Title

Better individual-level risk models can improve the targeting and life-saving potential of early-mortality interventions.

Permalink

https://escholarship.org/uc/item/6cp7t6wt

Journal

Scientific Reports, 13(1)

Authors

Hazlett, Chad
Ramos, Antonio
Smith, Stephen

Publication Date

2023-12-07

DOI

10.1038/s41598-023-48888-7

Peer reviewed

scientific reports

OPEN

Better individual-level risk models can improve the targeting and life-saving potential of early-mortality interventions

Chad Hazlett[1,2], Antonio P. Ramos[3,4] & Stephen Smith[2]

Infant mortality remains high and uneven in much of sub-Saharan Africa. Even low-cost, highly effective therapies can only save lives in proportion to how successfully they can be targeted to those children who, absent the treatment, would have died. This places great value on maximizing the accuracy of any targeting or means-testing algorithm. Yet, the interventions that countries deploy in hopes of reducing mortality are often targeted based on simple models of wealth or income or a few additional variables. Examining 22 countries in sub-Saharan Africa, we illustrate the use of flexible (machine learning) risk models employing up to 25 generally available pre-birth variables from the Demographic and Health Surveys. Using these models, we construct risk scores such that the 10 percent of the population at highest risk account for 15-30 percent of infant mortality, depending on the country. Successful targeting in these models turned on several variables other than wealth, while models that employ only wealth data perform little or no better than chance. Consequently, employing such data and models to predict high-risk births in the countries studied flexibly could substantially improve the targeting and thus the life-saving potential of existing interventions.

Goal 3.2 of the Sustainable Development Goals (SDG) seeks to "end preventable deaths of newborns and children under five years of age, with all countries aiming to reduce neonatal mortality to at least as low as 12 per 1000 live births and under-5 mortality to at least as low as 25 per 1,000 live births." Despite large overall improvements, including a 44% reduction in child mortality globally from 2000 to 2015, progress has varied widely both between and within countries[1–3], with under-5 mortality exceeding 80 per 1,000 live births in some countries of sub-Saharan Africa. Further, improvements in national averages can mask large and widening gaps between sub-populations, with more privileged groups improving their health status faster than others[4,5].

Many of these deaths are preventable as they can be avoided with low-cost interventions such as oral rehydration therapy, malaria prevention and treatment, vaccines (e.g., against measles), and nutritional support[6,7]. However, a critical feature of any life-saving intervention is that it can only save the lives of those who would have otherwise died. Even a "miracle drug" that can counteract any cause of death can only reduce mortality when given to children who would have died without it. Because of this, interventions that cannot be given universally must be carefully targeted to those at the highest risk of mortality (absent the intervention) to have an effect. While this point may seem obvious, in the context of statistically rare events, it has enormous implications because it can severely limit any intervention's theoretical maximum life-saving potential. Consider a country with an infant mortality rate of 80 deaths per 1,000 births. Normatively, this is far too high, yet this rate of 0.8% means that intervention given to a random set of individuals can save at most 0.8% of those to whom it is given.

Worse, in practice, such an intervention will only save the lives of a fraction of those who would have otherwise died. We call this fraction the efficacy rate (it is sometimes known as the "probability of necessity and sufficiency" in the causal inference literature). Suppose the efficacy rate is 10%, meaning the intervention would save the lives of one-tenth of those who would have otherwise died. An intervention given at random could save the lives of only 0.8% * 10% = 0.08% of those given to whom it is given. By contrast, if such an intervention with the same efficacy rate could instead be targeted to a group with a higher risk of mortality, say 10%, then it

[1]Professor, Department of Political Science and Department of Statistics and Data Science, University of California Los Angeles, Los Angeles, USA. [2]PhD Candidate, Department of Statistics and Data Science, University of California Los Angeles, Los Angeles, USA. [3]Research Scientist, Califonia Center for Population Research, University of California, Los Angeles, USA. [4]Principal Research Scientist, José Luiz Egydio Setúbal Foundation, São Paulo, Brazil. ✉email: tomramos@g.ucla.edu

would save the lives of 10% * 10% = 1% of recipients, saving 12.5 times as many lives. (We thank an anonymous reviewer for noting that such an analysis neglects possible spillovers, such as when saving one child's life could increase a nearby child's chance of survival. Such spillovers may be reasonably likely, for example, in the case of infectious diseases. Computing the net effect of an intervention under such spillover would require additional structure and assumptions on the problem. It is noteworthy, however, that because early mortality is statistically rare, the effect of such spillover effects on the net lives saved can only be non-negligible where that spillover occurs through clusters of births that share very high-risk levels. This is possible if, indeed risk levels might be highly concentrated (e.g., in particular villages). In this case, effective interventions well targeted to such high risk areas would enjoy an even greater life-saving benefit due to this spillover). In this way, the lives saved by a hypothetical intervention with a given efficacy rate are limited by and proportional to the mortality risk in the group to which it is given. Efforts to improve targeting should thus be treated with the same urgency as efforts to improve the efficacy of the interventions. Accordingly, our paper seeks to demonstrate the use of tools for better estimating which births are at the highest risk of early mortality.

