innovations
TECHNOLOGY | GOVERNANCE | GLOBALIZATION
mitpress.mit.edu/innovations

Boru Douthwaite

Enabling Innovation

Technology- and System-Level Approaches that Capitalize on Complexity

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."[1] 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, engineers ignore the innovation process at their peril. Enabling innovation means building on peoples' ingenuity and motivations, rather than working against them.

In this paper I describe the learning selection approach to enabling innovation that capitalizes on the complexity of social systems at different scales of analysis. In the first part of the paper I describe the approach and how it can be used to guide the early stages of setting up a "grassroots" innovation process. In the second part of the paper I look at how the learn selection model can be used "top-down" to guide research investments to trigger large-scale systemic change.

Boru Douthwaite is a technology policy analyst working for the International Center for Tropical Agriculture (CIAT) in Cali, Colombia. He previously worked as an agricultural engineer for eight years in the Philippines at the International Rice Research Institute.

This paper draws heavily on his book Enabling Innovation: A Practical Guide to Understanding and Fostering Technological Change *(London: Zed Books) published in 2002.*

© 2006 by Boru Douthwaite

Boru Douthwaite

WHY INNOVATION APPROACHES MATTER

In 1995 the Burmese 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, in just 6 months to come up with a rice harvester that could save the first crop by working in wet conditions.

By July 1995, when AMD's search had become frantic, somebody, and I still don't know who, gave the department 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. 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 the 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 then be distributed to the country's tractor stations. I did not find out about what was happening until production had already begun, and traveled to Burma soon afterwards.

Hardly any of the machines were ever used. Thankfully, only the first 1000 machines were 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.

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 that 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 threshers, 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 learned. 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 eight 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 early innovation process, 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. As Von Hippel comments: "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."[2]

These realizations motivated me to learn from other successful and unsuccessful attempts to introduce harvesting and drying equipment into Asia. I researched 13 cases in total and as a result developed a model of the early innovation process, called the learning selection model.

LEARNING SELECTION

The main finding from the research, and the most striking, was that the successful technologies were the ones which manufactures and users had modified the most. The research showed that engineers and designers were often 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 learned 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 to describe this process.[3] 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 detrimental changes.
- *Diffusion and promulgation.* These are the mechanisms by which the beneficial differences are spread to other areas.

The learning selection model is depicted graphically in Figure 1. It shows a technology, shown as a cogwheel, beginning as a "plausible promise" that motivates the key stakeholders to co-develop it. The technology then increases 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 might also be termed, 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 cogwheels.

Learning selection is shown inside the black box in Figure 1 and is responsible for the evolution. Learning selection is a process built on Kolb's 4-stage experiential learning cycle,[4] 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.

