

WORKING P A P E R

A Technical Supplement: Reducing the Global Burden of Acute Lower Respiratory Infections in Children

The Contribution of Improved Diagnostics

YEE-WEI LIM, FEDERICO GIROSI AND ROB BOER

WR-421-HLTH

December 2006

This working paper is the technical supplement to Lim YW, Steinhoff M, Girosi F, Holtzman D, Campbell H, Boer R, Black R and Mulholland K. Reducing the global burden of acute lower respiratory infections in children: the contribution of new diagnostics. *Nature*. S1 9-18 (2006)

This product is part of the RAND Health working paper series. RAND working papers are intended to share researchers' latest findings and to solicit informal peer review. They have been approved for circulation by RAND Health but have not been formally edited or peer reviewed. Unless otherwise indicated, working papers can be quoted and cited without permission of the author, provided the source is clearly referred to as a working paper. RAND's publications do not necessarily reflect the opinions of its research clients and sponsors. RAND® is a registered trademark.

PREFACE

This working paper contains technical material supporting the article by Lim et. al. “Reducing the global burden of acute lower respiratory infections in children: the contribution of new diagnostics” *Nature* S1; 9-18 (2006). It is intended to be read in conjunction with that article. This supplement includes additional material referred to in the published article. Although this technical supplement in its current form has not been formally peer-reviewed, an earlier version of this paper, which also contained material that appears in the corresponding *Nature* paper, was reviewed by two outside experts and was revised in response to their comments. The work was funded by the Bill & Melinda Gates Foundation to support the Global Health Diagnostics Forum.

ACKNOWLEDGEMENTS

The authors thank J. Hughes (Emory University, Atlanta, USA), M. Nelson (RAND Corporation, California, USA), K. Leuschner (RAND Corporation, California, USA), C. Boschi Pinto (World Health Organization Department of Child and Adolescent Health and Development, Geneva, Switzerland), T. Cherian (World Health Organization, Geneva, Switzerland), S. Graham (College of Medicine, Blantyre, Malawi), S. Madhi (University of Witwatersrand, Johannesburg, South Africa), E. Simoes (University of Colorado, Colorado, USA), M. Urdea (Halteres Associates, LLC, California, USA) and B. Schwartz (CDC, Atlanta, Georgia, USA).

The authors also thank E. Molyneux (College of Medicine, Blantyre, Malawi) and D. Zingmond (UCLA School of Medicine, Los Angeles, California, USA) for helpful comments on an earlier draft.

Table of Contents

Introduction.....	1
Methods.....	1
Modelling a new diagnostic test for bacterial ALRI.....	2
Modelling a new diagnostic test for severe ALRI	10
Model parameters.....	11
Model calibration.....	12
References.....	15
Figures.....	16
Figure 1. Decision tree for the bacterial acute lower respiratory infection (ALRI) model.	16
Figure 2. Decision tree for the severe acute lower respiratory infection (ALRI) model. .	17

Introduction

Acute lower respiratory infections (ALRI) are the leading causes of childhood mortality and morbidity. It is estimated that from 2000 to 2003, ALRI contributed annually to the deaths of >2 million children aged <5 years.¹ Despite some progress in reducing the number of ALRI-related deaths in developing countries, many children are still not being diagnosed or are not receiving adequate care. The analysis in Lim, et. al. “Reducing the global burden of acute lower respiratory infections in children: the contribution of new diagnostics. *Nature*. S1 19-28 (2006), and the additional materials included here are focused on understanding the potential reduction in mortality and morbidity associated with the introduction of hypothetical diagnostic tests for bacterial and severe ALRI respectively. The analyses consider test-performance characteristics and access requirements associated with new diagnostics, as well as availability of appropriate treatment.