Unfortunately, the importance of targeting has not been well recognized by recent and current public health practices in many cases. Although early mortality can have many causes and can be related to many risk factors, countries facing high early mortality often target related interventions based on one or a few factors, including income, wealth, or poverty, and employing simple statistical models or checklists. Similarly, academic research into models for early mortality that look beyond income, poverty, or wealth remains rare and not designed to demonstrate how practical, national-level systems could target interventions to high-risk births. A recent, partial exception finds, first, that poverty is a poor proxy for the risk of early mortality in Brazil and India, much as we find below for countries in sub-Saharan Africa[8]. Further, for births in India, they find that using just four risk factors in a linear model roughly doubles the fraction of infant deaths that can be identified in the highest risk quintile compared to models relying on a proxy for poverty.

We demonstrate the potential for tools that better estimate which births are at the highest risk of early mortality, using flexibly estimated, individual-level risk scores based only on a range of commonly available pre-birth data, which can improve the targeting of life-saving interventions against infant mortality. Specifically, we estimate mortality risk in individual infants (under one year of age) in 22 countries in sub-Saharan Africa, employing now-standard machine learning models to identify predictive relationships that may have eluded more rigid regression models. This flexibility is especially valuable as we provide these models access to a wider range of risk factors than in previous work. We also employ only risk factors that can generally be known or anticipated before a child's birth, as practical targeting systems may be limited to such data. As is standard in many machine learning applications, we employ cross-validation and held-out test sets to protect against overly optimistic performance claims due to over-fitting. However, rather than employing the mathematically complicated, non-intuitive performance metrics commonly employed in the forecasting literature, we rely on a simpler set of metrics designed to be easily interpreted by policymakers and to directly speak to the potential impact of a hypothetical policy intervention.

Materials and methods
Data sources
We employ data from the Demographic and Health Surveys (DHS), as augmented and prepared for analysis by the Integrated Public Use Microdata Series (IPUMS), henceforth referred to as the IPUMS DHS data. Each survey we used was the most recent one available via IPUMS, when our study began. Like other work studying early mortality using the IPUMS DHS data (e.g.,[9,10]), we use the "child module" which contains data on births for children born up to five years before the interview of their mother. Though interviews are conducted with mothers, one observation in the data represents each child's birth. This module restricts our attention to children under five, but contains more predictive variables than the "all births" module. More importantly, it would be inadvisable to use data collected on the mother, household, malaria prevalence, location, etc., in the year of the DHS data collection to predict early mortality from birth many years before.

Originally, the IPUMS DHS data contained surveys on 28 sub-Saharan African countries. We consider geographic information to be critical for these models since (i) even coarse geographic indicators can be highly predictive of risk and reduce the need for other, harder-to-measure variables associated with location, (ii) we considered malaria prevalence a potentially important variable, but because it is computed for a 10km area around the DHS sample cluster centroid, it is only available in surveys that collected geographic coordinates, and (iii) the ability to target interventions partly by geography is convenient for policymakers. Unfortunately, in six countries (Burundi, Chad, Ethiopia, Liberia, Namibia, and Sudan), the IPUMS DHS data did not include geographic indicators even at the district level. This limits us at present to surveys from 22 countries: Angola (2015), Benin (2011), Burkina Faso (2010), Cameroon 2011, Democratic Republic of Congo (2013), Cote d'Ivoire (2011), Ghana (2014), Guinea (2012), Kenya (2014), Lesotho (2014), Madagascar (2008), Malawi (2016), Mali (2012), Mozambique (2011), Niger (2012), Nigeria (2013), Rwanda (2014), Senegal (2017), Tanzania (2015), Uganda (2016), Zambia (2013), Zimbabwe (2015). Table 1 provides additional information on the available data within each country.

Non-geographic variables
Our outcome is infant mortality, indicated by death before reaching one year of age. For predictors or risk factors (we use the terms interchangeably) we limit ourselves to predictors that (i) would be reasonably available to policymakers and health workers in-country, and (ii) are measurable prior to the birth of the child, rather than depending on the health or other features of the child once borne. The latter constraint is in place so that

Country	N original	N complete	N deaths	Geography	Unavailable
Angola	11024	10829	472	Province (18)	Age first marriage, malaria
Benin	10490	10199	471	Department (12)	
Burkina Faso	11728	10905	752	Region (13)	
Cameroon	9041	8450	564	Region or city (12)	
Congo (D.R.)	14384	12731	785	Province (11)	Previous child death(s)
Cote d'Ivoire	5945	5211	421	None	
Ghana	4548	4203	189	None	
Guinea	5422	5223	396	Region (8)	Cooking fuel
Kenya	16557	15110	589	None	Previous pregnancy ended
Lesotho	2360	2127	145	District (10)	Bednets, malaria
Madagascar	9657	9229	459	Region (22)	Cooking fuel
Malawi	13633	13231	563	Region (3)	Malaria
Mali	8043	7913	475	Region (6)	Cooking fuel, malaria
Mozambique	8493	8026	537	Province (11)	
Niger	9650	9612	586	Region (8)	Religion, malaria
Nigeria	24222	23559	1686	Geopolitical zone (6)	
Rwanda	6055	5350	168	Province (5)	Cooking fuel
Senegal	9585	9371	401	Region (14)	Malaria
Tanzania	8007	7688	319	District (10)	Cooking fuel, religion, malaria
Uganda	12089	11619	484	District (112)	Cooking fuel, malaria
Zambia	10646	9811	483	None	
Zimbabwe	4852	4630	236	Province (10)	Malaria

Table 1. Data availability by country. *Note*: Summary of data by country. The DHS data initially include `N original` observations in each country, which falls to `N complete` when we remove missing values on the included variables. `Geography` indicates what type of geographic information, if any, is available and the number of units in parentheses. All other variables described in text are used in each country except where otherwise indicated under `Unavailable`.

these risk factors could be used to calculate risk in advance and deploy interventions to households, clinics, or regions with high risk.

