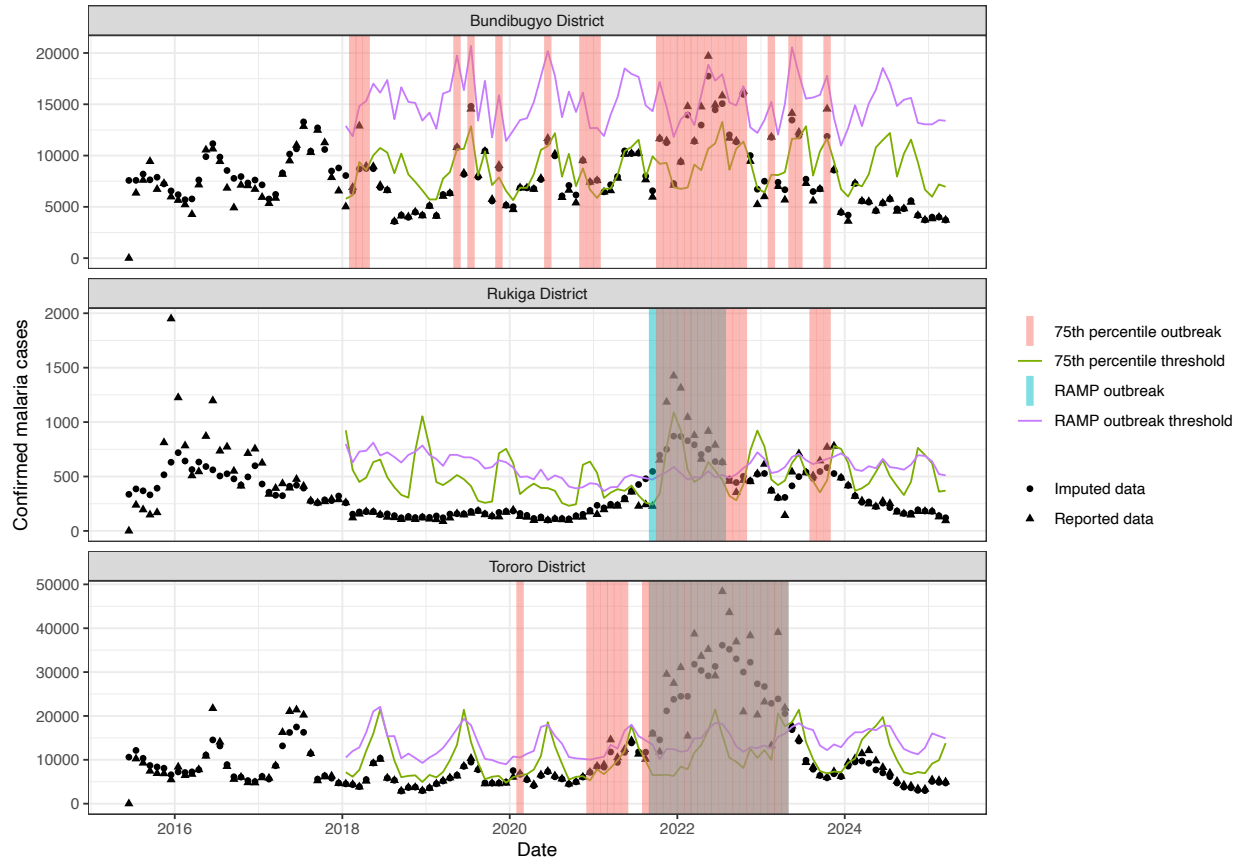


Figure 1: Detected outbreaks by method across three districts.

Periods of outbreak were matched with our simulation calibrated to location-specific data, and counterfactual outbreak responses were calculated (Figure 2). For each outbreak response, we subtracted the total number of cases in the counterfactual response from the cases that were estimated with the observed data to determine the number of cases averted through outbreak response.