Methods

We have developed decision-tree models to quantify the potential health benefits of new tests for bacterial and severe ALRI among children aged <5 years in Africa, Asia and Latin America. The analyses compare outcomes associated with current practice (that is, the status quo) to those associated with new tests. The outcomes include lives saved due to better diagnosis and reduction in overtreatment, and are a function of the test characteristics (sensitivity and specificity), the associated infrastructure requirements (advanced/moderate, minimal or none) and, when applicable, the level of access to effective hospital care. We have varied the assumptions about the model input parameters extensively in the sensitivity analyses to test the robustness of our results. The outcomes are estimated by calculating the incremental number of true-positive and false-positive cases of bacterial or severe ALRI relative to the status quo following the introduction of a new diagnostic test. One of our outcomes of interest is the number of children saved due to increased true-positive rates. We refer to this outcome as individual lives saved, to emphasize that it relates to individual children. The number of individual lives saved is computed using data from the published literature and the opinions of

experts from the working group on the risk of ALRI-related mortality in treated and untreated children.

We also calculate a societal outcome that considers the harm associated with unnecessary treatment. We quantify the harm of treatment in terms of the fraction of lives lost due to the treatment of one child. As these lives are lost as the result of indirect and unintended effects of treatment, we refer to them as indirect lives lost. The reduction in indirect lives lost associated with a reduction in the number of treatments is described as indirect lives saved. Indirect lives are a public-health concept: one indirect life cannot be matched to one particular individual. We also note that indirect lives are saved conditional on some behavioural assumptions (for example, that health-care workers will not find alternative reasons for prescribing antibiotics). Details of the method used to quantify the harm of treatment were reported by Girosi and co-workers.²

In addition to individual and indirect lives saved, we also consider the number of unnecessary treatments saved. By multiplying this number by the cost of treatment in US\$, we can derive the treatment cost savings associated with the introduction of the new test.

Modelling a new diagnostic test for bacterial ALRI

Defining the Population of Interest

The first step is defining the population of interest. Broadly speaking we are interested in children age zero to five with symptoms of ARI (cough and fever). Children around the world have four to six episodes of ARI a year, but it is important to recognize that, for a given child, not all episodes are alike, with most episodes being mild and self-limiting. Only a fraction of these episodes^a will be severe enough to trigger concern in the caregiver (usually the mother), and lead to healthcare seeking behavior. It is the set of children affected by these relatively more severe episodes that defines our population of interest. Denoting by I^{ARI} the incidence of ARI, and by p_{seek} the probability that an episode is severe enough to cause concern, the incidence of the cases which enter our model is $I \equiv p_{seek} I^{\text{ARI}}$. Henceforth, we will refer to these cases as “sufficiently severe cases of ARI”.

The condition which we want to detect in this population is bacterial ALRI. The incidence of bacterial ALRI is much smaller than incidence of ARI (between 0.03 and 0.05 episodes per child

^aSet to 1/3, based on conversations with field experts.

per year). We make the assumption that all episodes of bacterial ALRI fall in the universe of cases we consider. In other words, we assume that bacterial ALRI will always trigger health care seeking behavior. Denoting by I^{bALRI} the incidence of bacterial ALRI, this implies that the proportion of bacterial ALRI cases in our universe of cases is $p \equiv I^{bALRI} / I = I^{bALRI} / p_{seek} I^{ARI}$.

Defining the Clinical Path of a Child

The next step in the modelling process is to describe what happens to the child with a sufficiently severe case of ARI. In our model we divide children and their caregivers into three categories, according to the action taken by the caregiver. We refer to these categories as “access categories”, and they are as follows:

- The caregiver takes the child to a trained provider. A trained provider is not necessarily a medical doctor, but could be a nurse or equivalent medical personnel who has been trained to detect bacterial ALRI and is able to dispense treatment. Notice that we are making the explicit assumption that trained providers have access to the appropriate treatment of bacterial ALRI. Using data from UNICEF (www.unicef.pt/sowc06/statistics/tables.html) we estimated that in the developing world the proportion of children falling in this category is $p_{train} = 52\%$. This number however varies widely across regions, and it is known to be much lower in Africa (about 39%).
- The caregiver treats the child with effective antibiotics, acquired, for example, from a local pharmacy. This category includes caregivers who take their children to untrained or informal providers who always administer antibiotics. The proportion of children in this category varies greatly from region to region, and it will be denoted by p_{anti} .
- The caregiver is not able to provide any kind of effective care for the child. This category also includes caregivers who take children to traditional healers who do not make use of antibiotics, or buy antibiotics which are ineffective (for example because they are expired or have been inappropriately stored). We denote the proportion of children in this category by p_{nocare} .

