The life (only) you can save

Credible scientific evidence shows how a small charitable donation can save a child's life. So why haven't governments done those things already?

Premise: you can open a new browser tab, enter your credit card details, and save a life for $4,500

GiveWell is a charity fund that directs money to causes that save the most lives for the least amount of money. Over the past decade it has directed over $1 billion in charitable giving, which it says "will save over 150,000" lives. If these numbers are correct, those donations are incredibly cost effective. For example, GiveWell calculates that $4,500 directed to its top charity, the Against Malaria Foundation in Guinea, buys one thousand bed nets, which leads to an increase of 1,431 people sleeping under nets, and one additional life saved.

Many charities make bold claims about their impact. But a new crop of charities and charity advisory services associated with the "effective altruism" movement – including GiveWell, Open Philanthropy, and the inspiration for the title above, The Life You Can Save – distinguish themselves by their commitment to rigorous evidence and transparency in identifying the most cost-effective causes. In short, the best available evidence suggests their claims about lives saved are not just marketing spiel: statistically speaking, each $4,500 to the Against Malaria Foundation really does save a life.

To be sure, bed nets and vitamin A supplements represent the low hanging fruit in global health and development. But that is not a criticism; it's the whole point. This modern crop of evidence-based charities targets low-cost, scalable programs with solid data behind them and fairly simple, linear theories of change: where vaccination coverage is low, pay people to vaccinate their kids so they don't die of preventable diseases. Where intestinal worms are endemic, give kids deworming pills at school en masse so fewer kids get sick.

Paradox: why haven't public institutions dedicated to saving lives done those easy things already?

From a public policy perspective, however, the ability of organizations like GiveWell to credibly save thousands of lives with small donations poses the question: why weren't these things happening anyway? If delivering bed nets is such obvious low-hanging fruit, then why hasn't the government of Guinea or the myriad official aid agencies operating in Guinea completely saturated the market with free bed nets already?

The government of Guinea spends around half a billion dollars each year on health care, or USD $40 per capita. From 2017 to 2021, the country received about $195 million in annual foreign aid spread across its population of roughly 14 million, including about $12 million per year specifically for malaria control.

If bed nets are, in fact, the most cost effective way to save lives, then surely those $40 could've covered the $0.20 price tag. Was the government budget, or even just the foreign aid given specifically to fight malaria, spent on even more cost-effective programs? Do domestic and international policymakers disagree with GiveWell's assessments of its own giving and the lives it is saving? Are other bureaucratic or political factors preventing traditional actors from providing the life-saving commodities that these charities offer?

Hypotheses: contested facts, evidence standards, bureaucracy, politics, and implementation capacity

Fleshing out these questions a bit more, we can divide explanations for why governments and foreign donors don't already do what GiveWell does into four broad buckets, or groups of alternative (but not mutually exclusive) hypotheses:

1. **Facts and generalizability**. While GiveWell works hard to establish the scientific basis for its grants, disagreement may remain about the true cost-effectiveness of GiveWell charities. Even if we accept the evidence about the efficacy of insecticide treated bed nets or deworming pills, policymakers may question the external validity of these findings in their particular context or the ability of state institutions to implement them faithfully at scale.

2. **Evidentiary standards**. It's widely accepted that the kinds of evidence that GiveWell prioritizes -- e.g. randomized trials and microeconometric studies -- lend themselves towards atomistic, 'micro' interventions. Investments in systems infrastructure and national policymaking may be equally or more consequential but harder to evaluate with an equal degree of rigor. States and aid donors with a higher tolerance for ambiguity might prioritize these things without accepting the premise that GiveWell-style interventions are more cost effective.

3. **Donor constraints**. Another possibility is that all parties agree GiveWell charities are achieving more good per dollar than governments and traditional aid donors, but the latter are simply unable to fill that niche. For instance, political or bureaucratic constraints within donor agencies may influence aid allocation in ways which technical staff deem sub-optimal.

4. **Recipient preferences**. A similar phenomenon may exist on the recipient side. Public spending decisions are made through complex political processes, not technocratic cost-benefit analyses. Various constituencies may prefer other spending priorities above bed nets or vitamin A supplements. Even among aid donors, deference to country ownership or and the priorities of local actors may sway spending in the direction of other topics and interventions.