The resulting risk factors we consider are: sibling death at an age less than 1 yr., sibling death older than 1 year, mother education, birth month (sin and cosine), sex (female), a rural indicator, local malaria prevalence, maternal age, previous pregnancy terminated, clean water access, birth order (log), mother's age at first union, number of bednets (log(+1)), floor type, age of head of household, female household head, minority religion, christian, or muslim, unimproved toilet, pit toilet, clean cooking fuel, and within-country wealth percentile (and its square). The source data from which these are constructed are described in Table 3 of the Supplement.

Note that the IPUMS DHS data do not indicate how an individual or household is classified relative to the poverty line in a given country, nor does it attempt to measure income as such. Thus, rather than including any locally defined poverty measure, we employ the wealth-based indices provided by DHS. This measure is an index for within-country wealth, imputed by DHS using principal components analysis over several variables, including ownership of radios, televisions, and other domestic equipment; electricity and clean water; type of materials used in the walls, floor, and roof; and the type of toilet in the household. This score is calculated at the household level. The only requirement we make of this measure is that, within a given country, it provides an imperfect but potentially informative measure of relative wealth. Since this index is constructed within each country, and we train each country's risk model separately, there is no expectation or need for cross-country comparability.

Modeling

We estimate infant mortality risk for each birth in the data set. In each country, we separately run two sets of models. One is a "wealth model", described below, which reveals what can be learned from the available wealth data alone. The other set contains flexible machine-learning approaches that use the predictors described above to estimate individual-level mortality risk. These include random forest (rf), extreme gradient boosting over trees (xgb), kernel regularized least squares (krls), and elastic-net regularized logistic regression (elastic-logit).

Both rf[11] and xgb[12] are tree-based approaches. The simplest tree-based approaches are classification and regression trees (cart[13]), which recursively partition observations into groups based on one cutoff value of the predictor variables so that each group has similar values of the outcome. While such single-tree models are easy to interpret, they often perform poorly, and small input data changes can dramatically affect the model and results. The rf and xgb methods improve on single-tree approaches by combining many trees. For rf, this is achieved by effectively averaging the estimates of many trees, each employing only a random subset of the input variables and trained on a random subset of the observations. As applied here, each forest contains 500 trees, and the number of variables randomly sampled for each tree split (`mtry`) was set for each country's model using cross-validation.

Under xgb, this is similarly achieved by combining many trees, but by beginning with one tree, adding the predictions of another tree trained on the errors of the prior, and then repeating this process hundreds or thousands of times. We use cross-validation for each country to determine (i) the learning rate ($\eta \in (0, 1)$, where smaller values help to prevent overfitting), (ii) the maximum depth of the trees (`max_depth`), (iii) the proportion of the variables to be considered for tree construction (`colsample_bytree`), (iv) the proportion of observations from the training set used for modeling (`subsample`), and (v) the number of iterations for the boosting procedure (`nrounds`). The remaining parameters were set at common default values: the minimum requirement of prediction improvement before selecting a more complex model (γ) was set to 0, and `min_child_weight`, where higher values restrict the depth of each tree based on a measure of the homogeneity of the labels within the nodes (a sign of overfitting), was set to 1.

Kernel-based regularized least squares (krls[14]) is a regression and classification technique that employs kernels. Such approaches effectively build a model by leveraging information about similar observations. Each observation is treated as a potential "fence post" in the data, and the model learns how being closer to or further from each such location in the data influences the expected probability of mortality. We employ a Gaussian kernel as the measure of similarity using default values[14].

Elastic-net logistic regression (elastic-logit[15]) is similar to conventional logistic regression. While ordinary logistic regression chooses coefficients β_j on each variable j to maximize the likelihood by minimizing the negative log-likelihood, elastic-net logistic regression minimizes the negative log-likelihood *plus* a penalty term that helps to regularize or "shrink" the model, preventing overfitting. In our case that penalty term is equal to $\lambda(\frac{\alpha}{2}\sum_j |\beta_j| + \frac{1-\alpha}{2}\sum_j \beta_j^2)$. The parameters α and λ are chosen by cross-validation. This constitutes a compromise between what is known as ridge regression (which is penalized by the ℓ_2 norm, $\sum_j \beta_j^2$), and LASSO (penalized by the ℓ_1 norm, $\sum_j |\beta_j|$). By including the ℓ_1 norm in the penalty, elastic-net can set some coefficients to exactly 0, that is, selecting some variables while dropping others.

Finally, we construct a simple ensemble model in which each predicted probability is the unweighted average of the probabilities produced by elastic-logit, rf, xgb, and krls. This has two primary benefits. First, such ensembles often perform as well as the best sub-models and sometimes better. Second, it allows us to commit to reporting results for a single model—the ensemble model—as our premier estimate of the predictive power of these models on the test data. This is preferable to reporting the best-performing model on the test set, which can be misleading because such a choice effectively means "training on the test set".