What distinguishes these categories is the proportion of children who get appropriately and inappropriately treated with antibiotics. In other words, in each of the three categories above a

test is performed, which determines whether the child is treated with antibiotics or not. The characteristics of the tests performed vary as follows:

- For children taken to a trained provider we assume that a clinical diagnosis is performed, for example following IMCI guidelines. Studies of IMCI guidelines for the treatment of bacterial ALRI with antibiotics suggest that sensitivity of IMCI for detection of bacterial ALRI is about 85 percent.^{3,4} However, we know that the sensitivity of such a diagnosis increases with the severity of the case. Since we are only considering sufficiently severe ARI cases we argue that the sensitivity for the cases in our model should be higher than 85 percent, and we set it to 90 percent. The specificity of IMCI guidelines is lower, in the range of 70 to 80 percent.^{3,4} Since we are looking at sufficiently severe cases, we assume that providers may err on the side of caution for these cases, and therefore take the lower estimate for specificity. We denote the sensitivity and specificity of this clinical diagnosis as sens_{cd} and spec_{cd} respectively.
- When children are treated with antibiotics without any form of clinical diagnosis we can think of this as the outcome of a test which is 100 percent sensitive and 0 percent specific.
- When children receive no care we can think of this as the outcome of a test which is 100 percent specific and 0 percent sensitive.

The set of sensitivities and specificities listed above uniquely divides the children, in each access category, into the following four classes (which we refer to as “test outcome classes”)

- **True positives:** these are children with bacterial ALRI and who tested positive for bacterial ALRI. The probability that a child belongs to this category is:

$$P_{tp} = p[p_{train}\text{sens}_{\text{cd}} + p_{anti}]$$

- **False Positive:** these are children who do not have bacterial ALRI but are treated with antibiotics as a result of incorrect clinical diagnosis or because they belong the category of children who self-treat. The probability that a child belongs to this category is:

$$P_{fp} = (1 - p)[p_{train}(1 - \text{spec}_{\text{cd}}) + p_{anti}]$$

- **False Negatives:** these are children who do have bacterial ALRI but are not treated with antibiotics as a result of incorrect clinical diagnosis or because they belong the category of children who receive no care. The probability that a child belongs to this category is:

$$P_{fn} = p - P_{tp} = p[p_{train}(1 - \text{sens}_{cd}) + p_{nocare}]$$

- **True Negatives:** these are children who do not have bacterial ALRI and are not treated with antibiotics, as a result of correct clinical diagnosis or because they belong the category of children who receive no care. The probability that a child belongs to this category is:

$$P_{tn} = 1 - p - P_{fp} = (1 - p) [p_{train} \text{spec}_{cd} + p_{nocare}]$$

The probabilities listed above fully describe the status quo model. By multiplying these probabilities by the total number of cases entering our model (the incidence parameter I) we obtain the number of children falling into each of the test outcome classes. We denote these numbers by N_{tp} , N_{fp} , N_{fn} , and N_{tn} .

Defining Health Outcomes

The last modelling step is defining health outcomes corresponding to each test outcome class. This implies defining four quantities, that we denote by TP , FP , FN and TN , which quantify the “amount of health” assigned to a child in a given test outcome class. Once this is done, the total health outcome is computed as follows:

$$H = N_{tp} TP + N_{fp} FP + N_{fn} FN + N_{tn} TN$$

Several choices of health outcomes must be considered.

Counting Individual Deaths from Bacterial ALRI

The first outcome to consider is whether a child dies from bacterial ALRI. This implies assigning a value of 1 to instances in which the child dies of bacterial ALRI and a value of 0 when the child survives, or when the child does not have bacterial ALRI to begin with. The health outcome corresponding to a given test outcome class is then the average of these values, weighted by the appropriate case fatality. For example, a child in the “True Positives” class dies with probability m_{tp} , where m_{tp} is the case fatality of treated bacterial ALRI. Therefore we set $TP = m_{tp} \times 1 + (1 - m_{tp}) \times 0 = m_{tp}$. Similarly, we set $FN = m_{fn}$, where m_{fn} is the case fatality of untreated bacterial ALRI. Since children in the “false positives” and “true negatives” class never die of bacterial ALRI we set $FP = TN = 0$. Therefore, the total health outcome associated with the status quo, that is, the total number of deaths due to bacterial ALRI is:

$$D_{individual} = N_{tp} m_{tp} + N_{fn} m_{fn}$$

Using results from the literature and experts’ opinion, we determined that the case fatality for treated bacterial ALRI is in the range of 10-15%, while the case fatality for untreated is in the range 20-30%. In our model we choose the low end of both ranges, so that $m_{tp} = 10\%$ and $m_{fn} = 20\%$. If higher values are used then the number of deaths due to bacterial ALRI predicted in the model is not consistent with the current estimates of bacterial ALRI deaths.

Counting Indirect Deaths Due to Harm of Treatment

Whenever a child is treated (appropriately or inappropriately) some harm to society follows. In the case of bacterial ALRI harm derives from at least the following:

- Each time we treat a patient with antibiotics we increase the chance of development of antibiotic resistance. Antibiotic resistance implies that at some point in the future some people may die because of infection with a resistant strain of bacteria for which treatment may not be available. In addition, once resistance has built up, a new line of drugs has to be developed and administered, leading to additional cost, using up resources which would have been otherwise used to save lives.

- Each time we treat a patient with antibiotics there is a small possibility of an adverse drug reaction, leading to loss of lives.
- Each time we treat a patient with antibiotics we utilize scarce resources: we need to pay for the cost of treatment itself and to provide the labor necessary for administering the treatment. These resources could have been otherwise invested in cost-effective interventions, saving lives. Therefore, by providing treatment we miss the opportunity to save a number of lives, and therefore create a loss of lives.

We quantify the harm of treatment with a single number C , which represents the fraction of lives as a consequence of treating one child. For example, if we assume the only negative effect of treatment to be that once in 10,000 cases a child would experience a deadly anaphylactic reaction, then the harm of treatment would be equal to 0.0001. We refer to the total number of lives lost due to harm of treatment as “indirect lives”. The total number of indirect lives lost is computed as follows:

$$D_{indirect} = C(N_{tp} + N_{fp})$$

The calculation of C does not rely on an explicit calculation of the number of lives lost due to the sum of specific adverse effects of treatment, as the data necessary for such a detailed calculation are not available. Rather, C is calculated using a revealed-preference approach: if the medical community is in agreement that a diagnostic test with certain characteristics should be used, then the harm associated with treatment can be neither too low (otherwise it would be preferable to treat everybody) nor too high (otherwise it would be preferable to treat no one). Details on the method used to estimate C were reported previously by Girosi and co-workers.² The estimates for the bounds of C is the following:

$$\frac{p(1 - \text{sens}_{cd})(m_{fn} - m_{tp})}{p(1 - \text{sens}_{cd}) + (1 - p)\text{spec}_{cd}} \leq C \leq \frac{p \text{sens}_{cd}(m_{fn} - m_{tp})}{p \text{sens}_{cd} + (1 - p)(1 - \text{spec}_{cd})} \quad (1)$$

Plugging the values of the parameters shown in Table 1 into the equation above, we estimate that C is between 0.00026 and 0.0052. In the absence of other information, our estimate for C is the midpoint of this interval (0.0027). This implies that for every $1 / 0.0027 = 370$ antibiotic treatments administered, one indirect life is lost. In order to ascertain whether our estimates are of the right order of magnitude, we can calculate a lower limit to C using an alternative method, assuming that the financial cost of treatment is the only source of harm. Using this approach, if the cost of treating bacterial ALRI with antibiotics is 50 US cents, then for every 1,000 treatments administered, US\$500 is spent. If there is at least one intervention that can save one child at the cost of US\$500, then for every 1,000 treatments administered we miss the opportunity to save one child and the calculated harm of treatment is $C = 0.001$. This simple calculation gives us confidence that the harm of treatment is ≥ 0.001 , as other crucial components have not been included.

Counting Total Number of Deaths

The total number of deaths includes both the individual and the indirect lives lost. This is the criterion that we use to compare any test with the status quo, since it takes into account the relative weight of sensitivity and specificity, which is ultimately determined by the harm of treatment. The total number of deaths is computed as follows:

$$D_{total} = D_{individual} + D_{indirect} = N_{tp}(m_{tp} + C) + N_{fn}m_{fn} + N_{fp}C$$

Now that we have described in detail our representation of the status quo we can move to the description of the world with the new diagnostic.

