HOW DO WE DRIVE PRODUCTIVITY AND INNOVATION IN THE CHARITY SECTOR?

Lecture, hosted by the Royal Society for the encouragement of the Arts, Manufactures and Commerce (RSA), 8 May 2014.

Dan Corry

Introduction

Imagine you have a sector with an annual income of £60bn and nobody really knows what it achieves or scrutinises it very much.

And then imagine it is made up of 160,000 registered bodies, many enjoying rather handy tax-breaks, who are equally under-scrutinised, and employs over 700,000 people.

You might be a bit taken aback.

Well this is, broadly speaking, the charity sector. And this evening I want to talk about why it is left so much to its own devices and whether there is anything we can do about it, drawing on the way things work in the private and the public sector.

And I will touch on perhaps a more controversial topic, which is whether we should do anything about it.

The charity sector

Everyone involved in the not for profit or charity sector is driven by a sense of mission and passion. That is what makes it such an exciting and worthwhile place to be.

We all give money to charity for the same reason: we want to do good.

We donate the proceeds of our company, now sold as we go into retirement; or money from our family inheritance. We give a fiver by standing order once a month or sign up for a sponsored cycle ride.

We need not do it.

A similar thing applies to those setting up or working for charities and voluntary groups. Many could be doing something else, or enjoying more leisure time or getting paid better in a public or private sector job.

Yes, many in the ‘voluntary’ sector get paid (which some diehards feel lets the side down), but very few are involved to maximise profits, improve share prices or earn as much as they can.

Now, of course, this is not the story everywhere. There are those who give to increase their chances of meeting Prince Charles or to get a knighthood; those who are corrupt or looking for tax advantages; and those who run what I increasingly call ‘business charities’, where social purpose is one among other objectives. But these are the exceptions.

For most, it is about passion, mission, and doing good. And what fantastic motivations.

But good motives do not guarantee good outcomes, let alone as much good as you could achieve.

So when people say we should leave charities to it, satisfied by their objectives and their selflessness, and stop analysing them, I just cannot agree.

Yes we should admire charities for their dedication—but I don't think we should be happy with simply 'meaning well': it does not cut it.

Let me propose, instead, that we need some "creative destruction" in the charity sector.

There is some complacency out there, as I will go on to outline, just as there is some very good practice. If we want to identify one from the other, and to encourage one over the other, then things are well overdue a shake-up.

Charities seldom think about how effective they are at turning charitable funding and donated time into social outcomes. They are not generally under pressure from funders to start doing so, either. But this is surely what they are all about.

And the way they achieve this is broadly through what other sectors call productivity. Yet you rarely, if ever, hear this phrase used. It is little more than a four-letter word, as though sacrilegious to the very ethos of what it is to be involved in charities.

Yet productivity, and its close cousin innovation, are crucially important. These two things are about using resources in the best way possible. The alternative is to accept achieving weaker outcomes for the people and causes you care about.

And this needs to be dynamic. We don't just want every charity doing the best it can with the way it works today; we also want innovation.

This means constantly searching for new and better ways of, for example, organising the back office; or delivering interventions; or using social media for mentoring; or apps to help increase knowledge about constitutional rights in developing countries.

Impediments to productivity in the charity sector

So, what is holding the charity sector back?

Well, while passion and a fantastic sense of mission do drive activity, this is also an impediment to doing our very best at both the level of individual charities and for the sector as a whole. The desire to give and do often overwhelms the desire to achieve and constantly improve.

But many other factors also play a part and hinder the creative destruction that is needed. I will focus on six today: funding, feedback loops, competition, governance, measurement and collaboration.

1. Funding

First, is the constant search for funding—the obsession of every charity chief exec.

This ought to put an onus on charities to prove they do their job well. However, in practice, it mainly seems to drive innovation in fundraising techniques!

