

Opinion

Caroline Fiennes and Ken Berger



Caroline Fiennes is founder of Giving Evidence. Email caroline.fiennes@giving-evidence.com



Ken Berger is managing director of Algorhythm. Email ken@algorhythm.io

'Oops: we made the non-profit impact revolution go wrong'

The non-profit 'impact revolution' – over a decade's work to increase the impact of non-profits – has gone in the wrong direction. As veterans and cheerleaders of the revolution, we are both part of that. Here we outline the problems, confess our faults, and offer suggestions for a new way forward.

Non-profits and their interventions vary in how good they are. The revolution was based on the premise that it would be a great idea to identify the good ones and get people to fund or implement those at the expense of the weaker ones. In other words, we would create a more rational non-profit sector in which funds are allocated based on impact. But the 'whole impact thing' went wrong because we asked the non-profits themselves to assess their own impact.

There are two major problems with asking non-profits to measure their own impact

Incentives

The current 'system' asks non-profits to produce research into the impact of their work, and to present that to funders who judge their work on that research. Non-profits' ostensibly independent causal research serves as their marketing material: their ability to continue operating relies on its

The principle of self-assessment is flawed; it is easy to gain a distorted view of your own effectiveness.



persuasiveness and its ability to demonstrate good results.

This incentive affects the questions that non-profits even ask. In a well-designed randomized controlled trial, two American universities made a genuine offer to 1,419 microfinance institutions (MFIs) to rigorously evaluate their work. Half of the offers referenced a real study by prominent researchers indicating that microfinance is effective; the other half referenced another real study, by the same researchers using a similar design, which indicated that microfinance has no effect. MFIs receiving offers suggesting that microfinance works were twice as likely to agree to be evaluated. Who can blame them?

Non-profits are also incentivized to only publish research that flatters: to either bury uncomplimentary research completely or share only the most flattering subsets of the data. We both did it when we ran non-profits. At the time, we'd never heard of 'publication bias', which this is, but were simply responding rationally to an appallingly designed incentive. This problem persists even if charity-funded research is done elsewhere: London's respected Great Ormond Street Hospital undertook research for the now-collapsed charity Kids Company, later saying, incredibly, that 'there are no plans to publish as the data did not confirm the hypothesis'.

The dangers of having protagonists evaluate themselves is clear from other fields. Drug companies – who make billions if their products look good – publish only half the clinical trials they run. The trials they do publish are four times more likely to show their products well than badly. And in the overwhelming majority of industry-sponsored trials that compare two drugs, both drugs are made by the sponsoring company – so the company wins either way, and the trial investigates a choice few clinicians ever actually make.

Such incentives infect monitoring too. A scandal recently broke in the UK about abuses of young offenders in privately run prisons, apparently because the contracting companies provide the data on 'incidences' (eg fights) on which they're judged. Thus they have an incentive to fiddle them, and allegedly do.

Spelt out this way, the perverse incentives are clear: *the current system incentivizes non-profits to produce skewed and unreliable research.* ▶

Resources: skills and money

Second, operating non-profits aren't specialized in producing research: their *skills* are in running day centres or distributing anti-malarial bed nets or providing other services. Reliably identifying the effect of a social intervention (our definition of good impact research) requires knowing about sample size calculations and sampling techniques that avoid 'confounding factors' – factors that look like causes but aren't – and statistical knowledge regarding reliability and validity. It requires enough *money* to have a sample adequate to distinguish causes from chance, and in some cases to track beneficiaries for a long time. Consequently, much non-profit impact research is poor. One example is the Arts Alliance's library of evidence by charities using the arts in criminal justice. About two years ago, it had 86 studies. When the government looked for evidence above a minimum quality standard, it could use only four of them.

The material we're rehearsing here is well known in medical and social science research circles. If we'd all learned from them ages ago, we'd have avoided this muddle.

Moreover, non-profits' impact research clearly isn't a serious attempt at research. If it were, there would be training for the non-profit producers and funder consumers of it, guidelines for reporting it clearly, and quality control mechanisms akin to peer review. There aren't.

Non-profits should use research rather than produce it
Given that most operating non-profits have neither the incentives nor the skills nor the funds to produce good impact research, they shouldn't do it themselves. Rather than produce research, they should *use* research by others.

So what research should non-profits do? First, non-profits should talk to their intended beneficiaries about what they need, what they're getting and how it can be improved. And heed what they hear.

Second, they can mine their data intelligently, as some already do. Most non-profits are oversubscribed, and historical data may show which types of beneficiary respond best to their intervention, which they can use to target their work to maximize its effect.

Put another way, if you are an operating non-profit, your impact budget or impact/data/M&E people

First, non-profits should talk to their intended beneficiaries about what they need, what they're getting and how it can be improved. And heed what they hear.

probably shouldn't design or run impact evaluations. There are two better options: one is to use existing high-quality, low-cost tools that provide guidance on how to improve. The other is to find relevant research and interpret and apply it to your situation and context. A good move here is to use systematic reviews, which synthesize all the existing evidence on a particular topic.

For sure, this model of non-profits *using* research rather than *producing* it requires a change of practice by funders. It requires them to accept as 'evidence' relevant research generated elsewhere and/or metrics and outcome measures they might not have chosen. In fact, this will be much more reliable than spuriously precise claims of 'impact' which normally don't withstand scrutiny.

What if there isn't decent relevant research?

Most non-profit sectors have more unanswered questions than the available research resource can address. So let's prioritize them. A central tenet of clinical research is to 'ask an important question and answer it reliably'. Much non-profit impact research does neither. Adopting a *sector-wide research agenda* could improve research quality as well as avoiding duplication: each of the many (say) domestic violence refuges has to 'measure its impact', though their work is very similar.

Organizations are increasingly using big data and continuous learning from a growing set of non-profits' data to expand knowledge on what works. As more non-profits use standardized measures, they can make increasingly accurate predictions of the likelihood of changed lives, and prescribe in more detail the evidence-based practices that a non-profit can use.

In summary

Non-profits and donors should *use* research into effectiveness to inform their decisions; but encouraging every non-profit to *produce* that research and to build their own unique performance management system was a terrible idea. A much better future lies in moving responsibility for finding research and building tools to learn and adapt to independent specialists. In hindsight, this should have been obvious ages ago. In our humble and now rather better-informed opinion, our sector's effectiveness could be transformed by finding and using reliable evidence in new ways. The impact revolution should change course. 