Wealth model

To determine how much information is available in the wealth data alone and to provide a baseline performance level for comparison, we also model mortality using only wealth data. To allow a flexible model for wealth, we regress mortality on indicators for wealth, equivalent to finding the mean mortality level in each wealth decile. In evaluating model performance, we can consider how being in the highest-risk decile, quartile, or other groups is related to mortality risk. There is no presumption that risk is necessarily decreasing or even monotonic in wealth.

We note that this is a wealth model, not a poverty model, based on poverty assessments made by countries, usually based on income. Neither income nor poverty assessments themselves are available in the DHS data. This model therefore is not a direct proxy for poverty-line-based means-testing sometimes employed. That said, our wealth model may reasonably provide an optimistic indication of what can be predicted by income or poverty-based approaches. Household income is time-sensitive but is likely to generate risks of early mortality when it remains low on average over a long period, in which case wealth will also drop. Moreover, financial shocks such as large expenses or losses may have little impact on income but substantially on the quality of life, access to healthcare, and poverty as locally understood (see e.g.[16]). Our wealth model is also more flexible and less coarse than poverty-based targeting systems, which typically assign individuals through a binary system (e.g. above or below the poverty line). Again this is intended to be a favorable choice for the wealth model so that we can later consider its demonstrated performance likely to be over- rather than under-stated.

Sample splitting and cross-validation

Before fitting any model, we split the data into training and test sets. For each country, 80% of observations are placed in a training and validation set, while 20% are retained in a held-out testing set. For the models described above, we tune hyperparameters using either 10-fold cross-validation within the training set (for elastic-logit, rf, xgb) or leave-one-out cross-validation (for krls). The cross-validation procedure also gives a preliminary assessment of model performance. Once all modeling choices were finalized based on this cross-validation, those models were "frozen". We completed a paper draft using these cross-validation results to ensure all decisions were finalized. Only then did we run the model on the test set to obtain an honest out-of-sample performance assessment without "training on the test set" by revisiting any decisions.

Variable importance

Our models only seek to determine the best risk estimate possible given a set of predictors, and the answers say nothing of the causal effect of those predictors on the risk of mortality. Indeed, many risk factors, such as religion or the type of roof, are valuable not because they cause mortality but because they are expected to signal the presence of other, unobserved factors that cause higher mortality. In other words, confounding is expected and hoped for in these models.

It is valuable, nevertheless to understand which variables the models deem most important in estimating the risk to assess the model's credibility, improve future models, or focusing public health resources on data collection. Unfortunately, for many machine learning models, it is more difficult to characterize the model's behavior than with simpler regression approaches like OLS or logit, in which a small number of parameters fully

characterizes the model. That said, readily available measures of variable importance for different models can provide useful insights. Specifically, for elastic-net, we report how often the model chose a given variable and the proportion of time their coefficient is positive. For random forest, the (scaled) "variable importance" describes what fraction of the time a given variable was chosen to be included in the classification trees over which each random forest model aggregates.

Performance metrics

Finally, we employ simple, policy-relevant metrics rather than more esoteric and difficult-to-interpret metrics often relied upon in machine learning tasks to evaluate our approach. A model's *recall* (also referred to as sensitivity) is the proportion of actual deaths included in the group we consider to be at risk. What we term "recall10" is defined by first determining the group of births in the top risk decile (hence the "10") according to our model, then computing the fraction of all actual infant deaths accounted for by this group. In policy terms, this corresponds to a case in which we have resources to target 10% of the population and wish to know how much of the total mortality risk would be covered in that group. An ineffective risk assessment that randomizes risk would produce a recall10 of 0.10 (10%). By contrast, an effective system points us to a top risk decile that accounts for well over 10% of mortality. Similarly, we report recall20, which tells us the fraction of early deaths within the top risk quintile. For recall10 and recall20, we also report "efficiency gain", which tells us how many times more effective the ensemble model is than the benchmark wealth model. We consider hypothetical scenarios where it would be difficult for health ministries to target an intervention to more than 10% or 20% of the birth population due to budget constraints and practical exigencies. Finally, note that because we are interested in how these models would perform when applied across the population of births in each country, these recall measures are constructed using sampling weights (v005).

Results
Performance of the wealth model

Table 2 and 3 provide the primary results, showing the proportion of mortality among the top 10% (Tables 2) or 20% (Table 3) highest predicted risk individuals in the test set. The online supplement provides the same performance metrics but using cross-validated training set results. Those results are extremely similar to the test set results reported here, indicating little over-fitting in the training process.

Focusing first on wealth as a predictor, on the one hand, the average recall10 of the wealth model is only 0.12—slightly better than chance. In 15 of 22 countries, the recall10 is no greater than .12, and in some countries, it is very low (e.g., just 0.03 in Tanzania). On the other hand, there are two countries with recall10 of 0.16 (Benin and Kenya), and two with very promising recall10 values over 0.20 (Uganda, Rwanda).