As former Shell Foundation and Institute for Philanthropy boss, Kurt Hoffman, recently put it: *'when the best and brightest working in the charity sector are spending so much time fundraising, this happens at the expense of* *devoting conscious effort to innovating and improving performance—which is by far the most important source of value creation in every other product sector in our economy.'*

But the problem on the funding side goes further, since most private voluntary sector funders are not really trying to maximise return—in this case, social return—for a given level of resources.

They do want to have a positive social return, and may undertake brief checks to ensure money is not being totally wasted, sequestered or otherwise causing harm.

But this is very different from searching out the most efficient organisations or innovative parts of the sector and putting money there. So there is limited pressure on charities from 'investors' looking to invest in the best.

Charities also suffer because of the tightness of the funding they receive and the fact that much is restricted to particular uses. Donors unhelpfully look at metrics to do with the percentage of money going straight to the beneficiary, which militates against longer-term developments. All this restrains charities from having the capacity to think about productivity and innovation.

2. Feedback loops

The funding regime contributes to a second problem of weak 'feedback loops'—from funder to charity and vice versa, as well as between charities and their beneficiaries.

Think of it like this.

In most walks of life you know pretty well if your performance is not up to scratch. The guy selling coffee for £3 a throw soon feels the power of a feedback loop when his customers report that his coffee tastes horrible and take their custom elsewhere, and his investors pull out because he is losing trade and money.

For charities, the person who pays isn't the person who consumes, and charitable funders rarely move their funding because the charity is not doing as good a job as another.

All this means that if we think of the charity sector as a system, there are no feedback loops driving us in the right direction. There is no invisible hand of market forces—where beneficiaries go to a competitor if unhappy or funders support a clearly 'better' charity—to help calibrate the system in real-time towards improving productivity and innovation.

Nor does it happen as a result of citizens and users saying what they think very clearly and publicly, as they might in the public sector.

With honourable exceptions and despite the best efforts of many great charities, "voice", not very powerful in the public sector, is often very faint in the charity world, especially important when the "client" can't really go anywhere else.

In the absence of all this, charities focus on working out what the funder wants so as to guarantee their own survival, taking attention away from what they are achieving and what their beneficiaries need and want.

We at NPC often hear from charities about funders giving in haphazard ways, asking for slightly different information to another funder for no good reason, not giving enough certainty on future funding and so on. But this rarely reaches the ears of funders because charities are keen to maintain their support.

As Stanford Universities' Rob Reich puts it: *'people who become foundation officers are often surprised to find themselves transformed into the smartest and best-looking people in a room.'*

3. Competition and the process of entry and exit don't work to drive productivity

This contributes to a third problem for the sector: too many entrances and not enough exits.

The process of new charities starting up and others closing their doors does not work well in the sector; and when it does happen, it is not clear that the worst are replaced by new, better charities.

Why do too few charities exit?

Because it is possible to continue even if things are not going that well. As Professor Kimberley Scharf of Warwick University points out, although people start charities to solve a particular problem, they tend to try and keep them going even if another charity is doing a better job. The so-called "warm-glow" effect that you get from doing things via your "own" charity outweighs the original motivation of doing good.

So you will only set up a donkey sanctuary if you really care about looking after donkeys. But will you ever close down your donkey sanctuary, voluntarily, if on all objective measures the sanctuary down the road is doing a better job?

It is unsurprising, then, that mergers—never easy in any sector—are particularly rare in the charity sector. Without incentives to maximise value—in this case, social value—it's not hard to see why mergers rarely come to pass.

At the other end of the lifecycle, it is fairly straightforward to launch a charity; there are few barriers to entry and often little need for much start up capital. There is often less due diligence than a business start-up would get from a potential lender, so if someone is happy to give you the money, you are basically good to go.

It then becomes harder to scale up for a whole set of reasons, so charities often operate below optimal size.

One knock-on effect is that the bigger players have little reason to fear being pushed out by a bold new entrant, in the way that the arrival of Microsoft forced IBM to completely transform.

