

A conversation with Caitlin Tulloch, March 13, 2018

Participants

- Caitlin Tulloch – Technical Advisor, Best Use of Resources Initiative, International Rescue Committee (IRC)
- Josh Rosenberg – Senior Research Analyst, GiveWell
- Chelsea Tabart – Research Analyst, GiveWell
- Stephan Guyenet – Research Consultant, GiveWell

Note: These notes were compiled by GiveWell and give an overview of the major points made by Ms. Tulloch.

Summary

GiveWell spoke with Ms. Tulloch of IRC to discuss the work the IRC has been doing producing cost-efficiency analyses of its programs and its funding gap for community management of acute malnutrition (CMAM) programs. Conversation topics included IRC's cost-effectiveness estimates for CMAM, treatment protocol, monitoring and evaluation, and room for more funding.

Community management of acute malnutrition (CMAM) treatment protocol

Patients are admitted to CMAM treatment for severe acute malnutrition (SAM) when their mid-upper arm circumference (MUAC) is below 115mm and are considered "recovered" when their MUAC is at either 115mm or 125mm (depending on the program and whether follow-up moderate acute malnutrition (MAM) treatment is available). Patients continue to receive treatment until a) recovery, b) "default" (i.e., dropping out of the program), c) death, or d) referral out for complications.

MAM is characterized by a MUAC between 115mm and 125mm. Children in this category are still at a very high risk of mortality. Currently, SAM and MAM treatment are managed by two separate UN agencies, with different products and protocols used for each. These programs can receive very different levels of funding, and there is no continuous treatment path for patients who progress from MAM to SAM. As a result, children often progress all the way to SAM before receiving treatment, since there's not enough funding to treat them at the MAM stage.

Testing a unified MAM/SAM protocol

The Combined Protocol for Acute Malnutrition Study (ComPAS), an ongoing joint study between IRC, Action Against Hunger (ACF), and the London School of Hygiene and Tropical Medicine, is testing a unified protocol for MAM and SAM in two contexts (South Sudan and Kenya), which might be significantly more effective since, ideally, it will identify people in need of treatment early, before they develop SAM. This approach is also expected to be more cost-effective due to the simplified logistics involved in using only one product for both conditions (at tapered doses)

and treating all children in a unified system. The ComPAS study is scheduled to be completed in mid-July, with results expected by the end of 2018. IRC expects it to have high-quality monitoring data (including population-based surveys and good tracking of patient outcomes), as well as a comprehensive cost-effectiveness study.

Cost-effectiveness

Ms. Tulloch thinks IRC's CMAM programs appear especially cost-effective among IRC programs. The CMAM protocol is a particularly well-developed and understood intervention (though it has some limitations, particularly with achieving high coverage in certain contexts).

IRC estimates that CMAM costs roughly \$300 to \$500 per child treated (for children who did not default during the timeframe covered by IRC's report), with some variation in cost per child treated due to contextual factors. One review (http://www.who.int/nutrition/topics/Statement_community_based_man_sev_acute_mal_eng.pdf) estimates that the case fatality rate for children with SAM in sub-Saharan Africa is 20% without treatment. The mortality rate may be reduced to as low as 0% with treatment. At a cost of \$300 to \$500 per child treated, that would imply a cost-per-life-saved through CMAM of roughly \$1,500 to \$2,500.

It is possible some of the children included in this analysis continued treatment beyond the one- to two-year timeframe of the report and then defaulted; more comprehensive tracking protocols, which IRC is in the process of rolling out to more of its field teams (see below), would include methodologies to prevent that possibility from skewing the results.

Fixed center-level costs

Nearly all costs of IRC's CMAM programs are fixed at the center level (i.e., almost no costs vary with the number of children treated). This is in part because the ready-to-use therapeutic food (RUTF) used is mostly donated by UNICEF and other groups (the cost of RUTF is not included in IRC's analysis). CMAM programs' cost-efficiency is therefore constrained by the population density of the targeted area, and getting coverage rates as high as possible (e.g. by doing outreach to the local population) is a major component of maximizing cost-effectiveness.

