The UK as an Effective Altruist

Euan Ritchie and Ian Mitchell

The prime minister's most influential advisor, Dominic Cummings, is a champion of "effective altruism"—the use of evidence and careful reasoning to work out how to maximize the good with a given unit of resources (with "the good" understood in impartial welfarist terms), and applying the findings to improve the world as much as possible. Michael Gove, the Minister for the Cabinet Office, has spoken on similar lines about public servants' need to avoid group-think and the government's need to be "rigorous and fearless in its evaluation of policy and projects." He's urged the government to prioritise "what works" and to counter a bias against doing things differently, ideals reflected in the effective altruism movement.

With the UK government in the midst of a major "Integrated Review" of its foreign, development, and defence policy—and the recent the formation of a merged Foreign, Commonwealth and Development Office confirming it isn't afraid of change—now's a good time to consider whether the effective altruism movement can or should find great traction in UK aid programmes.

In this note we briefly explore the key tenets of effective altruism, consider where the UK government aligns with these tenets in its aid programme, and identify five ways the new Foreign, Commonwealth and Development Office could move further in that direction. Finally, we look at how a government differs from most effective altruists. In essence, effective altruism is about evidence-based approaches to maximise the expected value of an impact. But if a government took this approach seriously, it would also exploit opportunities that smaller individual donors may not have—thinking bigger and pursuing more systemic change where even if success is uncertain, the potential returns are huge.

WHAT ARE THE PRINCIPLES OF EFFECTIVE ALTRUISM?

In "Doing Good Better," William Macaskill (one of the leading recent figures in the effective altruism movement) states, "Maximising expected value is generally regarded as the best strategy for making decisions when you know the value and the probabilities of each option." Discerning these quantities is hard, but attempting to do so allows objective comparisons, or at least forces us to state assumptions that can then be contested (e.g., this discussion on alternatives to the disability-adjusted life year [DALY] metric widely used in development and public health circles).

At the heart of effective altruism is a concern with comparing the impact of different options, using as high-quality evidence as possible. For example, in health economics, the impact is translated into a common unit, such as DALYs, and then this impact is weighted by the probability of success (i.e., using an "expected value" framework). This is not to say that effective altruists have shied away from problems that are difficult to measure. Over the years, the movement has strongly endorsed systemic change and supported many causes such as biosecurity and pandemic preparedness, for which the impact is hard to quantify and no gold standard randomized control trial (RCT) evidence is available. When certainty of change is low, the impact of such a change may still mean that in expected value terms, it is worth a try.

Recognising that not all evidence is created equal, and for some problems the highest standard of evidence is unobtainable, effective altruists adopt a Bayesian approach: reviewing the evidence, weighting it by its quality, and incorporating it as part of a wider evidence base.

Naturally this is easier in some cases than others: assessing the impact of increasing research spend is hard to compare to lives saved by distributing bed nets but may be equally or more important. RCTs feature heavily in the evidence used to assess impact (some argue they are given too much weight) but not everything can be randomized: growth illustrates this well. But even if difficult, effective altruism recognises the need for some basis of comparison, even if these are "fermi estimates": back-of-the-envelope calculations to generate orders of magnitude and plausible ranges for different impacts. Budget allocations are unavoidably quantitative decisions; the alternative to explicitly trying to put a quantitative value on different projects is doing so implicitly. Better then to do the former so that we are forced to think carefully about assumptions and projections.

Therefore, cost-effectiveness yardsticks are essential to consider if money is going to be spent as effectively as possible, even if they are highly uncertain. Below we look at some of the principles that help guide effective altruists in assessing impact and what they could mean for the UK's development approach.

Scale

If the intervention we choose is successful, how much of the overall problem will that solve? For example, if we manage to eradicate malaria, how much will this improve overall health outcomes? Clearly this necessitates a precise statement of the problem; for example, "what percentage of total DALYs would this avert," or "by what percentage would maternal mortality fall"?

Scale is also important in another way: what is the total amount of giving that you can influence? If an individual improves the cost-effectiveness of their charitable donations, this is welcome but marginal. If the same individual could join a large organisation—perhaps the US Agency for International Development or the Global Fund—and make *that organisation's* spending more effective, that holds the potential for improvement on a much larger scale. Similarly, the UK is an important player in international development and so the effectiveness of UK aid matters, but if the UK can be a positive influence globally—through the multilateral system or leading other development actors, for example—this holds even greater promise.