This guides the argument that, at a sector level, we waste resources in aggregate because we have too many small charities. Of course, this needs to be set against the power of smaller charities, voluntary and community groups in dealing with specific issues; creating social capital; and securing the full commitment of their staff and supporters. But the basic point still stands.

Competition of the traditional neo-classical kind does not help us much in this sector.

4. Governance

So if these mechanisms do not work, are we rescued by governance?

Well, here we come to a fourth issue.

Governance in charities works in a very different way to other sectors. Questions about how the charity is performing and improving often get ignored in favour of topics which feel more urgent: what do our spread sheets show us about the next month? Does our spend all come within Charity Commission guidelines? Are we financially sustainable?

Add to this the voluntary, unpaid nature of trustee positions and the often unclear metrics around good performance, and the net result is that there is usually little pressure from the trustees of a charity (or of a charitable funder) to be high up what economists call the production possibility frontier or to be pushing the boundaries on innovation.

5. No decent metrics

The fifth systemic problem underpinning a number of these other issues, is that it is very hard to assess productivity in this sector. We do not have—and probably cannot have—an agreed measure of output akin to profitability, earnings per share or even GDP in the private sector. This lack of an agreed numeraire or metric is a major problem and one I will return to.

6. Collaboration is hard

Finally, and as consequence of some of the factors already mentioned, collaboration is almost viscerally discouraged. This makes the sector much less productive than it could be overall.

Most social issues are stubborn and complex, and finding a solution requires scale and concerted action from different groups working together. This means collaboration among charities and funders; yet we see very little of this in practice.

The intense competition for resources creates pressures to bilateral relationships between funders and charities, and therefore a reduced appetite for comparative assessments and cooperation.

What is our response?

I have explained a number of barriers to productivity and innovation.

So, what is our response?

One is of course, to give up on this quest. Maybe even to say that we do not care. The sector is about plurality; giving voice to marginalised people; the glue of civil society. It is about passionate people doing things of their own volition, outside the clutches either of the state or private for-profit interests. Surely charities have a right to exist and for funders to bank-roll them: as long as they do not cause harm, they should be left alone to act as they like.

In any case, when all is said and done, the sector is a bit marginal, a “nice to have”, and we should not beat ourselves up about its weaknesses.

But this easy answer is unacceptable.

Of course, the overall productivity of the charitable sector includes many things—going beyond the easily measurable aggregate impact of each charity. But this is not an excuse for each organisation to be less productive than it could be.

I don’t accept there is a necessary trade-off. The passion that is core to charitable works does not disappear just because we try and work as well as we can.

And the need is more urgent just now.

As the state pulls back, for good and bad reasons, from helping some of the most vulnerable in society, as well as funding some of the more “peripheral” areas of life like the arts and sport, the voluntary sector must deliver as much as it can with limited resources.

In addition, we must not forget that many charities and community groups receive public money, not only directly, but through important tax advantages. This carries responsibilities; resources are being given to the sector and we owe it to the public to use these resources in the best way possible.

So the search for doing better must be a priority.

Comparing sectors

Although I now work in the charity sector with NPC, I have experience thinking about these issues in other sectors. I want to consider if any of this could help us.

Private sector

Let me start with what drives private sector productivity and innovation—something I thought long and hard about especially in my years running the economics team at think tank ippr and then as a special adviser at the Department for Trade and Industry.

First, the dynamics of the market dictate that the best companies generally rise to the top.

It's hard to hide behind poor performance.

If you rest on your laurels, refuse to innovate, do not take up the latest technology or respond to the changing tastes of your customers and clients, or price keenly for what you produce, then someone else will take your place. You may begin by losing or gaining market share, but this will ultimately lead to what economists call "Entry and exit". In general, the "worst" firms leave the market—go bust, in normal language—while new ones replace them.

To put this another way: if you want to know why no one goes to Wimpy anymore, look no further than McDonalds and Gourmet Burger.

