Book Overview

How to Enable Innovation

Boru Douthwaite, *Enabling Innovation: A Practical Guide to Understanding and Fostering Technological Change*, Zed Books, London 2002
International Institute of Tropical Agriculture
Oyo Road, Ibadan, Nigeria
b.douthwaite@cgiar.org

Introduction

Making innovation happen is central to what many engineers do. However, when we finish our training most of us believe that it is our job to conceptualize designs, develop products and worry little about what happens after they have been introduced. Our courses are generally too practical to bother with theories about how innovation occurs, who it affects and how we might better manage the process. Diesel, inventor of the diesel engine, distinguished between two phases in technological progress: the conception and carrying out of the idea, which is a happy period of creative mental work in which technical challenges are overcome, and the introduction of the innovation, which is a "struggle against stupidity and envy, apathy and evil, secret opposition and open conflict of interests, a horrible period of struggle with man, a martyrdom even if success ensues" (as quoted by Mokyr, 1990, p. 155). Diesel is perhaps overstating the difficulties of managing innovation, but nevertheless as engineers we are still taught to prefer technical 'invention' and leave dealing with people and the 'innovation' side to others. However, I learnt from experience that we ignore the innovation process at our peril. The book, *Enabling Innovation: A Practical Guide to Understanding and Fostering Technological Change* is my attempt to explain why innovation approaches matter, and to develop an approach, based on a model, to managing innovation that builds on peoples' ingenuity and motivations, rather than one that fights against them.

Why innovation approaches matter

The book begins by describing a formative experience for me in Burma. In 1995 the military junta, the State Law and Order Restoration Council (SLORC) decided that, to boost production, the country's rice farmers should grow two crops of rice each year instead of one. There was a good reason why most Burmese rice farmers grew only one crop, however: growing two meant harvesting the second in the middle of the monsoon and, without very fast harvesting and drying, the grain would go moldy and spoil. The traditional single crop meant that the grain could be dried in the field after the rainy season and that there was far less rush. SLORC realized this, of course, and had asked the director of the Agricultural Mechanisation Department (AMD), part of the Ministry of Agriculture, just 6 months to come up with a rice harvester that could save the first crop by working in wet conditions.

B. Douthwaite. "How to Enable Innovation". Agricultural Engineering International: the CIGR Journal of Scientific Research and Development. Invited Overview Paper. Vol. IV. October, 2002.

By July 1995, when AMD's search had become frantic, the department was given the drawings of a rice harvester. These drawings were the fruit of five years of research and development I'd carried out with a team I'd led at the International Rice Research Institute (IRRI) in the Philippines, and with help from local manufacturers and the Philippine Rice Research Institute (PhilRice). The harvester my team had designed and built is known as a stripper-gatherer because, rather than cutting the rice so that it can be carried elsewhere for threshing to extract the grain, it moves through the field gathering the grain by stripping it from the standing stalks. Desperate for a solution, AMD set about building one immediately from IRRI's drawings. When it seemed to work they videotaped it in action and AMD's Director showed the footage to the Minister of Agriculture and then to the whole of SLORC. Four weeks after the drawings arrived, and without anyone using the machine more than twice, SLORC decided to build two thousand units, one thousand of which were to be ready within three months to be then be distributed to the country's tractor stations. IRRI did not find out about what was happening until production had already begun.

The details of what happened next are in book. In short, hardly any of the machines were ever used. Thankfully, only 1000 machines were eventually made, but all of these ended dumped in sheds or in the bush to rust away. In the rush to build the machines quickly, quality control had been scrapped and substandard materials had been used, making the machines inoperable without significant modification. Secondly, the few harvesters that were used were rejected by the farmers because the machines did not cut the straw but rather left it in the field making it unavailable for animal fodder and making subsequent land preparation much harder.

Why had this happened? When I asked the factory manager why there was no quality control he admitted that he knew there were problems with the machines but fixing them would mean he would not reach his quota. He was worried that any delays or negative reports from him would cost him his job, and was relying on the tractor station managers to keep quiet as well. When I visited a few tractor stations I quickly realized that this was the way things were done in Burma. I found that the stripper harvesters had been abandoned next to foot-operated rice mills, rice-hull stoves and other equipment that had been manufactured by AMD in previous years. Neither farmers nor the tractor stations had been asked if they wanted the equipment. It had just been assumed that the AMD engineers knew best and could develop what was needed with little consultation.

When I left Burma for the last time I learned that AMD was starting to build seven thousand mechanical rice reaper harvesters which were much more complicated that the stripper harvester, and so even less likely to work. Nothing had been learnt. I realized that the Burmese Ministry of Agriculture, AMD and the tractor stations were all locked into a top-down model of technology transfer that people said was working when it wasn't because they were too afraid of the consequences of feeding back stories of failure.