Solvability

By committing additional resources, how much more of the problem could be solved? For example, if we double resources spent on combating malaria, how many extra cases could we avoid? Alternatively, what probability do we think we have of solving the problem? In comments to the House, Boris Johnson rightly highlighted the importance of democracy (if in rather colourful language). But we don't really have a firm grasp of how to spend money in a way that brings about democracy in more places. In other words, while clearly desirable, it isn't solvable. This doesn't mean that we shouldn't try to find new evidence, or avoid all uncertain projects: Abhijit Banerjee and Esther Duflo (among others) argue that we still don't really know how to bring economic growth as outsiders, whereas Lant Pritchett argues that even if we only have a tiny probability of bringing about growth, the gains are so large it is still worth trying. But "solvability" must remain a key consideration.

"Neglectedness"

How many resources are already dedicated to solving this problem, and how much space is there for additional resources? If it's an important area that few people are working on, there may be low-hanging fruit, making the marginal pound spent more effective. Neglected tropical diseases are such an example. As their name suggests, they have received less attention as they only affect people in lower-income countries, meaning the incentives for privately funded researchers to develop treatments are low. While scale and solvability are uncontroversial, neglectedness should be interpreted with more care—it can rely on the notion that there are diminishing returns, and this is not always the case. A prime example is vaccines, whereby additional resources may have increasing returns up to the point at which we achieve herd immunity, and eradication becomes possible (so, this principle can be in tension with scale). Also, there may be good reasons why problems are neglected (e.g., they aren't solvable). Nevertheless, thinking about which problems are under-resourced relative to their importance is a useful framework and can surely help the UK to think about its unique selling point as a global development power.

Long-termism

Finally, you can't spend too long reading effective altruist material before encountering the idea that generally, we drastically undervalue the longer term in our decisions. There will be many, many more people in the future than there are currently alive (if all goes well) and there is no solid ethical reason to value their welfare less than our own (even if we discount for other reasons, such as uncertainty over consequences of our actions increasing in time or probability of greater resources in the future). So, we should be thinking not just about what matters now but also what might matter in the future. This does not mean forsaking any project that doesn't have a permanent impact—alleviating suffering today is clearly important—but recognising that interventions with permanent effects have the potential to affect many more people. For example, we should not give up on trying to remove barriers to growth and developing state capacity. Even if we can be less certain of success (it is hard to run RCTs to evaluate such programmes), the expected value could still be high given the long-run effects of success.

HOW DOES THE UK MEASURE UP?

The UK—and DFID in particular—has long been well-regarded for its use of evidence, its transparency, and a principled focus on the poorest or most fragile regions. All new spending projects require business cases that set out a clear framework for how impact will be achieved, and attempts to quantify the benefits and costs (though they have on occasion exaggerated expected results, and we expect the aid watchdog to pay close attention to claimed results).

Another UK strength is in being the largest contributor of official development assistance (ODA) to the multilateral system in absolute terms, recognising the benefit of collaboration and achieving economies of scale. With the multilateral aid review, the UK leads evidence-based assessments of multilaterals—such as the World Bank—for effectiveness, which can inform allocation decisions of all donors.

The UK has used its influence in the multilateral sphere to push for higher cost-effectiveness and greater use of evidence, for example in helping shift humanitarian responses towards cash transfers. Just as effective altruists advocate joining influential institutions where an individual's impact can be amplified, the UK has had an outsized impact on the effectiveness of the development system through its interaction with multilaterals.

DFID has funded high-risk, high-return projects which have had a huge expected value. For example, a joint DFID-Vodafone M-Pesa project transformed the banking and payments system in Kenya. Such projects are not just likely to be neglected—as most donor agencies are risk-averse—but also have the potential for scale. This appetite for risk and scale return should be nurtured in the new department.

Even if in these respects the government has adhered to principles an effective altruist might endorse, we'd argue that there are clearly areas where the UK could move further in that direction:

1. **Consolidate around the best programmes.** Givewell, a charity evaluator in the effective altruism space, compares the impact that other charities have and recommends only a handful of charities that fall above a stringent effectiveness benchmark. Conversely, the UK has thousands of projects, which vary a great deal in terms of their effectiveness. Each project may be valuable in its own right, but with (effectively) a fixed budget that is too small to fund everything we might want, being "valuable" is not sufficient; the marginal impact of spending more on existing projects might be higher.

