

When Good Intentions Are Not Enough: a think piece on Knowledge Management in CARE International

Matt Bannerman, June 2014

This think piece has been developed as part of a collaborative program of action research, supported by CARE UK and involving colleagues from CARE DRC, CARE Rwanda and the East and Central Africa Regional Management Unit, exploring how the capacities and resources of CI members could be marshalled in a more targeted, intentional and systematic way to support the effective functioning of CARE's country programs and the oversight of programming and operations that they provide. All views expressed are, however, the author's own.

Why do we find managing our knowledge so hard? For many years, CARE has been aware of the need to improve in this area, and hundreds of internal and external reviews and evaluative processes at multiple levels have exhorted the organisation to improve the way that it promotes reflection and learning, and captures, stores and shares information and knowledge. More recently, KM has emerged as a central focus of the transformational change process towards the CARE 2020 vision and a CI Knowledge Management Working Group has been established to try and map a way forwards.

The persistence of the problem suggests however that this is not merely a 'technical' problem – but that instead it is a systemic and cultural problem. As CARE staff we have not developed the habits and behaviours associated with knowledge sharing because there are very limited incentives – positive and negative, carrot and stick – for us to do so.

Why intentions fail and incentives succeed

As human beings we all respond to incentives all the time, and the balance of incentives is critical in determining the choices and decisions we make – a fact that policy makers in the rich world seem to have discovered relatively recently, but that development organisations such as CARE have known for far longer, given that we frequently include behaviour change objectives and activities as part of our programming. Yet when it comes to the kinds of behaviours that we seek to promote in ourselves and our colleagues in order to improve the way we manage knowledge, incentives don't seem to figure much in our thinking. Instead, we seem to assume that communicating and explaining our KM strategy, reminding ourselves of the centrality of KM to our mission - essentially relying on the undoubted good intentions of ourselves and our colleagues - will work. Sadly, experience suggests that they haven't in the past and won't in the future. Below we offer some suggestions as to why, and what we might do differently.

There are costs to managing knowledge, and those costs are born in large part by individuals or teams when they invest time in either seeking and searching for knowledge and learning relevant to programs that they are currently designing, or in capturing and sharing knowledge and learning that may be relevant for others. Of course, such costs can be reduced through effective systems, such as a really user friendly online platform. They remain real costs, however, and the balance between those costs and the

benefits – in the form of the incentives that exist or that we create as an organisation for seeking and sharing knowledge – is what ultimately determines the choices that we make to turn good intentions into concrete actions. To use a simple analogy: at home we may have the best intentions to recycle our household waste, but it is the subtle balance of incentives – feeling virtuous, gaining peer approval, fearing financial penalties if we fail, and simply the level of effort required – that determines the extent to which we act on our intentions.

From altruism to expectation

So what are the incentives that work on us as we make our choices around KM? Blogger from private sector KM specialists KnoCo [Nick Milton](#) provides a useful list, from which we draw on in what follows. It is helpful to state from the outset that we are not equating incentives with financial rewards - interestingly, even in the private sector, these come a long way down the list, with widespread scepticism that monetary rewards can ever prove an effective incentive for knowledge sharing. Encouragingly, altruism comes in higher – a sense of ‘doing the right thing’, either directly to support a colleague, or behaving in a way that feels in line with a broader social goal or mission – for example, with our shared vision in CARE 2020. Those of us working in organisations with social missions such as CARE feel most comfortable with these kinds of incentives, and it is thus unsurprising that this is where much of our existing effort to promote KM is focused. But, as we noted above, though altruism can be an effective enabler, it rarely ‘clinches the deal’ – most of us need our altruism to be reinforced with other more powerful incentives before we make the right choice.

Similar but subtly different is the important driver of professional pride – in posting an evaluation report that shows the amazing results attained by our project, my team and I are publically sharing our sense of achievement. Of course, this implies a hazard that I will only share the good news, when in fact much of the most critical learning can be gained when things turned out less well, or differently than expected – this is a problem with our frequent focus on ‘success stories’ (of which more below).