It would be easy to dismiss what happened in Burma as the inevitable outcome of having a military junta, running a centrally controlled government through fear. This, however, would be a mistake, because the only way this story differs from others I came across in the nine years I worked in Asia is that it is more extreme and its lessons are consequently clearer to see. The fact is that similar centrally-made decisions about what is 'good' for farmers have led to even greater wastage of resources in other countries.

My Burma experience, as well as the realization that it was not isolated, led me to two conclusions: firstly, the way people think about and plan for innovation is vitally important; and secondly, an adequate model of the innovation process, particularly for early innovation where products move from concept to initial manufacturing, did not exist. I discovered that most people thought little about how innovation would happen, and when they did, tended to assume a model that had worked well for distributing the high yielding plant varieties responsible for the Green Revolution. This is a top-down model, very much like that used by SLORC on the stripper-harvester, which sees formal Research and Development (R&D) laboratories as the source of an innovation which is then passed on to others to implement. The key stakeholders—the people who will reproduce and use the technology—are not seen as sources of innovations or ideas in their own right. And I also found out that a similar model is also mistakenly used in the developed world. Von Hippel (1988) in his influential book the *Sources of Innovation*, writes: “It has long been assumed that product innovations are typically developed by product manufacturers. Because this assumption deals with the basic matter of who the innovator is, it has inevitably had a major impact on innovation-related research, on firms’ management of research and development, and on government innovation policy. However, it now appears that this basic assumption is often wrong.” (von Hippel, 1988, p. 3).

These realizations motivated me to enroll for a Ph.D. to look at case studies of successful and unsuccessful postharvest equipment in the Philippines and Vietnam. I subsequently developed a model of the early innovation process, called the learning selection model. The model, and its application to technologies other that agricultural equipment—wind turbines, computer software, local currencies and seed—is the basis of the book *Enabling Innovation*.

Developing the Learning Selection Model

In my Ph.D. I constructed case studies of the early adoption of 13 attempts to introduce postharvest equipment (Table 1 and Table 2) from both the public and private sectors.

Table 1: The case study technologies

Technology	Description	Adoption status	Cost $
Stripper-gatherer (SG) harvester	Walk-behind harvester	140 units sold in 5 years (Philippines)	2000
Mechanical reaper	Walk-behind harvester	1071 units sold in 8 years (Philippines)	3000
SRR dryer	Low temperature dryer	700 units sold in 3 years (Vietnam)	100
Flatbed dryer	Heated air dryer with manual mixing	1000 units sold in 17 years (Vietnam)	2000
Flash dryer	High temperature dryer	2000 units donated in 4 years (Philippines)	3500
Recirculating dryer	Heated air dryer with mechanical mixing	1500 units sold in 6 years (Philippines)	15,000

Table 2: Case study technologies

Technology	Source of innovation	Introduced? Philippines	Vietnam
Harvesting			
SG	Public	✔(x2)	✔
Reaper	Public	✔	✔
	Private	✔	
Drying			
SRR	Public		✔
Flatbed	Public	✔(x2)	✔
Flash	Public	✔(x2)	
Recirculating	Private	✔	

The case studies I developed described the innovation histories of the technology, following a normative view of the innovation process shown in Figure 1. This view incorporates Roger's (1994) classification of adopters into five categories ranging from innovators, the first group to adopt a new technology, through early adopters, early majority and late majority to laggards, the last. He described innovators as venturesome, enjoying the technical challenges posed by new technologies and actively seeking them out. Laggards, by contrast, are the last people to adopt because they do not like taking risks and are conservative in their outlook.