   There are good reasons to not focus solely on a handful of the most effective projects: calculations of cost-effectiveness might be wrong, and funding a wider array of projects reduces the risk of having no impact. It may also generate ideas about where future best performing projects may come from. But it's hard to believe that the current array of projects reflects an optimal allocation, rather than internal politics, path dependency, and a desire to be seen as doing everything everywhere. In 2018, the average size of bilateral aid disbursement was around £1 million, but the median was only £139,000. DFID's chief economist points out that the most effective interventions can be an order of magnitude better than other "good" interventions, so we should cut out the latter. Of course, the best programmes may not always scale as effectively. But where there are order-of-magnitude differences in effectiveness, we should at least try.

2. **Focus on key partnerships and aid orphans.** The same logic applies to countries. In 2018, 135 countries received UK aid. When we focus on cross-border aid (omitting debt-relief, scholarships, refugee costs, and other administrative costs), there were 124 countries, ranging from Pakistan receiving around £330 million, to Vanuatu receiving £2 thousand. Fifty-five countries received under £1 million in project-aid in 2018. Most of this was miscellaneous spend by the Foreign Office in middle-income countries "in line with UK objectives." Perhaps these are high-priority projects for recipient countries, but it may have far more impact to consolidate this long tail of expenditures into a smaller group of countries and focus on proven, cost-effective programmes.

   As the UK prioritises its partners, it should also consider those that are relatively neglected. For example, Nigeria received just $18 for each person in extreme poverty from all donors combined, while other countries like Rwanda receive $100 for each person. There may be good reasons for such disparities, but considering global aid allocation as a whole suggests some countries fall through the gaps, not having strong colonial or geographic ties. We'll be looking more at this in the coming months.

3. **Don't be enticed by illusory "win-wins."** Attempting to kill two birds with one stone is seductive, but more often what happens is that effectiveness in one area is compromised by attempts to simultaneously address another, to the detriment of both. This was illustrated by our CGD colleagues in the case of climate mitigation and poverty. For the same cost of a climate mitigation/poverty reduction programme in Tonga, the Green Climate Fund could have reduced the same amount of CO_2 by spending a fraction of the amount elsewhere—and had enough left to give everyone in Tonga $450.

   Tagging additional objectives onto clearly focussed programmes is a path to ineffectiveness across all of them. To get a bit more technical: multi-objective optimization is generally harder than single-objective optimization. If there's a genuine synergy (for example, on air quality and climate), there should be evidence to support it.

4. **Use and lead through multilaterals for economies of scale.** There are several international organisations that consistently score well on measures of aid quality, such as use of evidence, systems for evaluating projects, and focus on the poorest. The UK is already the largest user of the multilateral system in absolute terms, and this allows it to benefit from outsized influence: enabling it to improve the cost-effectiveness of spending far greater than its own. As well as pressuring the humanitarian system to move towards cash transfers as mentioned above, and it has pushed for better measurement of impact for the Global Fund. However, the UK could still do far more to hold such funds to the performance commitments they make, and to encourage them to produce real outcome data that can allow their effectiveness to be assessed more easily. An ICAI report found that "DFID needs to hold agencies to account less for how they work and more for what they achieve." A continued and expanded role in the multilateral system used to push for these changes could be the UK's most effective way of having impact. Multilateral spending is also an efficient way of spending aid, requiring little overheads or feet on the ground, benefitting instead from the reach of large organisations such as the World Bank.

   With the UK's withdrawal from the European Union, a major part of its multilateral spend will cease. It would be disappointing, and detrimental to the UK's effectiveness, if this heralded a broader withdrawal from engagement with multilaterals, where the UK could amplify its values.

5. **Rethinking research.** Longer term, new research and technology is likely to be an important factor for developing countries' ability to prosper. So the UK's large R&D aid budget—now around 10 percent of bilateral aid—holds promise. But an effective altruist would want it spent in ways that maximise impact. The way that most R&D is allocated currently—scattered among universities with little transparency and coherency, and often little to ensure a development focus—does not fit this description. The new government's manifesto suggests a "a new agency for high-risk, high-payoff research"—learning lessons from Advanced Research Projects Agency (ARPA) models—and more recently this has been made explicit in the UK R&D Roadmap. This thinking should be applied to aid too. In recent work we've discussed instruments that might increase the productivity of research aid, and several ODA R&D pots in need of serious reform.