Work by Jonathan Haskel at Imperial College shows that this is incredibly important for productivity. It is not just about the weak leaving in a sort of Darwinian, X-Factor way, but to do with the fact that productivity is often developed by bold new entrants who smash the previous paradigm by bringing in new approaches, new innovations. These are the cutting edge start-ups we should want to see coming through—the charities with big ideas and a commitment to measuring how well they are doing, and publicly sharing their knowledge of what works and what doesn't.

Second, shareholders put their money at risk and want to see a good financial return, so they have strong incentives to push the firm as far as it can go.

And finally, all firms operate according to a general infrastructure, by which I mean education, training and physical infrastructure, often provided by or enabled by the public sector, but also agreed and policed standards, metrics and regulatory structures.

All these elements, external to the firm itself, give the right signals and enable innovation.

This is hard stuff to reproduce in the non-profit sector.

One answer is to say that we need charities and others to behave more like their private sector counterparts. This explains why some are attracted to more commercial segments of the sector: social enterprise; social and impact investing; and 'for profits with a social purpose'.

Here we have different governance structures and encounter more of a commercial mind-set, which means that, on the whole, they are better on risk, more customer-facing, and often take loan or quasi-equity type funding, which they have to pay back, bringing an element of discipline lacking when funding is via pure grants.

The question still unanswered is whether social enterprises and the like are really an alternative to for-profit outfits, and not an alternative to charities? I will leave that thought hanging.

So instead of capitulating to it, let's see what we can draw from the private sector experience.

I previously mentioned that infrastructure matters in the private sector—and government sees a role for itself in helping. It would be nice if it did that in the voluntary sector too—a few grants and capacity-building projects here and there don't go very far!

In government, the business department and the Treasury worry about private sector productivity endlessly, while academics probe for better insights and ideas for action. There is a similar concern for public sector productivity.

But almost nobody really cares about it in the charity sector.

Impact

More profoundly, the private sector dynamic works partly because there is an ability to compare firms and their products and services, and allocate resource accordingly, whether you are a consumer or investor.

In the charity sector, the inability to know much about productivity and other metrics that allow you to see how a charity is doing is inevitably a great stumbling block. At its worst, charities have little or no idea of the effect they are having on the people they exist to help.

Given the lack of internal forces pushing towards productivity, the concept of impact has grown up in the charity sector to try and provide grit to those looking to put productivity centre stage.

Impact is not productivity, but it acts as a kind of proxy because it focuses on results and so allows you to work out the cost per unit of impact.

Impact at the level of the individual charity starts (or should start in NPC's view) with what is known as a theory of change—a process of backwards mapping from your ultimate goal to the outcomes and activities needed to achieve it.

After this, you can start to get a handle on how to measure what you achieve. This may start with the outputs you produce along the way: although not outcomes per se, even counting this can be a big step forward for many, many charities.

Measuring "softer" outcomes becomes harder still, but you can usually have a go. And even if you are an advocacy or campaigning organisation, it still makes sense to think about what doing a good job would look like and measure yourself against that.

Then there are issues around how much you actually contribute towards that outcome. This introduces the language of attribution, additionality, deadweight, displacement, and substitution.

Purists say you can only address any of this with top of the range methodology, copied from the medical world—randomised control trials or RCTs. But these are not only expensive and often disproportionate, but can only get you so far in social policy.

While a beautifully proven pill that works in one setting will probably work in another, an intervention for loneliness that works for a retired urban professional, may not work for a retired rural farm labourer.

When you have assessed your impact, it becomes possible to look at trends across time, and compare across different settings (youth clubs, schools, towns, cities, prisons) and against other interventions and organisations.

All of this is vital to any charity that seeks to do better; if you don't know how you are doing and how this is changing over time, you can never learn or improve.

But it also starts to allow the external pressures that exist in the private sector to play some sort of role in the non-profit sector.