Figure 1: Phases in the innovation process

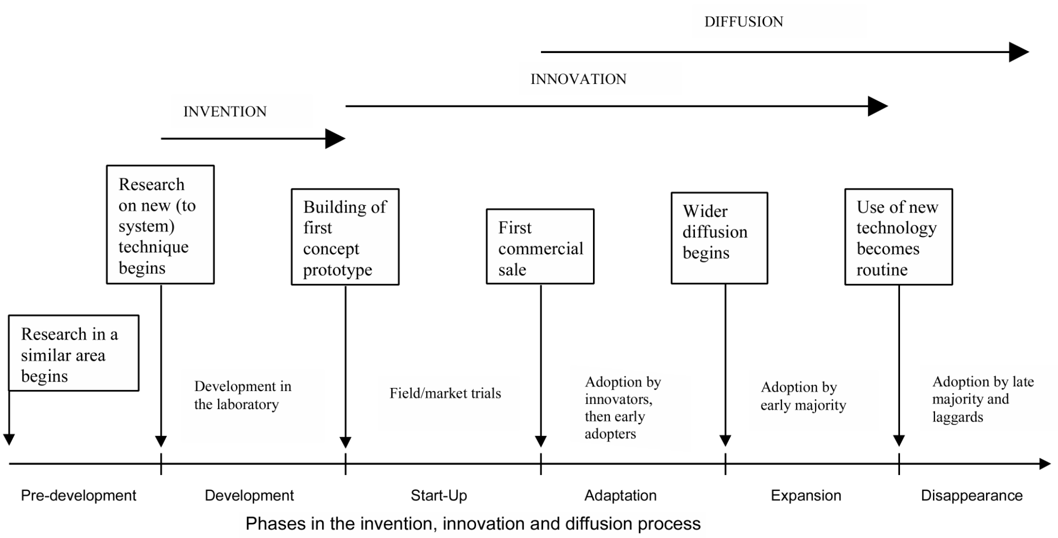

The main finding from the work, and the most striking, was that the successful technologies were the ones in which manufactures and users had modified the most. This was in complete contradiction to my engineering training which had lead me to believe that machinery was developed by engineers, not by end users, and good designs would need few, if any, subsequent modifications. What had become clear from my results was that engineers and designers were singularly unable to develop machine designs that people adopted, without a great deal of further co-development with the manufactures who would build the machine and the people who would use it. This co-development occurred when manufacturers and users believed that the first commercial prototype made a 'plausible promise' of being of benefit to them, thus motivating them to become co-developers. In the co-development process the key stakeholders learnt about the equipment and developed their own procedures and protocols that often increased the performance of the equipment in ways that the engineers had not envisaged. In short, the successful equipment *evolved* after launch through adaptations made by the key stakeholders, increased in fitness as a result, while unsuccessful equipment did not evolve.

I developed the learning selection model, shown in Figure 2, to describe the early evolution of postharvest equipment that I had observed. As the name suggests, the learning selection model is based on an analogy with natural selection, which is the algorithm that drives biological evolution. Natural selection consists of three mechanisms. These are:

- *Novelty generation:* As a result of random genetic mutations and sexual recombination of differing genetic material, differences between individual members of a species crop up from time to time.
- *Selection:* This is the mechanism which retains random changes that turn out to be beneficial to the species because they enable those possessing the trait to achieve better survival and breeding rates. It also rejects harmful changes.
- *Diffusion and promulgation*: These are the mechanisms by which the beneficial differences are spread to other areas.

The learning selection model is shown in Figure 2. It shows a technology, depicted as a cogwheel, beginning as a 'plausible promise' that motivates the key stakeholders to co-develop it. The technology then increasing in fitness by gaining knowledge and becoming 'meshed in' to existing systems through the adaptation and learning that takes place. Here, fitness is taken in the biological sense to mean improvements in the likelihood that the technology will be adopted and promulgated. The 'meshing in' of the technology, or its 'social construction' as it is also called, is represented by the move from a single cogwheel to three inter-locked ones. The increase in knowledge is represented by the increase in size of the cogwheel(s).

Figure 2: The Learning Selection Model

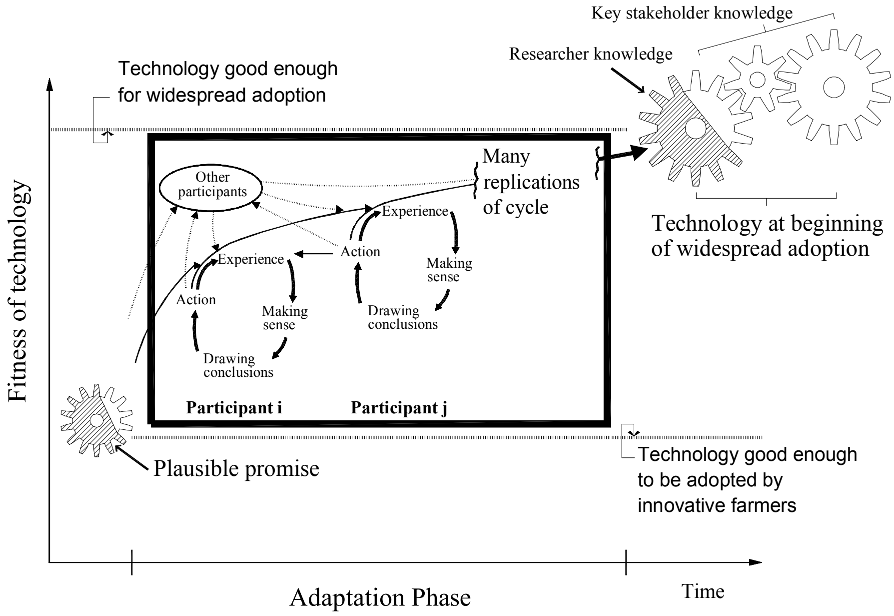

Learning selection is shown inside the black box in Figure 2 and is responsible for the evolution. Learning selection is a process built on the 4-stage experiential learning cycle, and is perhaps best explained using an example.