Modelling the World with the New Diagnostic

We assume that when a new test is introduced, children who have access to a trained provider in the status quo receive it first, followed by those who self-treat and finally those who have no access to effective care.² In other words, access to the new test is correlated with levels of access to care, as found in the status quo. This model of access to care is explained in more detailed in the work of Girosi et al.² Let's take Africa as an example; from the statistical tables of UNICEF "State of the World 2005" we estimate that approximately 39% of the population will

access a trained provider upon an episode of ARI. A very small percentage of the population has the option of self-treating with effective antibiotics, about 5% according to experts' opinion, leaving the remaining 56% with no access to effective care. If we postulate that less than 39% of the population will have access to the new test then only a fraction of those who currently have access to a trained provider will have access to the new diagnostics. If we postulate that 50% of the population will have access to the new test then all the children who have access to a trained provider, all those who self-treat and some who currently have no access to care would have access to the new test.

Once the populations that have access to the new diagnostic have been identified using the hierarchical access model, then it is easy to quantify the health benefits resulting from such an innovation. For each of the populations affected we assume that the new test is adopted only if it is better than the status quo test. We use indirect lives to evaluate which of the two tests is better, using the estimate of the harm of treatment obtained in the previous section.

Modelling the Attributable Benefit

The benefit of the new diagnostic is then computed as the difference between the health outcome in the world with the new diagnostic and the health outcome in the status quo. Since health outcome is expressed in terms of lives, the health gain of the new diagnostic over the status quo is expressed in terms of lives saved (individual and indirect).

More formally, using superscripts *new* and *sq* to denote quantities in the world with the new diagnostic test and in the status quo, the number of lives saved by the new diagnostic test is expressed as follows:

$$\Delta D \equiv D_{total}^{sq} - D_{total}^{new} = \Delta N_{fn} (m_{fn} - m_{tp} - C) + C \Delta N_{fp}$$

where superscripts *new* and *sq* denote quantities in the world with the new diagnostic and in the status quo, and Δ refers to difference between quantities in the status quo and quantities in the world with the new diagnostic.

Modelling a new diagnostic test for severe ALRI

This model has a structure similar to the one for diagnosis of bacterial ALRI, and therefore it will be described in less detail, highlighting the crucial differences.

Defining the Population of Interest

The universe of cases entering the model is, as before, the set of sufficiently severe cases of ARI, with an incidence of $I \equiv p_{seek} I^{ARI}$. Within this set, we are interested in detecting the severe ALRI cases. Denoting by I^{ALRI} the incidence of ALRI, and by p_{severe} the probability that an ALRI episode is severe, the incidence of severe ALRI is $I^{sALRI} = p_{severe} I^{ALRI}$. The median incidence of ALRI among children 0 to 5 years of age is $I^{ALRI} = 0.28$ (per child per year)^b, with a range [0.2, 0.7]. Severe cases are estimated to be a proportion $p_{severe} = 9.9\%$ of all cases, so that the proportion of severe ALRI cases among all sufficiently severe ARI cases is $p = p_{severe} I^{ALRI} / (p_{seek} I^{ARI}) = 1.66\%$.

Defining the Clinical Path of a Child

The possible actions taken by a caregiver, or “access categories” are similar to those in the model for bacterial ALRI. Here are the options we consider

- The caregiver takes the child to a trained provider. This has been described above.
- The caregiver takes the child to an untrained provider (for example a pharmacist). We set the probability of this event at $p_{untrain} = 28\%$, based on experts’ opinion. We note here that this parameter does not critically affect any of our results.
- the caregiver is not able to provide any kind of effective care for the child.

Children with severe ALRI must be referred to a hospital. Within each access category children will be referred to hospitals with different probabilities, depending on the sensitivity and specificity of the clinical diagnosis.

- The clinical diagnosis for severe ALRI made by a trained provider has high sensitivity (95%) and moderate specificity (80%). Sensitivity is high because severe cases are relatively easy to diagnose.