Survey	Wealth	Elastic-logit	rf	xgb	krls	Ensemble	Efficiency gain
Angola 2015	0.12	0.22	0.22	0.19	0.19	0.22	1.91
Benin 2011	0.16	0.23	0.31	0.20	0.22	0.23	1.47
Burkina Faso 2010	0.12	0.15	0.21	0.16	0.15	0.20	1.67
Cameroon 2011	0.07	0.16	0.23	0.15	0.15	0.22	3.13
Congo (D.R.) 2013-14	0.14	0.19	0.19	0.19	0.16	0.20	1.41
Cote d'Ivoire 2011	0.12	0.29	0.18	0.29	0.27	0.26	2.20
Ghana 2014	0.14	0.11	0.22	0.16	0.08	0.19	1.40
Guinea 2012	0.11	0.14	0.22	0.15	0.13	0.19	1.67
Kenya 2014	0.16	0.21	0.16	0.12	0.19	0.20	1.21
Lesotho 2014	0.14	0.17	0.14	0.14	0.10	0.17	1.25
Madagascar 2008	0.09	0.10	0.15	0.10	0.13	0.21	2.38
Malawi 2016	0.10	0.11	0.20	0.19	0.14	0.17	1.73
Mali 2012	0.07	0.18	0.21	0.23	0.22	0.23	3.14
Mozambique 2011	0.07	0.07	0.18	0.11	0.10	0.15	2.00
Niger 2012	0.12	0.15	0.13	0.15	0.13	0.15	1.29
Nigeria 2013	0.12	0.20	0.21	0.19	0.24	0.25	2.15
Rwanda 2014	0.24	0.21	0.09	0.21	0.21	0.21	0.88
Senegal 2017	0.11	0.17	0.28	0.29	0.16	0.32	2.89
Tanzania 2015	0.03	0.19	0.14	0.14	0.16	0.21	6.50
Uganda 2016	0.20	0.14	0.24	0.23	0.25	0.25	1.26
Zambia 2013	0.10	0.19	0.21	0.14	0.20	0.20	1.90
Zimbabwe 2015	0.09	0.09	0.21	0.19	0.17	0.15	1.75
Average	0.12	0.17	0.20	0.18	0.17	0.21	2.05

Table 2. Recall10 results, test set.

Survey	Wealth	Elastic-logit	rf	xgb	krls	Ensemble	Efficiency gain
Angola 2015	0.31	0.41	0.34	0.36	0.38	0.44	1.41
Benin 2011	0.29	0.36	0.47	0.34	0.39	0.46	1.59
Burkina Faso 2010	0.23	0.35	0.27	0.31	0.41	0.36	1.54
Cameroon 2011	0.24	0.25	0.32	0.24	0.27	0.29	1.22
Congo (D.R.) 2013-14	0.24	0.28	0.32	0.30	0.31	0.35	1.45
Cote d'Ivoire 2011	0.20	0.48	0.30	0.48	0.40	0.48	2.35
Ghana 2014	0.22	0.16	0.22	0.30	0.19	0.24	1.12
Guinea 2012	0.18	0.27	0.38	0.24	0.34	0.33	1.86
Kenya 2014	0.26	0.27	0.32	0.23	0.31	0.32	1.23
Lesotho 2014	0.34	0.31	0.28	0.21	0.17	0.24	0.70
Madagascar 2008	0.19	0.26	0.31	0.24	0.26	0.30	1.59
Malawi 2016	0.24	0.21	0.30	0.28	0.22	0.28	1.15
Mali 2012	0.16	0.35	0.34	0.34	0.35	0.39	2.47
Mozambique 2011	0.19	0.17	0.30	0.20	0.25	0.24	1.30
Niger 2012	0.19	0.23	0.20	0.27	0.24	0.26	1.41
Nigeria 2013	0.21	0.39	0.35	0.37	0.40	0.40	1.93
Rwanda 2014	0.48	0.30	0.33	0.36	0.36	0.36	0.75
Senegal 2017	0.26	0.36	0.38	0.39	0.30	0.41	1.57
Tanzania 2015	0.11	0.32	0.29	0.33	0.25	0.27	2.43
Uganda 2016	0.26	0.29	0.33	0.38	0.35	0.42	1.60
Zambia 2013	0.26	0.32	0.32	0.29	0.30	0.29	1.12
Zimbabwe 2015	0.17	0.26	0.23	0.30	0.21	0.34	2.00
Average	0.24	0.30	0.31	0.31	0.30	0.34	1.54

Table 3. Recall20 results, test set.

Similarly, Table 3 shows that in the countrywide average, only 24% of infant mortality is captured in the top quintile of risk according to the wealth-only model (average recall20 is 0.24). Rwanda is again an exception, with the top quintile accounting for 48% of mortality. Lesotho similarly has a recall10 of 0.34. In these two countries the wealth model out-performs the more flexible models attempted. On the other hand, the recall20 of the wealth model is no better than chance (≤ 0.20) in Cote d'Ivoire, Guinea, Madagascar, Mali, Mozambique, Niger, Tanzania, and Zimbabwe.

Beyond wealth: Gains from richer models

The elastic-net logit, rf, xgb, and krls models all augment the wealth model by adding the additional variables described above. Overall these models perform similarly to each other, and in all but a few cases, considerably better than the wealth-only model. According to these models, averaging by country, the top risk deciles account for 17-21% of mortality (recall10); the top risk quintile accounts for 30-34% of mortality (recall20). On average, the ensemble model performs better than any of the individual models capturing 21% of the mortality in its top risk decile (recall10), and 34% in the top risk quintile (recall20) when averaged across countries.