Experience—Suppose a farmer finds that the rice miller pays her a low price for the grain dried in her dryer because some of it is not properly dried.
Making Sense—She reflects and makes sense of the experience. She realizes that uneven drying is loosing her money and that it might be sensible to try and improve the dryer's performance.
Drawing Conclusions—She then develops personal explanations of what happened from her own or others previous experience or theories. She hypotheses that if she reduces the amount of paddy she loads into the dryer then drying will be more uniform.
Action—She then decides to test her hypothesis, and in so doing *generates a novelty*.

Testing the novelty begins another learning cycle. Her *selection* decision to adopt or reject the novelty will depend on whether the rice miller pays her more for her product. The miller will make this price decision after going through his own learning cycle when he tests a sample of her rice for milling quality. If the farmer is participant *i* in Figure 2 then the miller represents participant *j*.

So far the third component of evolutionary system—the *promulgation and diffusion* mechanism—is missing. In the example promulgation of the novelty occurs when the farmer tells people in her social network, represented in Figure 2 by the 'other participants' box, about the benefits of her novelty and they select to adopt it. Moreover, many of these people may be going through their own learning cycles creating the conditions for the *recombination* of differing observations and experiences that can lead to the generation of novelties that have 'hybrid vigor.' In the process the technology evolves and with it the participants' opinions and knowledge of it and the way they organize themselves to use and promote the technology. These processes are all involved in learning selection.

Testing the Learning Selection Model on other types of technology

In the book I go beyond my Ph.D. research to see the extent to which the learning selection model applies to other technologies. The book shows that the model is applicable over a range of technologies from wind turbines, to computer software to local currencies and seed, where a high degree of uncertainty exists over the outcome. In other words, the LS model is useful when 'learning by using' and 'learning by doing' predominate over 'learning by modeling' in the early adoption phase. The LS model is also applicable if users are able to modify the technology, and if there are ways of evaluating changes.

Wind turbines

The wind turbine industry, described in Chapter 4, particularly shows the applicability of the model. Excitingly, it also shows that the democratic user-led type of innovation that it describes is able to harness the innovative potential of the people who are directly affected by the technology. A grassroots development process in Denmark was able to produce a wind turbine industry with a 55% share of a billion dollar a year world market, beating the US who spent over 300 million dollars funding a top-down development program led by the National Aeronautics and Space Administration (NASA). The origins of the Danish industry were a few agricultural machinery manufacturers and ideologically motivated 'hobbyists' who began building, owning, and tinkering with wind turbines (generating novelty). There were many early teething problems but the owners organized themselves into a group who lobbied successfully for design improvements (selection), working closely with manufacturers to solve the problems. The owners' group developed a co-operative ownership model and pressured politicians to support the sale of their electricity to the national grid at a fair price (promulgation and diffusion). In contrast, the NASA led a top-down science development approach that implicitly assumed that scientists could develop the 'perfect' wind turbine with little input from the owners and users. NASA's approach failed.

Computer Software – Linux and Windows

Another example of the power that a grassroots innovation model can harness is the development of the computer operating system Linux, which is a "a world-class operating system" that has coalesced "as if by magic out of part-time hacking by several thousand developers all over the planet connected only by the tenuous strands of the Internet" (Raymond, 1997). Linux started life when a Finnish computer science student started to write a Unix-like operating system that he could run on his PC; he had become tired of having to queue for hours to gain access to Unix on the University's main frame. When he finally got the core of an operating system working he posted it on the Internet so that others could try it out. Importantly he gave the source-code so other people could understand the program and modify it if they wanted. Just like the first Danish wind turbines, early versions of Linux were not technically sophisticated or elegant, but they were simple, understandable, and touched a chord with 'hackers'—people like Torvalds himself who got a kick out of generating novelty for the sake of being creative, not for money.

Torvalds' main role in the development of Linux after the first release was not to write code for features people wanted but to select and propagate improvements to the system from the ideas that streamed in. Ten people downloaded version 0.02 and five of these sent him bug fixes, code improvements and new features. Torvalds added the best of these to the existing program along with others he had written himself and released the composite as version 0.12. The rate of learning selection accelerated as the number of Linux users increased and, to cope with the volume of hacks (novelties) coming in, Torvalds began choosing and relying on a type of peer review. Rather than evaluate every modification himself he based his decisions on the recommendation of people he trusted and on whether people were already using the patch (modification) successfully. He in fact played a similar role to that of an editor of an academic journal who makes sure submitted articles are reviewed but retains final control over what is published and what is not. This approach has allowed Torvalds to keep the program on track as it has grown, as Figure 3 shows, from 10,000 lines of code to 1.5 million, all written by volunteers.