^b Rudan I, Tomaskovic L, Boschi-Pinto C, Campbell H; WHO Child Health Epidemiology Reference Group. Global estimate of the incidence of clinical pneumonia among children under five years of age. Bull World Health Organ. 2004 Dec;82(12):895-903.

- The clinical diagnosis for severe ALRI made by an untrained provider is less sensitive and less specific than the one made by a trained provider. Following experts' suggestions, we set these test characteristics to be 10 percentage points less than the corresponding characteristics for a trained provider.
- If children receive no care, this is equivalent to a test that is 0% sensitive and 100% specific.

Counting Indirect Deaths Due to Harm of Treatment

For the severe ALRI model, we estimate $C = 0.0048$, with a lower limit of 0.001 and an upper limit of 0.006, so that for every $1 / 0.004 = 250$ children referred to a hospital and receiving effective care, one indirect life is lost. It has been suggested by field experts that, in this case, the main component of the harm of treatment is a lack of hospital beds: If a child uses the last bed, the next child coming to the hospital cannot be admitted because it is at full capacity. Parameter C also captures other costs of hospitalization, which include increased risk of nosocomial infections, risk to siblings of having their mother busy in the hospital taking care of the sick child, and risk of family resources being exhausted for the episode of hospitalization. We can use this argument to estimate the order of magnitude of the harm of treatment. For example, if we assume that 10% of the time a child hospitalized with severe ALRI would occupy the last bed, and that on average a child not admitted to a hospital would experience a 5% increase in the risk of dying because of the missed treatment, then the harm of treatment would be $C = 0.1 \times 0.05 = 0.005$, which is within our estimated limits.

Model parameters

The parameter estimates used for the two analyses are shown in Tables 1 and 2. These estimates were obtained from a variety of sources, including the published and unpublished literature, databases of international agencies (such as the United Nations Children's Fund and the WHO) and expert opinion. Expert consensus was reached for the parameter estimates and the plausible ranges of all key inputs for both models.

The ability of a health facility to use a new diagnostic test is characterized by its infrastructure requirements, which follow those described by Olmsted and colleagues.⁵ Advanced/moderate infrastructure implies consistent access to running water and electricity, a need for minimal laboratory equipment and a trained provider (such as a nurse or laboratory technician). Minimal infrastructure implies limited access to water and electricity, a physical location with no laboratory equipment and a minimally trained health provider (such as a pharmacist or village health worker). No infrastructure refers to settings with no reliable water or electricity, and makes no assumptions about the training or literacy of the caregiver; a test that can be performed in these settings is essentially a home diagnostic tool.

The estimates of access to different levels of health-care and infrastructure used in these analyses were derived from a multinomial logit model using data from the Demographic and Health Surveys conducted from 2000 to 2005 for 17 African, six Asian and six Latin American countries.⁵ For example, in sub-Saharan Africa, a test that requires advanced/moderate infrastructural support would be available to only 28% of the population.⁵ If the test requires only minimal infrastructural support, an additional 47% of the population might have access to it (75% in total). By contrast, a test that requires advanced/moderate infrastructure would be already accessible to 58 and 90% of the Asian and Latin American populations, respectively (Table 3).⁵

Model calibration

In order to ensure that the model parameters of Tables 1 and 2 provide a meaningful picture of the status quo we go through a calibration exercise: we estimate the current numbers of bacterial and severe ALRI deaths and compare them with the number of deaths predicted by the model. While we do not expect perfect agreement, the models should reproduce the observed number of deaths with reasonable accuracy.

Calibration of the model for diagnosis of bacterial ALRI

The number of bacterial ALRI deaths is obtained by first computing the number of treated and untreated bacterial ALRI, and then multiplying these numbers for the appropriate case fatality. Denoting by *pop* the population of children age 0 to 5 in the developing world, the

number of bacterial ALRI cases is $popI^{bALRI}$. Bacterial ALRI cases are treated only if the child has access to a trained provider and is diagnosed positive, or if the child belongs to the “self-treating” category. Therefore the number of deaths is:

$$\begin{aligned}
 D^{bALRI} &= popI^{bALRI} \left[(sens_{cd} p_{train} + p_{anti}) m_{tp} + (1 - sens_{cd} p_{train} - p_{train} - p_{anti}) m_{fn} \right] \\
 &= popI^{bALRI} \left[m_{fn} - (m_{fn} - m_{tp}) (sens_{cd} p_{train} + p_{anti}) \right]
 \end{aligned}
 \tag{2}$$