5. **Implementation capacity and transaction costs.** A common argument for philanthropy is that it has the flexibility and private-sector entrepreneurial know-how to identify opportunities and take risks that public sector institutions can’t pull off. On the ground, privately-funded and privately-led implementation may be able to execute projects at a speed and cost that public entities can’t. By these hypotheses, it’s not that public actors disagree with GiveWell’s

assessment of its top charities, nor that they preferred to do other things, it's simply that public donors can't do what GiveWell's charities do. It would cost them more per bed net delivered, oversight of staff in the field would be weak, corners would get cut, and fewer lives would be saved.

Proposal: ask policymakers directly

We'd like to make this exercise as concrete as possible. So we'd propose starting with GiveWell's list of Top Charities, and selecting interviewees with knowledge of these topics and contexts.

1. "Medicine to prevent **malaria**" via the Malaria Consortium, working in Nigeria, Chad, Burkina Faso, Guinea, Mali, and Niger.
2. "Nets to prevent **malaria**" via the Against Malaria Fund, working in DRC, Ghana, Guinea, Malawi, Papua New Guinea, Togo, Uganda, and Zambia.
3. "Supplements to prevent **vitamin A** deficiency" via Hellen Keller International, working in Burkina Faso, Cameroon, Côte d'Ivoire, the Democratic Republic of the Congo (DRC), Guinea, Kenya, Mali, Niger, and Nigeria.
4. "Cash incentives for routine childhood **vaccines**" via New Incentives, working in North West Nigeria.

The bigger challenge is to identify the organizations which could or should, one might argue, have filled the niche that GiveWell occupies. Who exactly ought to be paying for bed nets for poor households in Guinea is open to debate, but we'd like to identify actors who fulfill some combination of two criteria: (a) normatively they should be concerned for saving lives in the relevant population, and (b) they have the funds and capacity to do so; e.g., they're operating health programs in the same geographic area or the same disease silo.

As a starting point, we propose the following list of interviewees:

- Government: Ministry of Health officials in Nigeria and the DRC, which together span all four of GiveWell's top charities.
- Multilaterals: Global Fund, World Bank, UNICEF, the European Commission
- Bilaterals: US Agency for International Development, UK Foreign and Commonwealth Development Office, Agence Française de Développement
- Other philanthropies: Gates Foundation
- Grantees: staff at the Malaria Consortium, Against Malaria Fund, Hellen Keller International, and New Incentives.

Methodologically, the work would include three main components:

1. A review of the evidence establishing the impact of GiveWell top charities, including (e.g.) clinical trial evidence on the effectiveness of interventions, evidence that the scope for impact is similar in the actual program locations, and direct evidence of ongoing impacts.

2. A review of the public spending portfolio of government ministries of health and foreign aid donors in selected countries, as well as an assessment of the evidence base and likely health/welfare impact of these alternative spending choices.
3. Semi-structured interviews with government and aid agency staff. These interviews would combine a small quantitative survey and open ended questions. Topics would include:
    a. Descriptions of the allocation process guiding the deployment of public health monies in these contexts.
    b. Respondents views on the likely impact of GiveWell charities
    c. Their view of the development effectiveness of alternative spending priorities in their own agencies.
    d. Political and bureaucratic feasibility of both allocating money to GiveWell's causes and implementing programs of the type which GiveWell charities implement.

Stepping back, we see two main audiences for this work. The main audience is policymakers in developing-country governments and officials in foreign aid agencies making allocation decisions. If various hypotheses above hold true – especially #3 – then using GiveWell's portfolio as a reference point and fleshing out the obstacles to public institutions filling the gaps currently filled by private philanthropy could help to improve how they gather and evaluate evidence and how they make resource allocation decisions.

In the same way that cash transfers increasingly serve as a benchmark for other aid programs (if you can't beat cash, why not just do that?), we think charities like GiveWell could serve as a more nuanced benchmark for a broader swath of official aid programming.

The second audience is private philanthropy itself. If other hypotheses noted above prove more relevant – especially #1 and #2, and possibly #4 – that would seem to imply some case for a course correction in the allocation of charitable giving. Those lessons might (or might not) push in the direction of concrete examples of where thinking more in terms of holistic health systems rather than discrete interventions might be justified, or engaging more with locals' preferences over resource allocation.