Figure 3: The phenomenal growth of Linux shown in terms of users and lines of code

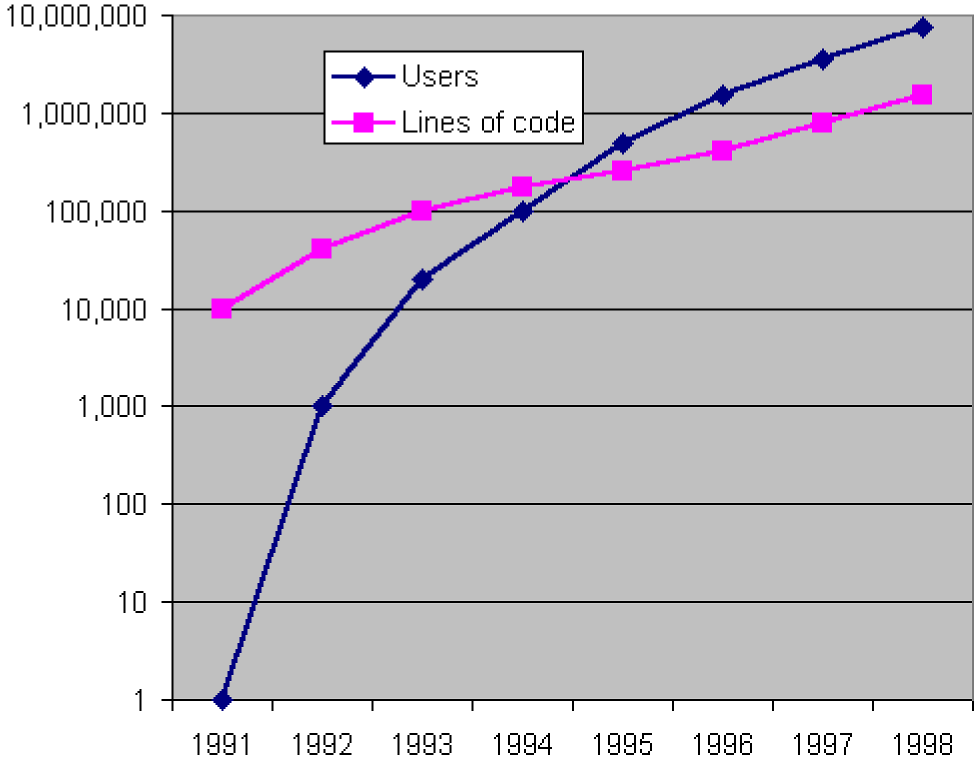

Such has been the success of Linux that Microsoft, which until recently was the richest company in the world based on market capitalization, is privately worried. Vinod Valloppillil, a Microsoft engineer, analyzed the open source software movement in a confidential memorandum that was leaked and posted on the web. Valloppillil (1998) wrote: "Linux could win … The ability of the open source software process to collect and harness the collective IQ of thousands of individuals across the Internet is simply amazing". Microsoft jealously guards its own source code to make sure they remain closed and users cannot modify it. While Linux is not yet seriously threatening Microsoft's 90% domination of the PC market by the end of 1998 Linux was installed on 17% of servers—the computers that run networks including the Internet—up from 7% the previous year. Windows NT, the market leader, was fairly static at 36% (Gomes 1999).

The fact that a grassroots, communitarian development model can lever more creative talent than one of the richest companies in the world has, I feel, an exciting resonance particularly for engineers working in the rural development field.

The learning selection approach to co-developing innovations with users

The learning selection model, and the wind turbine and Linux examples in particular, show that the learning selection model can provide both a better way of understanding the research, development and early adoption process and of managing it. Hence, in the last chapter I develop the learning selection approach to understanding and fostering innovation, which is based on the map of the innovation process shown in Figure 4.