This number must match the total number of bacterial ALRI deaths observed in the developing world. A direct estimate of this number is not available from WHO data, and so we proceed as follows. From ref. (1) we know that ARI is responsible for 19% of the deaths of children under age 5. However, this number does not take in account the fact that 26% of all neonatal deaths are due to pneumonia or sepsis, and that neonatal deaths account for 37% of all deaths under age 5. Splitting neonatal deaths equally between sepsis and pneumonia, we find that ARI deaths account for 24% of all deaths under age 5 ($24\% \approx 19\% + 1/2 \times 26\% \times 37\%$). We estimate the number of bacterial deaths as 81% of all ARI deaths. The value of 81% is obtained by assuming that 90% of all ARI deaths are due to severe ALRI, and 90% of all severe ALRI deaths are due to bacterial ALRI. These last 2 figures have been determined in conjunction with experts. Using the number of deaths reported in the latest WHO World Health Report, we estimate the number of bacterial ALRI deaths to be 2 million. Plugging the values of the parameters of table 1 in the equation above, we predict a number of deaths equal to 2.01 million, obtaining a perfect calibration. If we assume that none of the neonatal deaths were attributable to sepsis the number of bacterial ALRI deaths would climb to 2.4 million, while if we assume that none of the neonatal deaths were attributable to pneumonia it would shrink to 1.6 million, creating errors of about 20% in the calibration. This potential margin of error seems reasonable, given the uncertainty associated with the mortality figures. Clearly the calibration is affected by all of the parameters appearing in equation (2) above. In order to provide the reader a sense of how the different model parameters affect the predicted number of deaths, we have implemented the calibration procedure in Excel. In the interactive Excel spreadsheet `bacterial_world_calibration.xls` the users can vary the parameters of equation (2) at will, and see how they affect the calibration process. Overall, if we allow a calibration error of

about 20%, we find that the model parameter can vary in reasonable ranges and still provide an acceptable calibration.

Calibration of the model for diagnosis of severe ALRI

The calibration of the model for a diagnosis of severe ALRI follows the same steps outlined above, with few differences. The number of deaths by severe ALRI predicted by the model is

$$D^{sALRI} = pop I^{sALRI} \left[m_{fn} - (m_{fn} - m_{tp}) (sens_{train} p_{train} + sens_{untrain} p_{untrain}) p_{hosp} \right]$$

Here $sens_{train}$ and $spec_{train}$ are the sensitivity of a clinical diagnosis of severe ALRI by a trained and untrained provider respectively, while p_{hosp} is the probability that a child in need of hospitalization has access to a hospital that can provide appropriate treatment. The number of deaths predicted by the model, using the parameter values of table 2, is 2.1 million, while the number of deaths estimated from WHO data is 2.2 million, providing excellent agreement. The interactive Excel spreadsheet used for the calibration can be found in the file `severity_world_calibration.xls`.

For results and discussion, please see Lim et al, Reducing the global burden of acute lower respiratory infections in children: the contribution of new diagnostics. *Nature*. S1 9-18 (2006)

References

1. Bryce, J., Boschi-Pinto, C., Shibuya, K. & Black, R. E. WHO estimates of the causes of death in children. *Lancet* **365**, 1147–1152 (2005).
2. Girosi, F et al.. *Determining the Priority Global Health Needs and Quantifying the Health Benefits resulting from the Introduction of New Diagnostics in the Developing World. WR-418-HLTH* (RAND Corporation, Santa Monica, 2006). Available online at http://www.rand.org/health/feature/research/0612_global.html (as of December 7, 2006).
3. Factor, S. H. *et al.* Diagnosis and management of febrile children using the WHO/UNICEF guidelines for IMCI in Dhaka, Bangladesh. *Bull. World Health Organ.* **79**, 1096–1105 (2001).
4. Kalter, H. D. *et al.* Identifying sick children requiring referral to hospital in Bangladesh. *Bull. World Health Organ.* **75 (suppl. 1)**, 65–75 (1997).
5. Olmsted, S. S., Derose, K. P. & Beighley, C. *Determining Access to Care and User Requirements for Diagnostic Tests in Developing Countries. WR-423-HLTH* (RAND Corporation, Santa Monica, 2006). Available online at http://www.rand.org/health/feature/research/0612_global.html (as of December 7, 2006).