Since costs to run a center are roughly fixed, cost efficiency might also decrease over time if fewer people end up needing treatment, e.g., if food security increases, or fewer people progress from MAM to SAM.

Government contributions

It can be difficult to incorporate government contributions into cost estimates, in part because IRC operates its CMAM programs under a variety of support models. In some countries, CMAM is almost entirely funded and managed by IRC, while in other places government contributions are more significant. In some contexts, governments' contributions are difficult to track. The ComPAS study's research protocol includes improved documentation of government contributions; IRC

explored including full ingredients costing in the CompAS study (rather than only tracking costs that appear in IRC's finance system).

IRC is unsure about the extent to which its programs have "crowded in" government funding for CMAM.

Monitoring and evaluation

The basic monitoring and evaluation package used by most of IRC's CMAM programs is program-based rather than population-based — i.e., it monitors people who pass in and out of the program, but does not include a full prevalence estimate of SAM in the general population, or track mortality rates in the general population vs. among program participants. IRC does currently use a population-based survey protocol in three or four countries, which it plans to roll out more widely.

There is ongoing discussion at IRC about the core set of indicators that all field teams will be required to track going forward; these will likely include tracking people who a) were treated and did not default, b) defaulted, c) were referred out of the program, d) recovered, or e) died.

Ideally, when comparing mortality inside and outside of a CMAM program, the "denominator" (mortality outside the program) ought to capture mortality attributable to SAM; if it is capturing mortality from secondary complications (of the kind which, in standard CMAM procedure, might result in a patient being referred out for inpatient treatment), that could skew the ratio. There are methods for dealing with that issue. In contexts in which more comprehensive monitoring protocols have been rolled out, IRC tracks outcomes for patients who are referred out of CMAM to inpatient care (who are also generally more likely to die). IRC typically tracks mortality data among people who have ever passed through the program, including people who default.

Key indicators

IRC generally does not construct a counterfactual each time it runs a program; instead, it monitors key indicators that align with components of a particular theory of change. The key indicators tracked most often include:

- Coverage rate for treatment
- Mortality rate
- Recovery percentage
- Default percentage
- Prevalence of SAM in the population (which is also tracked because a certain threshold of SAM prevalence in a population is what triggers an official declaration of "famine").
- Co-morbidity with other conditions (e.g. HIV, tuberculosis), which is typically the reason that people are referred out of the program for inpatient treatment. (Until the 1990s, all SAM cases received inpatient treatment, which seems to have led to increased co-morbidity.) The final cause of death for patients in inpatient care is now carefully tracked.

Budget and room for more funding

IRC had about \$740 million in revenue in fiscal year 2016. IRC could likely absorb tens of millions of dollars in additional funding for its CMAM programs.

With more funding, it might make sense to invest in smaller-scale crises that have received less attention (in, e.g., Niger, Chad, and the Central African Republic). In some of these places, UNHCR has paid the fixed costs of basic infrastructure for country offices, so increasing funding to achieve scale could have a particularly high return on investment. (Availability of food is rarely a bottleneck for CMAM programs, since RUTF is fairly inexpensive, produced locally in many places, and usually donated by UNICEF.)

The five major famines in the past year (in Somalia, Nigeria, Yemen, Niger, and Sudan) did not activate large responses from donors and still have considerable funding gaps.

The UN's Financial Tracking Service (FTS) puts out a humanitarian response plan (HRP) that estimates the total funding need in a given context. FTS then tracks funding commitments to those contexts and monitors the remaining gap; last year, funding of the FTS's HRPs was at roughly 50%. Ms. Tulloch thinks that data can likely be broken down by cluster in order to find, e.g., the unfunded portion of HRPs dedicated to nutrition by country and year. Ms. Tulloch's impression is that that gap is generally large, particularly for the "forgotten" crises in e.g. the Sahel and South Sudan.

All GiveWell conversations are available at <http://www.givewell.org/conversations/>