Much of this work is up to the sector itself. Charities need to up their game when it comes to the measurement agenda, and act on what it tells them, even when tough to stomach and resisted by trustees, volunteers, funders and even some of the user group. But so too do foundations, and their representative bodies, plus the philanthropy advisers in private banks, in accountancy and in legal firms.

And if funders insisted on more impact measurement and made funding decisions based on whether a charity was the best in class and/or improving, then we would have a mechanism that drove productivity and effectiveness.

Certainly there are, and always will be, funders who want to do minimal checking and give money out in a reactive way, but many are at least asking the questions and at times even helping to support evaluation work.

But government can also help press the case for decent measurement.

It can support work to proselytise for impact measurement. It has already done this with the Inspiring Impact programme—where NPC leads a wide ranging group of sector bodies building support for and tools to enable good impact measurement; and for the G8 [or is it now G7?] Impact Investing Task Force work on measurement of social impact—a first stab at trying to draw up some international norms—which NPC are deeply involved in.

More indirectly, government can insist on some kind of impact measurement as it gives out money. And it can open up its warehouse of data to the charity sector to help them find out what happened to those they worked with—or, even more excitingly, create a control group.

The work we did to help launch the Ministry of Justice's excellent Justice Data Lab is world leading and may turn out be more fruitful and important than the plethora of What Works Centres the government is currently encouraging.

Government and other funders must also think about how their own activities, including their commissioning practices, affect the ability of individual charities and the system to have impact.

I think it would be a salutary kick up the pants if they were required to publish a report on how they tried to do this each year.

So there are steps we can take with concerted effort from charitable funders, commissioners, data holders, the evaluation community and from government to make impact measurement the norm. This would undoubtedly drive improvement in the sector.

Public sector

What might we then learn from the public sector in our quest for improving the charitable sphere?

Nothing, you may think, since its record on productivity and innovation is not always sterling. Anyone who has worked in public service knows the danger of admitting this to their barber or hairdresser—a litany of complaints about bus times and post office queues soon follows.

But this is unfair and unhelpful.

I have been intrigued and frustrated by public sector productivity as I have worked across a whole range of services from schools to transport, including in my spell in Downing Street.

Here the drivers of productivity, innovation and progress are closer to those in the charity sector, since the main ways have to be ‘non market’.

First, in public services, a number of incentives stem from democratic accountability. Politicians at a local and national level do—in my experience—genuinely want to make services better for citizens and to do so cost effectively. They also know that if they do not they are at risk of being turfed out of office by the electorate.

The charity sector cannot copy this straight. But it could improve its governance by putting its aim to improve services cost effectively up in lights, and by having perhaps better links with user groups.

The Charity Commission could help push trustees in this direction too, by demanding more information, and making clear that improvement is part of their role.

In the public sector, faith has rightly been put in the so called "Public service ethos"—the belief that those working in the public sector just want to do good (sound familiar?). Over the years, however, people have become suspicious. Was it powerful enough? Did all public sector workers really have a selfless desire to do good? Did they not often put the needs of themselves, the producers of the service, above those of the users?

So with waning faith in the idea that services would naturally improve through internal pressures, a number of different elements came in to drive improvement, innovation, modernisation and productivity.

One is the attempt at benchmarking, which allows some comparison of costs per unit or user satisfaction across organisations and activities, either informally or through mechanisms like the now scrapped Audit Commission. There seems little reason—in either principle or practice—not to apply a similar approach to the charity sector.

a. Competition among charities for esteem

More dirigiste were the top-down approaches, which encouraged public services to get better through targets, with rewards for hitting them and so on. I think there is something we can learn from this too.

One of the more attractive, top-down approaches I've long thought about would be a "star rating" system to help drive what I call competition for esteem.

Let me outline my thinking.

Several times in my career, I worked on versions of assessment and league tables. They were not always popular with those subjected to them and involved arguments about whether they measured the right thing, accurately and took into account all of the relevant factors.