Figure 4: The learning selection view of the innovation process

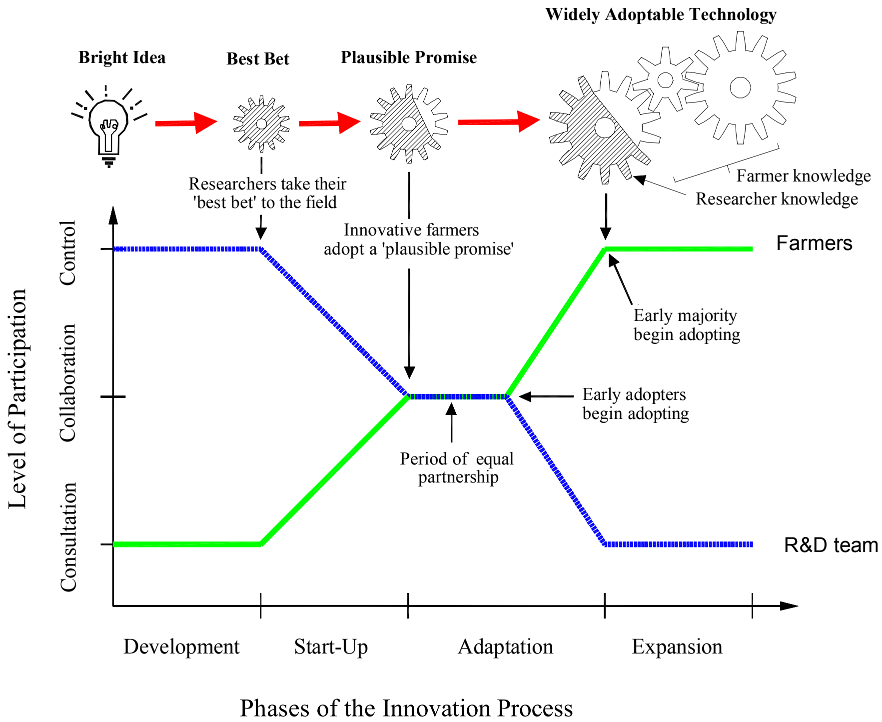

The figure shows an innovation process beginning with a bright idea that individuals or small teams of researchers then develop in relative isolation. While the R&D team may ask the key stakeholders—the people who will ultimately take ownership of their idea, replicate it and make it work—for some advice, they are driving the process. Mokyr (1990, p. 9) believes it has to be this way because the process of inventing 'plausible promises' is by its nature something that 'occurs at the level of the individual'. He says creating a plausible promise is 'an attack by an individual on a constraint that everyone else has taken for granted'. It is not something that lends itself to a broad consensus approach.

At some point the R&D team crystallizes the knowledge they have generated into a prototype: their 'best-bet' of what the key stakeholders want. Then, in what marks the beginning of the start-up phase, they begin to demonstrate their best-bet to the key stakeholders. It may take several prototype iterations before the R&D team has received and incorporated sufficient feedback for at least a few innovators to adopt it. It is this adoption, based on the belief that the new technology makes a 'plausible promise' of bring benefit, which marks the beginning of the adaptation phase. It also marks the beginning of a period of co-development and learning selection in which the technology evolves and its fitness improves, as shown in Figure 2.

Learning selection is analogous to natural selection in Darwin evolution. The process works when people make changes to a technology and then select and promulgate the ones that they find beneficial. This improves the adoptability of the technology—its suitability to the environment in which it is used—and hence its market appeal. At a certain point the attributes of the technology are good enough for the second category of adopters, Rogers' early adopters, to start to show an interest. This marks the point at which the key stakeholders begin to take over ownership of the technology.

However, the analogy between natural selection and learning selection is not perfect. One important difference is that natural selection is blind and learning selection is not--genetic mutations occur at random but technology and system change can be directed. Hence, learning selection does not necessarily happen. It only comes about if the key stakeholders are sufficiently motivated to modify it and carry out sensible learning selection on it. They must also understand the technology well enough to do so themselves. Consequently, at least one stakeholder who understands the technology is essential as he or she must champion it and fill knowledge gaps until the other stakeholders have learned enough to take over. This take-over marks the end of the early adoption process and is the point at which market selection begins to work.

The take-over also marks the beginning of the expansion phase when the technology becomes mainstream. As this happens, the people adopting the technology change from hackers (innovators) and early adopters to people who want the technology to work reliably and profitably. Increasingly in this phase, manufacturers and researchers are able to gather and codify more and more information that can be used to build predictive models. This allows them to move from 'learning by using' which requires adopters to be co-developers, to 'learning by modeling', where learning comes from virtual tests carried out on computer rather than field experience. In so doing, our learning selection model of the innovation process becomes less relevant and the conventional assumption that manufacturers or R&D departments can and do develop finished technology begins to fit better.

A practical 9-point guide to enabling innovation

I also develop a practical nine-point guide to catalyzing an innovation process, written for R&D managers working in the public or private sector.

Start with a plausible promise

The first step to induce change through learning selection is to produce a 'plausible promise'; something that convinces potential stakeholders that it can evolve into something that they really want. Experience shows that it is difficult to enlist co-developers if the whole project is abstract and up in the air.
The plausible promise does not need to be refined or polished: it can be imperfect and incomplete. In fact the less finished it is, the more scope there is for the stakeholders to innovate and thus gain ownership of the technology. On the other hand the more problems there are then the greater the chances that the key stakeholders will give up in frustration. A delicate balance must be found.