BIGGER BUDGET, BIGGER BETS: HOW GOVERNMENTS CAN GO FURTHER THAN INDIVIDUALS

Reviewing GiveWell's recommended charities, one might be struck by the preponderance of "micro" interventions: delivery partners implementing specific programmes such as delivering antimalarial drugs or support for de-worming initiatives. These are important, and given they tend to be relatively easy to randomize, the evidence for their effectiveness is strong. But these recommendations are primarily aimed at individuals. Effective altruists focus on the margin: where should the next pound be spent? This is the right approach, but in this context, it does not make sense to treat the government just like another individual, as the scale is completely different, and this opens up new opportunities for effective assistance. (Some effective altruists argue that even individuals should move in the direction of trying to exploit scale, using such devices as "donor lotteries.")

We have evidence that investments in energy infrastructure, transport infrastructure, and market integration have an impact on growth, and even if costs are also high, some estimates nevertheless suggest high benefit-cost ratios. But these are likely to need the scale, coordination, and political leverage that the government can provide. So while micro-interventions should be part of the government toolbox, so too should attempts to bring about systemic change. There is clearly still a place for high-probability, (relatively) lower-impact interventions—those affecting individuals and backed by gold-standard evidence. But by focusing only on these and ignoring low-probability, high-impact interventions, aid will not live up to potential longer-term impact. If such interventions don't feature on GiveWell's top charity list (for example), this is not because they are inconsistent with effective altruist thinking, but because the target audience is different. Effective altruists are clear that systemic thinking is crucial.

EFFECTIVENESS SHOULDN'T BE LIMITED TO AID OR ALTRUISM

Our colleagues have suggested that "impact" could be a general principle for the new FCDO. But the government should also bear in mind that some of the policies that could have the biggest impact on development lie outside the purview of the new organisation. The government has made clear its intention to see how aid could be used to address development and the national interest simultaneously. It should also ensure the reverse: that national policies are conducive for international development.

For example, research into antimicrobial resistance will be crucial globally and undoubtedly deserves attention in the new ARPA-style R&D agency the new government covets (i.e., an agency modelled on the US's successful Advanced Research Projects Agency). Switching to greener energy reduces our dependence on oil imports and could mitigate the scale of catastrophe awaiting developing countries on our current trajectory. Generally, some estimates suggest that the benefit-cost ratios of good policies (such as reducing trade restrictions) dwarf those achievable with aid. As the recent outbreak has shown, investing more in a global system of pandemic preparedness is as important for the UK as for other countries. Our colleagues have pointed out the importance of being driven by impact, whatever the objective.

It is a common misconception that effective altruism is only about giving money. In practice, impact can be achieved through numerous means, such as through one's career—for example, the way that people pursue their own interest (making a living) can also bring about positive impact. The government should take note, and think more about impact on development from policies beyond aid.

BUT THE UK COULD DO MUCH MORE TO BE EFFECTIVE

The Foreign Secretary has confirmed that the UK will retain its commitment to the poorest and conflict-affected countries and ensure that aid is *"invested in a way that can deliver the most effective results for the strategic objectives of alleviating poverty for the most vulnerable and delivering on climate change and on the wider international agenda."*

If the UK is going to spend aid effectively, then it should move further in the direction of effective altruists: responding to evidence about where and how aid is best spent and by which agencies, focusing on issues of scalability, and thinking about where its unique strengths lie, instead of trying to do everything everywhere. Crucially, a government can act at much greater scale than most altruists—it can avoid the caricature of focussing only on randomised controlled trials—and use its scale to lead systemic change.

We are very grateful to Dr. Hauke Hillebrandt of Let's Fund, Ranil Dissanayake, and Charles Kenny for comments on a draft. All views and errors remain our own.

Center for Global Development

Ideas to Action: Independent research for global prosperity

WWW.CGDEV.ORG

EUAN RITCHIE is a research associate at the Center for Global Development.

IAN MITCHELL is co-director of development cooperation in Europe and a senior fellow at the Center for Global Development.