Compared to the wealth model, the ensemble model, on average, captures over twice as much of the mortality in its top risk decile (efficiency gain for recall10) and 1.5 times the mortality in its top risk quintile (efficiency gain for recall20). No single country always shows the highest recall10 across models, nor the lowest. Moreover, within given countries, the ensemble model outperforms the constituent models in many cases, does as well as the top one or two models in many others, and is never the worst-performing model in that country. Thus, relying on the ensemble model yields better results than any other model on average and safeguards against cases where individual models perform poorly in a given country.

Variable importance

Here we characterize the sets of variables most important to forming the risk forecasts, as indicated by variable selection choices in the elastic-net logistic regression and the established "variable importance" measure in the random forest (Table 4). In the elastic-net logistic regression model, every variable was included in at least 23% of the country-level models. A few variables, though, are included especially often: most models included the previous death of a sibling in the first year (86%) or the death of any previous sibling at any point (76%). 77% of country-level models included the child's gender; 73% included an urban/rural indicator. Over two-thirds (68%) of models included the mother's years of education, and many included the birth month (73% use sine-transformed birth month; 55% use cosine-transformed birth month). Maternal age was included in 64% of models.

For random forest, the "scaled variable importance" describes what fraction of the time a given variable was chosen to be included in the classification trees over which each random forest model aggregates. Compared to elastic-net logistic regression, the random forest was somewhat more selective, with over half of the variables

Variable	Elastic survival	Proportion positive	rf variable importance
Sibling death (< 1 yr.)	0.86	1.00	0.11
Mother education	0.77	0.06	0.39
Sin(month)	0.73	0.69	0.46
Female	0.73	0.00	0.17
Sibling death	0.71	1.00	0.11
Rural	0.64	0.71	0.10
Malaria	0.62	0.75	0.92
Maternal age	0.59	0.08	1.00
Prev. pregnancy ended	0.57	0.75	0.12
Clean water	0.55	0.58	0.14
Cos(month)	0.55	0.33	0.48
Log(birth order)	0.55	0.25	0.54
Mother age first union	0.52	0.45	0.63
Log(num bednets)	0.52	0.09	0.38
Safe floor	0.50	0.18	0.11
Age household head	0.50	0.64	0.77
Female household	0.50	0.55	0.11
Minority religion	0.50	0.70	0.06
Christian	0.35	0.43	0.08
Muslim	0.35	0.71	0.06
Unimproved toilet	0.32	0.86	0.09
Clean cooking fuel	0.31	0.20	0.05
(Wealth percentile)^2	0.27	0.00	0.94
Pit toilet	0.23	0.60	0.10
Wealth percentile	0.23	0.60	0.93

Table 4. Variable Importance.

appearing in fewer than 15% of the regression trees. A few variables were rarely included, such as clean cooking fuel (3%) and minority religion (5%). By contrast, maternal age was included nearly every time (99.6%), followed closely by wealth percentile and its square (93%, 94%). The latter set is particularly notable, as wealth percentiles were among the *least* commonly employed in the elastic-net logit models. Other important variables included malaria incidence rate in the area (92%), age of household head (77%), and mother's age at first marriage or union (63%).

Discussion

Across 22 countries in sub-Saharan Africa, wealth-based measures alone were fairly predictive of early mortality risk in a few countries (notably Rwanda, Uganda, and Lesotho) but weakly predictive in others and on average. On average, individuals in the highest 10% and 20% risk groups, according to these models, accounted for only 12% and 24% of infant mortality. When wealth information is combined with other risk factors in flexible models, it often aids in predicting individual-level risk of early mortality. On average across countries, the 10% of births with the highest risk, as predicted by the ensemble model, account for 21% of deaths, and the 20% at highest risk account for 34%. Compared to the wealth-only models, each country's ensemble model, on average, would have identified about twice as many deaths in the top risk decile and 1.5 times as many in the top 20% risk group. We note that the relatively weak relationship found between mortality risk and wealth in many countries is indeed a *within*-country phenomenon: it is quite possible that taken across countries, differences in wealth would more strongly predict mortality. However, the nature of the analysis and the wealth data in the DHS (which is indexed to be valid only within-country) leave us examining how risk and wealth relate within a given country.

The various models used to incorporate these risk factors—a regularized logistic regression (elastic-logit), random forest (rf), extreme gradient boosting (xgb), and kernel regularized least squares (krls)—have similar average performance. However, some models exceed others in particular countries. Using the average of predictions from these (ensemble) generates better performance while protecting against cases where individual models performed poorly in particular countries. Particularly important to many models were variables regarding the previous death of siblings in the first year, the mother's age and education, malaria prevalence, and the child's gender. The value of these variables as predictors says nothing of their causal influence or importance. Rather, they are likely to be predictive not because they cause or prevent mortality but because they signal the presence or absence of other important unmeasured factors that do.