Find a product champion

The next step is to identify the innovation or product champion. He or she needs to be highly motivated and have the knowledge and resources to sort problems out. Someone from the R&D team is likely to be suitable because he or she will probably have both the necessary technical knowledge and the motivation as they already have a stake in the technology. He or she must also have good people and communication skills as, in order to build a development community, they will need to attract people, interest them in what they are doing, and keep them happy working for the common cause. The product champion's personality is therefore crucial.

Keep it simple

Don't attempt to dazzle people with the cleverness and ingenuity of the prototype's design. A plausible promise should be simple, flexible enough to allow revision, and robust enough to work well even when not perfectly optimized. The critical comments of your colleagues don't matter. Your potential co-developers' needs and knowledge levels do. For example, if you are designing a combine harvester and you know the manufacturers and farmers you'll be working with are familiar with a certain type of thresher, then use that in your design, even if it is technically not the most elegant solution. As John Gall said: "A complex system that works is invariably found to have evolved from a simple system that worked." (As quoted by Quoteland, http://192.41.61.35/quotes/author/182.html)

Work with innovative and motivated partners

Allow the participants in your learning selection process to select themselves through the amount of resources they are prepared to commit. Advertise or write about your plausible promise in the media, by doing field demonstrations, or on the Internet and then wait for people to make the effort to contact you. Don't give inquirers anything with a resale value for free. For example, if your prototype has an engine, then charge the market value for it. Otherwise people may be motivated to adopt in order to get something for nothing. In addition, people generally value something more highly if they have paid for it and they will be more committed to sort out the problems that emerge.

On the other hand you must make it clear to the first adopters that they are adopting an unperfected product and that they are working with you as co-developers. You need to reassure them that you will be contributing your own resources to the project and will not abandon them with a lemon. You should be prepared to offset some, but not all, of the risk they are taking in working with you. Getting the balance right is very important here too.

Work in a pilot site or sites where the need for the innovation is great

Your co-developers will be influenced by their environment. Their motivation levels will be sustained for longer if they live or operate in an environment where your innovation promises to provide great benefits. In addition, they are more likely to receive encouraging feedback from members of their community.

Set up open and unbiased selection mechanisms

<u>(i) The product champion/selector</u>

Once you have the key stakeholders working with you and generating novelties, you need ways of selecting and promulgating the beneficial changes. Initially the product champion usually plays this role. An effective selector must be able and prepared to recognize good design ideas from others. This means that, if he or she is also the inventor, they must be suitably receptive and thus able to accept that others might have better ideas.

Very few people are capable of being effective at both championing their product and selecting novelties simultaneously. This is because to be good at the former they need to believe deeply in the product's benefits and be able to defend it against criticism. To be effective selectors, on the other hand, they need to keep an open mind and be able to work with others to question fundamental design decisions.

If a product champion defends the technology too strongly, or shows bias, then ‘forking’ occurs and the disaffected person or group branches off on its own to do what they felt prevented from doing by the selector. It is good to have people test alternative design paths but if it is done in frustration or spite then cliques form, making any comparison and subsequent selection between rival branches difficult. Creative talent is split and energies can be dissipated in turf wars.

(ii) Alternative selection mechanisms

Even if the product champion can be open-minded and unbiased he or she may have problems convincing others. One option is to set up a review mechanism that is well respected by your key stakeholder community. There are a number of ways of doing this. Three that work are: (i) review by an independent organization; (ii) peer review; and (iii) providing potential adopters with enough information to make informed selection decisions themselves.

Don’t release the innovation too widely too soon

For the innovation to evolve satisfactorily, the changes the stakeholders make to it need to be beneficial and, as those generating the novelties will have gaps in their knowledge, product champions should restrict the number of co-developers so that they can work with them effectively. When people show enthusiasm for a prototype it is very tempting to release it as widely as possible but this should be resisted. The technology will always be less perfect than one initially thinks.

However promising the technology might appear, there are many things that can and will go wrong. First adopters need to be aware of this and have ready access to the product champion. Otherwise, their enthusiasm will quickly turn to frustration and the product champion will end up defending the technology against their criticisms when the problems appear. Once the product champion becomes defensive, he or she will be far less useful at sorting out problems.

Don’t patent anything unless it is to stop someone else trying to privatize the technology

In learning selection, people co-operate with each other because they believe that all will gain if they do. The process is, therefore, seriously damaged if one person or group tries to gain intellectual property rights over what is emerging. Firstly, the communitarian spirit is damaged. Secondly, patents are monopolies that immediately reduce the novelty generation rate and thus slow down future development and the flow of ideas.