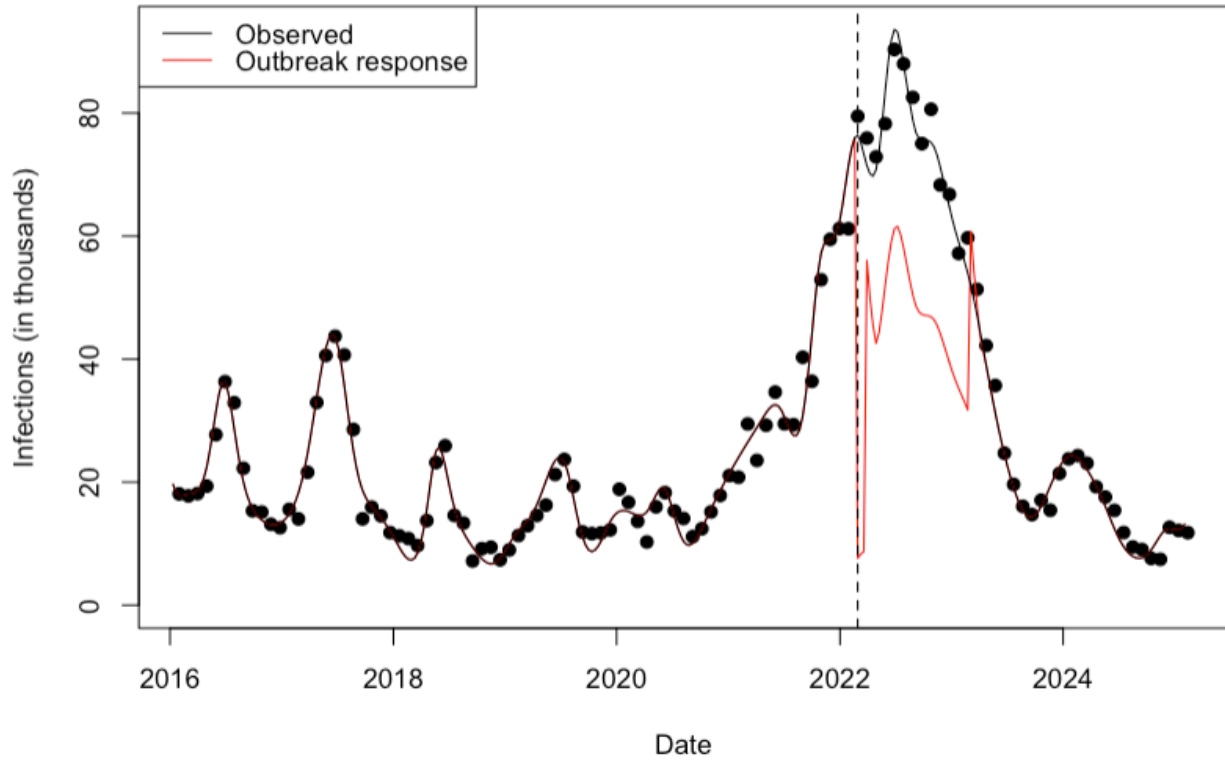


Figure 2: Outbreak response vs observed in example district

For each district, we calculated the impact of outbreak response in terms of average averted cases across all outbreak responses triggered by each respective method in the district. We then compared the outbreak response impact to the median response impact for an outbreak response delivered at any time point across the whole period being assessed (Figure 3).




Figure 3: Calibrated model with data for Tororo District

In Bundibugyo, the average outbreak response impact for the 75th percentile approach was 337.9 thousand cases averted which was less impactful than a response delivered randomly, with a ratio of 0.98 (Table 1). The RAMP outbreak approach did not detect any outbreaks, so there were no averted cases arising from RAMP outbreak response. In Rukiga, the average outbreak response impact for the RAMP approach was 32.6 thousand cases averted, which was 2.20 times more impact than a randomly delivered response. For the 75th percentile approach, the average outbreak response impact was 26.6 thousand cases averted, which equated to a relative response impact of 1.79. In Tororo, the RAMP outbreak response led to 1037.9 thousand cases averted with a relative response impact of 2.78. The 75th percentile approach led to 665.2 thousand cases averted and relative response impact of 1.78.

		District		
		Bundibugyo	Rukiga	Tororo
Number of outbreaks	RAMP	0	1	1
	75th percentile	10	2	3

Response impact*	RAMP	NA	32.6	1037.9
	75th percentile	337.9	26.6	665.2
Relative response impact**	RAMP	NA	2.20	2.78
	75th percentile	0.98	1.79	1.78

Table 1: Number of outbreaks detected and outbreak response impact for outbreaks detected by each method across districts

* Cases averted by outbreak response averaged over all outbreaks in the district in thousands

** Ratio of the average response impact to the median impact of response across all time periods

Discussion

These results demonstrate the value of a clear framework for characterizing the effectiveness of responses triggered by different outbreak detection approaches. With these tools for evaluating how dynamic policies would have performed in the past, national malaria control programs can calibrate the use of these dynamic policies in supplementing the policies that are implemented over longer timescales and that do not change according to the observed level of disease in a location.