Figures

Figure 1. Decision tree for the bacterial acute lower respiratory infection (ALRI) model.
 This figure has been drawn using the TreeAge Pro Suite software (TreeAge Software, Inc.). The outcome here is the sum of individual and indirect lives. Individuals are assigned an outcome of 1 if they survive, 0 if they die and $-C$ if they have been treated, where C is the harm of treatment.

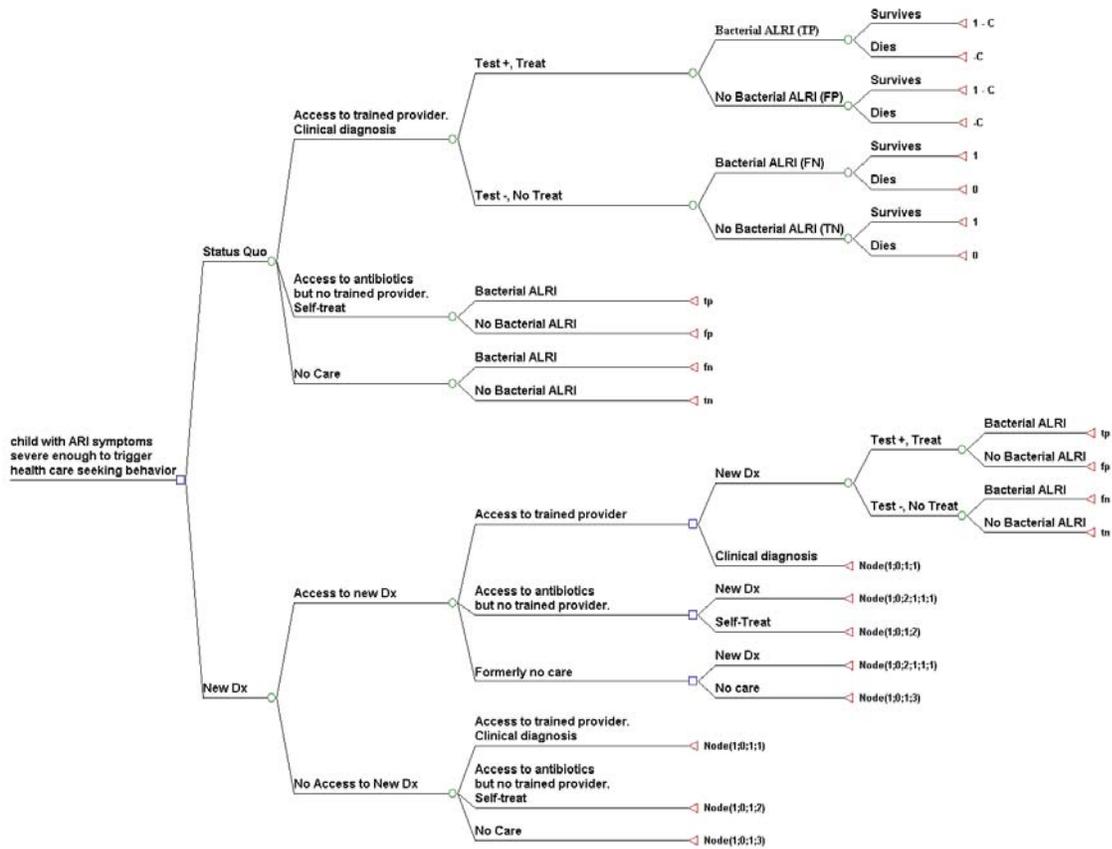


Figure 2. Decision tree for the severe acute lower respiratory infection (ALRI) model.

This figure has been drawn using the TreeAge Pro Suite Software (TreeAge Software, Inc.). The outcome here is the sum of individual and indirect lives. Individuals are assigned an outcome of 1 if they survive, 0 if they die and $-C$ if they have been treated, where C is the harm of treatment.