Realize that culture makes a difference

Culture can influence the degree to which knowledge is guarded within a particular group, or spread around. Learning selection is going to be greatly impeded in cultures where new knowledge is carefully guarded, either through secrecy or the taking out and enforcement of intellectual property rights.

Know when to let go

Product champions need to become personally involved and emotionally attached to their projects to do their jobs properly. This makes it easy for them to go on flogging dead horses long after it has become clear to everyone else that the technology is not going to succeed. Equally, project champions can continue trying to nurture their babies long after they have grown up and market selection has begun. It is, therefore, a good idea to put a time limit on the product champion's activities.

Using the learning selection model for monitoring and evaluation

Since I wrote the book I have been working at the International Institute of Tropical Agriculture (IITA) and applying the learning selection model to monitoring and evaluation (M&E) of on-farm research aimed at fostering innovation processes. M&E is essential to help the management of these projects adapt in time to unexpected events and unintended consequences. M&E is also needed to help the key stakeholders learn about the project outputs and thus motivate them to continue with the learning selection process. The M&E approach is simply to follow the technology introduced by a project and then asking the journalistic questions of what? why? who? when? where? and how? about the elements of an evolutionary process, that is, the novelties generated; selection decisions made; and the promulgation mechanisms used.

In this work I have found that identifying the changes and modifications that farmers are making to a technology or practice gives very good insights into their perceptions of the technology, the opportunities they see for it, and the constraints they face (Douthwaite et al., 2001). Some of the modifications made have improved the fitness of the technology and this has provided invaluable insights into the research process. The value of using the learning selection model in M&E is that it gets away from 'cookbook' approaches that come with pre-conceptions of how the project will unfold, usually hardwired into the indicators chosen. Focussing on what people are doing, or not doing to a technology, avoids measuring only what is expected.

Conclusions

The value of a book such as Enabling Innovation is that people find the ideas in it useful. The feedback I have received so far has been very encouraging, and I have listed some.

> **'A breath of fresh air. Here is an engineer looking critically and creatively at technological change and raising both practical and philosophical questions about the nature of innovation. Hopefully, engineering departments will include it in their courses. And for those already engaged in setting up companies and introducing new products there are useful, practical guides such as the section 'How to launch a learning selection innovation process''** - Mike Cooley, International authority on human-centered systems design

'A brilliant book… Douthwaite makes a very cogent case for using neo-Darwinian thinking for what he calls 'Learning Selection' as a preferred mode of technology invention' - Richard Jefferson, Center for Applied Molecular Biology to International Agriculture (CAMBIA), Canberra, Australia

'A refreshing approach to innovation as a complex adaptive multi-agent system. Innovation emerges as different agents learn and select improvements. Hence it is not the experts that generate knowledge and technology for us. We do it best ourselves in self-organizing networks of interaction. Based on a compelling use of examples from agriculture, industry, economy and IT, the book is relevant for a wide audience of people who look for ideas on which to base the management of innovation' - Niels Röling, Professor of Communication and Innovation Studies, Wageningen University, The Netherlands

'This is an original and important book. It challenges the conventional wisdom of many development assistance programs. Douthwaite leads us to a better understanding of the factors that determine people's willingness and ability to adopt new technologies' - Professor Jeff Sayer, World Wildlife Fund

'For me the book confirms once again that it's not enough for us to have bright ideas about what needs to be done. We need to build a theory about a process for achieving successful change. That theory is, I believe, going to have many of the components of 'learning selection'. The book contains many insights that I think can help with this wider agenda. "Enabling Innovation" has much to contribute to enabling change. It will, I predict, be much cited over the coming years. It's certainly a book I know I will keep going back to.' - John Jopling, Founder of Foundation for International Environmental Law (FIELD), University of London

References

Douthwaite, B. N. de Haan, V. M. Manyong and D. Keatinge. 2001. Blending 'Hard' and 'Soft' Science: The 'Follow the Technology' Approach to Catalyzing and Evaluating Technology Change. *Conservation Ecology* 5(2). Available from: http://www.consecol.org/Journal/vol5/iss2/art13/index.html

Mokyr J. 1990. The Lever of Riches: Technological Creativity and Economic Progress. Oxford University Press, Oxford.

Raymond, E.S. 1997. The cathedral and the bazaar. Downloaded from www.openresouces.com/documents/cathedral-bazaar/index.htm on 15th August 1999.

Rogers E.M. 1995. *Diffusion of Innovations*. 3rd edition. Rev. ed. of: Communication of innovations. The Free Press, New York, USA

Valloppillil, V. 1998. *'Halloween Memorandum' 11 August*, downloaded from http://www.opensource.org/halloween/halloween1.html

von Hippel, Eric, 1988. *The Sources of Innovation*. Oxford University Press, New York and Oxford.