These findings are consistent with previous findings that births at high risk of death exist across all socioeconomic groups and that, for various purposes, combine multiple risk factors and show that doing so improves our ability to identify the higher-risk births[8,17,18]. They also speak to a large body of literature in medicine and

public health that develops risk scores to identify those at risk of some event (e.g.,[19,20]). To our knowledge, flexible machine learning approaches using numerous pre-birth predictors have not previously been used to develop birth-level mortality risk scores in sub-Saharan Africa or elsewhere.

What recall10/20 represents

One reason to use recall10 or recall20 is that it lends itself to a simple interpretation of model performance, rooted in a hypothetical scenario involving an intervention with some fixed efficacy rate and fixed budget. Consider an intervention that a given country can afford to administer to only 10% of births. Ideally, this would be targeted to the 10% at the highest risk of early mortality. Let us suppose it has a fixed efficacy rate, which, as defined above, is the probability of saving a life that would have otherwise been lost. In a country with N_{births} births per year, the number of lives saved per year would be efficacy $* Pr(\text{death}|\text{high risk}) * N_{births}/10$. Applying Bayes' rule, this is equal to recall10 $*$ efficacy $* Pr(\text{death}) * N_{births}$ where $Pr(\text{death})$ is the overall mortality rate expected for children born in that year (absent the intervention). Notice that this (hypothetical) number of lives saved is proportional to recall10. By contrast, measures commonly used for classification accuracy in forecasting work, such as "f-scores" or or the "area under the curve" on a receiver operator characteristic curve (AUC-ROC), are more difficult to connect to a substantive policy-oriented interpretation.

To complete this hypothetical example, using recent estimates of the number of births and baseline mortality rates in each of these countries[21], the number of lives saved by an intervention with an efficacy of 5% would be 841 thousand for the wealth model, but 1.61 million for the ensemble model. In simpler terms, the efficiency gain of roughly two for these models compared to the wealth model (or, roughly, random targeting)) translates into approximately twice as many lives saved by an intervention with any particular fixed efficacy. Note that the average efficiency gain (2.3) is not exactly the ratio of lives saved because the latter takes the relative birth rates and mortality in each country into account.

Implications and implementation

The good news demonstrated here is that using flexible and inclusive modeling techniques can greatly improve our targeting ability and consequently produce multiplicative increases in the potential lives saved by a given intervention. Nevertheless, our analysis was only a demonstration in the general case, using widely available data. For a country to realize these gains, it will encounter practical exigencies that can increase and decrease the potential life-saving gains relative to our example. On the one hand, the real-world implementation of such an approach would be better and more appropriate than ours because the agency charged with constructing such a model can best choose to collect the variables that (i) are most likely to be predictive; (ii) are most feasible to collect widely; and (iii) are practical for purposes of targeting through existing administrative systems. On the other hand, such systems may not be prepared to collect some of the variables used here at high frequency. Further, the nature of an intervention and logistical limitations ultimately restrict the degree to which an intervention can be targeted. Nevertheless, targeting any intervention at any level can benefit from improved individual-level risk estimation since these individual-level risks can be aggregated to form the best predictions at any level. Concretely, at one extreme, suppose village health workers distributing a given intervention are positioned to enter specific data about individuals and determine their eligibility for an intervention. This would allow fully individualized risk-based targeting. Similarly, health workers could be instructed to use instead a heuristic scoring system (e.g. a checklist), which could be designed to approximate a well-fitted individual-level risk model. On the other hand, many interventions cannot be individually targeted, and are targeted to geographic areas (e.g., the placement of health facilities, decisions about where to spray for mosquitoes or to improve sanitation, etc). In these cases, individual-level risk scores still provide the ideal basis for targeting, as they can be aggregated at will to determine "area average risks" or "area total risk" for targeting. Our results suggest that in many countries, such an approach would offer larger improvements over making such decisions using local poverty or wealth measures.

Limitations and opportunities

Our work has several important limitations. First by focusing on individual-level risk of early mortality we demonstrate how the life-saving potential of *any* intervention—even one that can unfailingly prevent death by every cause—is severely limited by the accuracy of the targeting system used to determine who gets it. While improving risk modeling has proven here to be an effective way to amplify life-saving potential in this sense, future research could look beyond this theoretical maximum life-saving potential, considering, for example, the need for interventions that address specific causes of death and how this limits the life-saving potential of targeting or requires targeting systems sensitive to the cause of death addressed by each of several interventions. This would require a dataset that records the cause of death, which the IPUMS DHS data does not.

Second, while we aimed to rely only on the types of predictors that are already widely available through existing IPUMS DHS data, unfortunately, a few countries did not have surveys even in recent decades (Sudan) or lacked important basic variables such as district-level or lower geographic information and malaria prevalence (Burundi, Chad, Ethiopia, Liberia, Namibia). This limits generalization of our findings to areas with the lowest data availability but also points to the need to identify or build suitable data sources for these countries.

Third, the DHS sampling procedure attempts to produce a sample representative of the population. If these are not as representative as hoped, then while the results still reflect the predictive power of these models in some population of each country, the results may not be representative of these countries as a whole.

Finally, and most broadly, the central limitation is that not all deaths can be predicted, at least with the data we considered readily available in most countries. In general, we find our models for the highest 10% risk group capture only 20-30% of mortality. This is beneficial and worth using for targeting since the effectiveness of an intervention is proportional to this recall rate, as shown above. However, it leaves much unpredicted mortality

spread across the lower-risk groups. Thus, when resources can be made available, programs must still be targeted to a much wider group to capture a large fraction of births that will die. Relatedly, in the future and with appropriate data, approaches like ours could be refined if information on the cause of death were available. Associating each birth to a particular cause of death could generate models that suggest not only which births to target but also which intervention(s) to make available to them.