But more interesting, is the way they worked to drive improvement and innovation.

League tables and tough Ofsted inspections and reports introduced for schools were coupled with greater choice mechanisms and a view that parents could use the data to decide where to send their children. With money following pupils, you had a quasi-market system all ready to go.

And while on the margin there was some of this, the strongest effect was on heads and governing bodies—and their pride. They did not want to slip down the league table or have to face colleagues judged better than them. It was a real driver and had real effects—perhaps the ultimate "nudge" approach.

Of course, it sometimes went further than a nudge, with heads likely to face the chop for failing.

But the mechanism of comparison did drive the system and achieved it through the effects it had on providers and bosses rather than through consumers and users responding to what the comparisons showed. Either way, as analysis by Professor Gwyn Bevan and colleagues at LSE shows, it certainly seemed to work to raise standards.

The Comprehensive Performance Assessment (CPA) to establish a star rating system for local councils worked in a similar way. It was not that citizens voted on the basis of whether their council got a one or a three star rating, but that chief execs and councillors did not like to be worse than their peers. Chief execs were sacked if they failed and those who had achieved three stars elsewhere were “bought”, much like star footballers.

It all went crazy for a time, while councils tried to sort out the Gianfranco Zolas from the Eric Djemba-Djembas.

Clearly things are not exactly the same in the charity sector, but the lesson is that forcing greater comparison between charities on the basis of impact is likely to be a good thing, even if not entirely perfect, as it helps create “competition for esteem”.

We also need some sort of impact leadership group to show the way, create a bit of competition and put (implicit) peer pressure on non-performing boards and CEOs.

b. Contracts

Another way to sharpen our focus on productivity and outcomes in the public sector relates to contracts.

Separating the purchaser and the provider (even where both are from the public sector) and having an explicit contract between them was used to force efficiencies and other desirable outcomes. The ultimate end of this approach is the more recent move to payment by results—so that outcomes are specified and providers are pushed to compete over who can provide these outcomes most efficiently.

This kind of approach is starting to have some impact on charities, and other non-profits, especially those involved in public sector delivery. Despite the many dangers with contracts, and especially payment by results, there is clearly something useful in the focus on outcomes it brings.

Social impact bonds may or may not be a big player in the future, but they are at the very least making people think harder about outcomes, and working together and managing performance to produce these outcomes.

Social investment and impact investing is related to all this too, since the aim is to take a reduced financial return because you are looking to achieve a social return. Exactly how much most social investors are *really* going to be focused on the precise achievement of social goals rather than the financial part is a matter that is now starting to emerge.

c. Working together

Finally, in pubic services it is clear that everyone needs to work together if we are to achieve outcomes—however hard joined-up government proves to be. This is very different from the private sector where, broadly, each individual firm doing the best it can on its own leads to the best outcome overall.

The charity sector is much more like the public than the private sphere on this dimension, so it is vital that the search for impact does not stop at the doors of a particular charity.

The real aim is the impact we achieve overall and whether we can learn, improve and scale up. This means using collective theories of change for a whole area of work; common metrics and approaches; sharing information about evaluation on what works as well as what doesn’t; and collaborating and working together.

Conclusion

So while driving productivity in the non profit sector is a challenge, I have suggested some ways we can go forward.

Impact measurement, competition for esteem, more collaboration, greater pressure and support from the sector itself and from government—all would help.

The charity sector will never be like other sectors. Indeed its special role means we don’t want it to be.

But if we can get 20 or even just 10 per cent more out of the sector by doing what we already know better, and even more if we can get innovation motoring, then the world would be a better place.

It only needs 10 per cent of charities to step it up a bit, and suddenly 16,000 organisations are better equipped to help their beneficiaries, bring in new ideas, and innovate to survive. That’s likely to help hundreds of thousands of people who rely on those charities. And that’s the effect of just one in ten improving.

That is surely something worth fighting for.