More powerful still is recognition. Pride in my achievements may motivate me to post my findings in a public space – yet my motivation is further reinforced if I then get feedback from others in the organisation recognizing what my team has done, linking our learning with a broader body of knowledge, or providing new perspectives or ways of understanding and applying what we have learned¹. Sharing my ideas in a conversation with an engaged, informed community is naturally vastly more rewarding than shouting into a void – yet in the absence of any form of recognition, this is exactly what our efforts at knowledge sharing can feel like.

Getting a response when I take action is good, but what if I don’t? Sometimes I may need a little encouragement, and this is where the role of the informed Knowledge Manager can be critical: a

¹ The Minerva platform offers many opportunities for this, including the ‘like’ or ‘rating’ functionality and the ‘Twitter-esque’ Pulse IM function. Uploading a file to an online folder is not rewarding – but getting real time feedback from colleagues appreciating the content, asking further questions and linking to other resources is. This would also enable ‘crowd sourcing’ a measure for quality using the ‘like’ or ‘rate’ functionality, to make it easier for users to identify high quality material

colleague who takes note of, for instance, the date of an upcoming conference relevant to my work or learning milestones established within my project plan, and follows up at the appropriate time to nudge me to share outcomes, results and learning². Milton calls this the 'badgering' function, and almost all successful knowledge management enterprises use it in some form. This is a very difficult function to automate or distribute; instead this sort of 'guiding intelligence' is logically located within a technical team working across the organisation, able to create links and connections between innovative activities in the field and the global debates and dialogues around a particular sector or thematic area.

Finally, we come to expectation – when active seeking for and sharing of knowledge becomes the default way that we behave, to the extent that when it does not happen we notice and take action. This needs those of us in leadership roles to consistently and visibly communicate those expectations – both explicitly in what we say but also implicitly in what we do. Sadly the two can contradict and cancel out each other – for example if a senior leader writes an eloquent forward to a KM Strategy for her unit or team, but then does not mention KM performance in an annual appraisal of senior staff.

Metrics and accountabilities

Which brings us to another key issue – and gap – for KM: the absence of any widely accepted metrics by which we can measure the performance of teams or individuals. Flawed they may be, but the metrics of pipeline and burn rate are real and have meaning for all involved in resource mobilization and program implementation. Targets can be set and followed up, benchmark comparisons can be made across units and a (mostly) healthy sense of peer competition can be engendered. But where are the equivalents for KM, sitting as it does at the heart of our vision for our future organisation? How can we get to the stage where team leaders regularly say to their teams: 'our knowledge numbers for this quarter are fantastic – how can we keep it up and further improve?' Or, just as importantly 'we are lagging behind in our knowledge performance; what do we need to do to get back on track?' At the moment, if KM enters those management conversations at all, it does so in terms of what we aspire to: 'wouldn't it be great if...', 'I wish we could get better at...' To shift these conversations on to what we expect of ourselves and others, metrics are essential.

There is another problem here (which this think piece may inadvertently be aggravating) – in our long slog to improve KM across CARE, we have been successful in creating an (almost) universal acknowledgment that this is an organisational weakness – and in so doing weakened our individual accountabilities for doing something about it. As a CARE Country Director, I was in no doubt that meeting donor commitments and keeping within UNR budgets was my responsibility – but if KM came up in conversation, I could sigh and say: 'we really need to get better at this as an organisation.' I was quite sincere in saying this, as of course are many other concerned colleagues when they make similar

² Learning 'loose ends' or unclosed learning loops are common within CARE's knowledge base. Many reports or evaluations will include statements such as 'it is too early to assess X' or 'future studies should examine Y'. Yet operational factors that impose a short attention span, along with a knowledge ecosystem that is entirely supply driven, mean this is often where the trail ends. An intelligent knowledge management system should be able to prompt for follow up at the appropriate time in order to close the learning loop. Of course, establishing the conditions in which anyone will respond to such a prompt is the subject of the rest of this article.

statements, but in doing so we collectively dilute the accountability for taking action to the extent that it becomes lost in the many other things we intend to get around to some day when the in-box is empty, or hope that someone else might fix.