Data availability
The data used for this analysis can be found in the Integrated Public Use Microdata Series, available at IPUMS). Data need to be requested but are publicly available. Code to reproduce all results in this paper can be found at [repository information to be linked when accepted for publication].

References
1. Rajaratnam, J.K., Marcus, J.R., Flaxman, A.D., Wang, H., Levin-Rector, A., & Dwyer, L, *et al.* Neonatal, postneonatal, childhood, and under-5 mortality for 187 countries, 1970-2010: a systematic analysis of progress towards Millennium Development Goal 4. The Lancet. (2010)
2. Stuckler, D., Basu, S. & Mckee, M. Drivers of inequality in Millennium Development Goal progress: a statistical analysis. *PLoS Med.* **7**(3), e1000241 (2010).
3. Deaton, A. The Great Escape: Health, Wealth, and the Origins of Inequality. Princeton University Press; (2020)
4. Bendavid, E. Changes in Child Mortality Over Time Across the Wealth Gradient in Less Developing Counntries. *Pediatrics* **134**, e1551–e1559 (2014).
5. Moser, K., Leon, D. & Gwatkin, D. How does Progress toward the Millennium Development Goal affects Inequality between the Poorest and the least Poor? An Analysis of Demographic and Health Survey. *BMJ* **7526**, 1180–1182 (2005).
6. Black, R., Morris, S., Bryce, J. & Venis, S. Where and why are 10 million children dying every year?. *Commentary. The Lancet.* **361**, 2226–34 (2003).
7. Jones, G., Steketee, R., Black, R., Bhutta, Z. & Morris, S. How many child deaths can we prevent this year?. *The Lancet.* **362**, 65–71 (2003).
8. Ramos, A.P., Weiss, R.E., & Flores, M. Leave No Child Behind: Using Data from 1.7 Million Children from 67 Developing Countries to Measure Inequality Within and Between Groups of Births and to Identify Left Behind Populations. PlosOne. (7):25–45. (2020)
9. Akachi, Y., Steenland, M. & Fink, G. Associations between key intervention coverage and child mortality: an analysis of 241 sub-national regions of sub-Saharan Africa. *Int. J. Epidemiol.* **47**(3), 740–751 (2018).
10. Aheto, J. M. K. Predictive model and determinants of under-five child mortality: evidence from the 2014 Ghana demographic and health survey. *BMC Public Health* **19**(1), 1–10 (2019).
11. Breiman, L. Random Forests. *Mach. Learn.* **45**, 5–32 (2001).
12. Chen, T., & Guestrin, C. XGBoost: A Scalable Tree Boosting System. In: Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. KDD '16. New York, NY, USA: ACM; p. 785–794. (2016) Available from: http://doi.acm.org/10.1145/2939672.2939785.
13. Breiman, L., Friedman, J. H., Olshen, R. A. & Stone, C. J. *Classification and Regression Trees* (Wadsworth and Brooks, Monterey, CA, 1984).
14. Hainmueller, J., & Hazlett, C. Kernel regularized least squares: Reducing misspecification bias with a flexible and interpretable machine learning approach. *Political Anal.* p. 143–168. (2014)
15. Zou, H. & Hastie, T. Regularization and variable selection via the Elastic Net. *J. Roy. Stat. Soc. B* **67**, 301–320 (2005).
16. Krishna, A. Escaping poverty and becoming poor: who gains, who loses, and why?. *World Dev.* **32**(1), 121–136 (2004).
17. Ramos, A.P., Heymann, J., & Weiss, R. Improving Program Targeting to Combat Early-life Mortality by Identifying High Risk Births in India. *Popul. Health Metr.* 16(17) (2018).
18. Houweling, T.A.J., Klaveren, D.v., Das, S., Azad, K., Tripathy, P., & Manandhar, D., *et al.* A prediction model for neonatal mortality in low- and middle-income countries: an analysis of data from population surveillance sites in India, Nepal and Bangladesh. *Int. J. Epidemiol.* 48(1):86-198 (2019).
19. Mpimbaza, A. *et al.* Admission Risk Score to Predict Inpatient Pediatric Mortality at Four Public Hospitals in Uganda. *PLoS ONE* **10**(7), e0133950 (2015).
20. Beymer, M., Weiss, R.E., Sugar, C., Bourque, L.B., Gee, G.C., & Morisky, D., *et al.* Are Centers for Disease Control and Prevention Guidelines for Preexposure Prophylaxis Specific Enough? Formulation of a Personalized HIV Risk Score for Pre-Exposure Prophylaxis Initiation. *Sex. Trans. Dis.* 44(1) (2017).
21. World Bank Development Indicators. Birth rates, early mortality; Data retrieved from World Development Indicators, http://data.worldbank.org/. (2021)

Author contributions
A.P.R. designed the project and collected the data. S.S. wrote the computer code and produced tables. C.H. designed the statistical analysis. All authors wrote the paper and approved the final manuscript.

Competing interests

Additional information