Each set of district results highlights a unique outbreak detection situation. In Bundibugyo, the data is highly variable, which leads to the 75th percentile approach frequently triggering ineffective outbreak responses. The RAMP outbreak detection approach is more robust to these variation due to the use of smoothing and data cleaning and imputation. In Rukiga, the malaria burden is relatively small, so the impact of outbreak response not as impactful as in Tororo, where the burden is very high. Nonetheless, both outbreak approaches trigger outbreak responses that are more impactful than a randomly delivered response. In Tororo, the 75th percentile approach triggers two less effective outbreak responses, while both approaches detect the large outbreak that occurred due to failure of IRS (Epstein et al., 2022).

The outbreak detection algorithm developed by our team for malaria in Uganda exhibits several positive traits that we believe translated to the higher effectiveness of response when triggered by our outbreak detection algorithm in contrast to the 75th quantile approach. The first is facility-level data cleaning and imputation. Outliers and missingness can plague outbreak detection, as heterogeneity in reporting frequency or large outliers can deceive an algorithm. Executing this data cleaning at the facility-level is critical, because aggregation to an administrative region can mask many of the data issues and make them appear like real signal.

The second is the use of bandwidth smoothing to execute partial pooling across time. This is not a novel concept in outbreak detection algorithms, but our implementation allows for the bandwidth used to vary according to context, which creates a degree of freedom for tuning outbreak detection to meet program needs. For example, in instances where there is lots of heterogeneity over time in transmission, a longer bandwidth may be preferred so that longer run trends are used to trigger response. In contrast, areas of low transmission may prefer a much shorter bandwidth, which is equivalent to being more sensitive to fluctuations in the data on a shorter time interval. This partial pooling also has the benefit of increasing the effective sample size informing the baseline level of cases relative to the 75th percentile approach, which leverages one data point from every year of observation.

The example analysis of outbreak detection in Uganda highlights how the transition from evaluating algorithms through the lens of outbreak frequency to response effectiveness provides an objective measure to improve the application of outbreak detection to the goal of optimal resource allocation during policy formulation. A primary limitation of this analysis is that the outbreak response timing and intervention components that were delivered in simulation are hypothetical, not necessarily what the program would actually do in the case of an outbreak. Future work carried out in partnership with malaria control programs should investigate how to stratify outbreak detection algorithms to address the variety of contexts present within a country. In addition, inclusion of real outbreak response costs should allow for the dynamic policy of outbreak response to be compared directly with the ICERs for static policies like delivery of LLINs and IRS. We believe that funding dynamic policies is critical for addressing a dynamic pathogen like malaria, and improved frameworks for directly evaluating the cost-effectiveness of outbreak response should bolster the case for funding these activities when developing budgets and strategic plans.

Citations

- Brady, O. J., Smith, D. L., Scott, T. W., & Hay, S. I. (2015). Dengue disease outbreak definitions are implicitly variable. *Epidemics*, *11*, 92–102. <https://doi.org/10.1016/j.epidem.2015.03.002>
- Epstein, A., Maiteki-Sebuguzi, C., Namuganga, J. F., Nankabirwa, J. I., Gonahasa, S., Opigo, J., Staedke, S. G., Rutazaana, D., Arinaitwe, E., Kanya, M. R., Bhatt, S., Rodríguez-Barraquer, I., Greenhouse, B., Donnelly, M. J., & Dorsey, G. (2022). Resurgence of malaria in Uganda despite

sustained indoor residual spraying and repeated long lasting insecticidal net distributions. *PLOS Global Public Health*, 2(9), e0000676. <https://doi.org/10.1371/journal.pgph.0000676>

Nakkazi, E. (2025). Sudan Ebola virus disease outbreak in Uganda. *The Lancet Infectious Diseases*, 25(4), e206. [https://doi.org/10.1016/S1473-3099\(25\)00174-4](https://doi.org/10.1016/S1473-3099(25)00174-4)

O'Neil, E. A., & Naumova, E. N. (2007). Defining Outbreak: Breaking Out of Confusion. *Journal of Public Health Policy*, 28(4), 442–455. <https://doi.org/10.1057/palgrave.jphp.3200140>

Weinstein, M. C., & Stason, W. B. (1977). Foundations of Cost-Effectiveness Analysis for Health and Medical Practices. *New England Journal of Medicine*, 296(13), 716–721. <https://doi.org/10.1056/NEJM197703312961304>

World Health Organization. (2018). *Malaria surveillance, monitoring and evaluation: A reference manual*. World Health Organization. <https://iris.who.int/handle/10665/272284>

Zalwango, M. G., Zalwango, J. F., Kadobera, D., Bulage, L., Nanziri, C., Migisha, R., Agaba, B. B., Kwesiga, B., Opigo, J., Ario, A. R., & Harris, J. R. (2024). Evaluation of malaria outbreak detection methods, Uganda, 2022. *Malaria Journal*, 23, 18. <https://doi.org/10.1186/s12936-024-04838-w>