The (uncomfortable) importance of making judgements

When bemoaning the state of our KM, we often hear CARE colleagues say something along the lines of: ‘we do so much great work in the field, but we are hopeless at documenting it and sharing more widely’. This statement is both true and misleading. Let’s be clear: of course CARE does much wonderful, innovative, impactful work across the world. However, like any other organisation, there is variation: some of our work is merely good, a reasonable effort with the resources to hand, whilst it is an uncomfortable truth that some of our work will inevitably be poor, hurried, ill thought out and on occasion actively harmful. I’ve seen and been involved in all three, as has any other field manager with any length of service with CARE, or indeed any other development organisation. We are human, and working on complex challenges in a complex – and highly time and resource limited – world. As a thought experiment, if we could plot the distribution of our projects against some generic measure of ‘program quality’, we would likely end up with something resembling the familiar bell curve, with relatively small numbers of exceptional and poor outliers either side of a larger mass of solid, pretty good initiatives. In fact, private sector companies, particularly in the manufacturing sector, often plot exactly these kinds of curves and use them as the basis of highly sophisticated, statistically-driven quality control systems, of which Six Sigma³ is perhaps the best known.

Of course, at CARE we cannot manage the quality of our development work in the same way that Toyota manages the fit of its sun roofs, and nor should we try. The point of the bell curve idea is to illustrate that one of the purposes of an effective knowledge system should be to help us identify and learn from the outliers: those interventions where the results of our efforts were unusually successful (or unusually unsuccessful). Recognising these can help us to guide our practice and our investments across our organization towards strategies that deliver the greatest impact on poverty and injustice and away from those we have learned to be ineffective or harmful: in other words, shifting our quality curve.

Our current knowledge system is challenged in several ways in trying to fulfil this critical function, but most can be summarised in one statement: we find it very difficult to make judgements on quality. We face the ‘apples and oranges’ dilemma in much of our data, with a lack of uniformly defined indicators or metrics making comparison difficult. The problem of incentives rears its head again: it can be very rewarding to share unusually positive outcomes, but few incentives exist to share our failures, rich as they may be with learning⁴. Instead, our knowledge base risks becoming clogged with examples from the middle ground of our bell curve, where what happened was pretty much what we would have expected given our previous experience with that kind of input in that kind of context. Identifying in this

³ Six Sigma, developed by Motorola in 1986, is a disciplined, structured program aimed at delivering near perfect products and services by analyzing processes to discover where and how defects occur, measuring them and eliminating the problem areas. See e.g. www.bqf.org.uk/performance-improvement/about-lean-six-sigma

⁴ The frequent requests we make to field colleagues for ‘success stories’ reinforces this expectation.

haystack of success stories and case studies the needle of the truly surprising – data which, if validated and repeated under a range of conditions, could prompt us to change our whole approach to a particular development challenge – is the work of an expert. Or rather, of a network of experts, connected in a community, empowered to examine, test, compare and challenge, and forge consensus around what really works and drive uptake across the organization. It is critical that the knowledge system we develop as CARE supports this kind of ‘crowd-sourcing’ of expert judgements on quality, both in terms of the platform that makes it possible and the culture that empowers us to use it.

Where we came in

This brings us back to where we came in: the vital importance of creating incentives around knowledge and learning that support our ambitions as an organisation. The arguments presented here suggest four priorities for further investment in knowledge management at CARE, of which only the first is a regular topic in our organizational conversations on this theme:

1. Optimize our systems and platforms to minimize the costs of sharing knowledge
2. Increase the rewards, including building in systems for recognition and encouragement from the knowledge community.
3. Develop simple metrics and use them to establish expectations and accountabilities around knowledge that influence leadership behaviours across the organization.
4. Invest in an intelligent knowledge management function in each of our areas of programmatic focus – which can be supported by an excellent platform but ultimately depends on a connected community of experts with the breadth of view to link new ideas to existing knowledge structures, the bandwidth to follow-up, prompt and badger, and the courage to make judgements on quality.

The prize is significant. CARE’s 2020 Vision describes a world in which *‘CARE’s knowledge of successful approaches to addressing the underlying causes of poverty, in particular gender discrimination, is systematically captured and utilized through an interconnected information and knowledge management culture and system. This has become the key strategic resource to support CARE’s mission.’* Now there’s a powerful incentive.