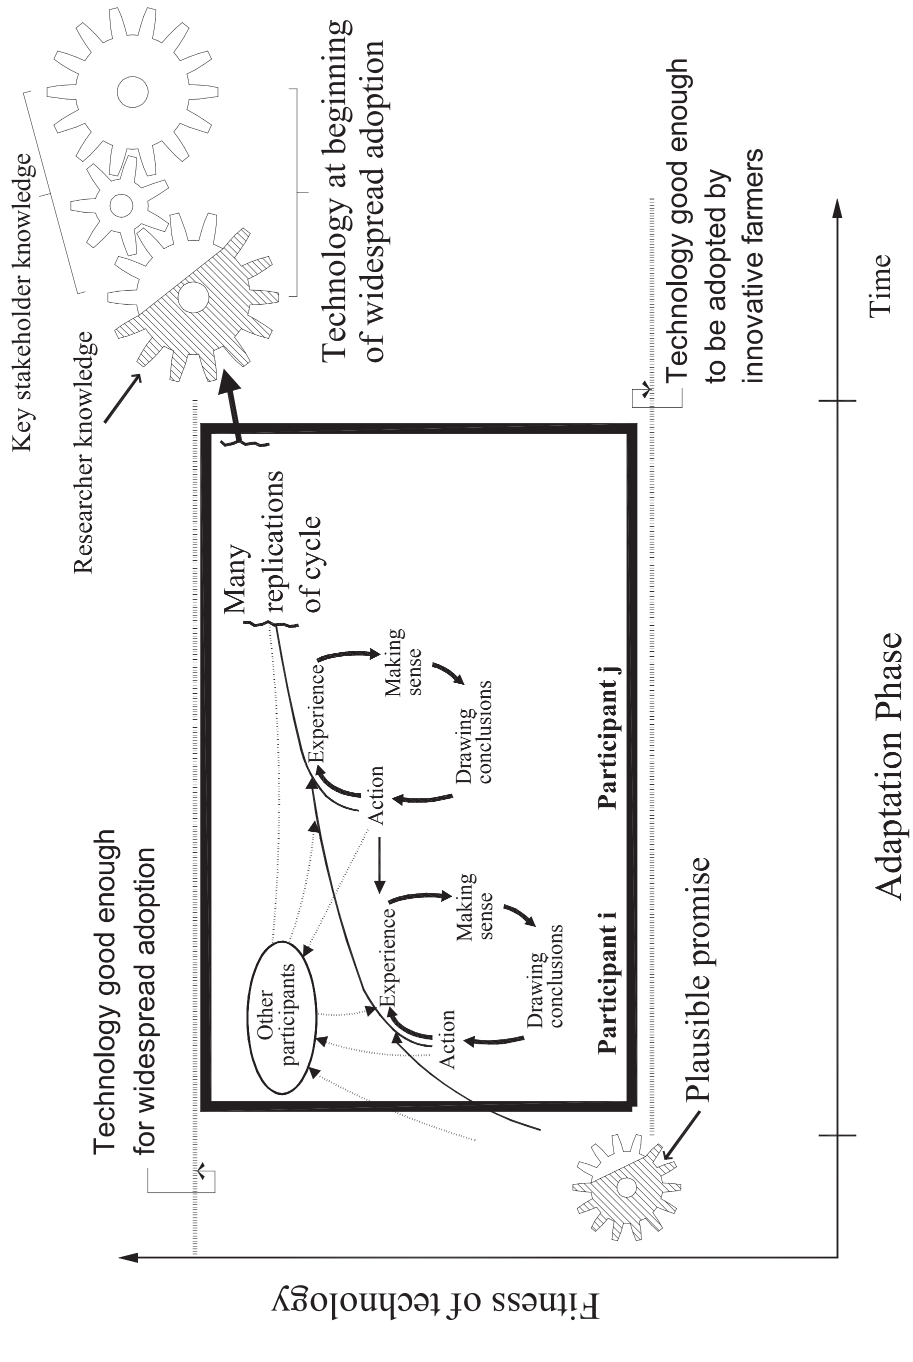

Figure 1. The Learning Selection Model

- *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 rice 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 1 then the miller represents participant *j.*

So far the third component of the 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 1 by the "other participants" box, about the benefits of her novelty and they choose to experiment with it themselves.

The farmer, the miller and the people they are connected to them through their social networks will be going through their own learning cycles creating the conditions for the recombination of differing observations and experiences that can lead to 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 it. These processes are all involved in learning selection.

The learning selection model is most useful when key stakeholder "learning by using" and "learning by doing" are important in the early adoption phase, which is the case for technologies that open up new markets. The learning selection process works best when users are able to modify the technology, and if there are ways of evaluating changes.

Wind turbines

The wind turbine industry is a good one for describing the applicability of the learning selection model. Excitingly, it shows that learning-selection-type innovation processes are able to harness the innovative potential of the people who are directly affected by technology. A grassroots development process in Denmark was able to produce a wind turbine industry that had a 55% share of a billion dollar a year world market in 2000, beating the U.S. 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

Another example, *now very well-known*, of the power that a grassroots innovation model can harness is the computer operating system Linux. Linux 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."[5]

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 allowed Torvalds to keep the program on track as it grew.[6]

The learning selection approach to co-developing innovations with users

The wind turbine and Linux examples show that the learning selection model can provide a powerful way of understanding the research, development and early adoption process and of managing it. Figure 2 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 argues 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.[7]

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 bringing 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, through the process shown in Figure 1.

Learning selection works when people make changes to a technology and then select and promulgate the ones that they find beneficial. This improves the fitness 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,[8] 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 adopt and modify a technology 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, often from the R&D team, 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 the market becomes the main selection mechanism. 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, the 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 TEN-STEP APPROACH TO ENABLING INNOVATION

It is possible to derive a ten-step approach to enabling grass-roots innovation based on the learning selection model.

1. 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.

2. 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.

3. 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. To quote John Gall, "A complex system that works is invariably found to have evolved from a simple system that worked."[9]

4. 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.

5. 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 benefit. In addition, they are more likely to receive encouraging feedback from members of their community.

6. Set up open and unbiased selection mechanisms

(i) The product champion

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 receptive and 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

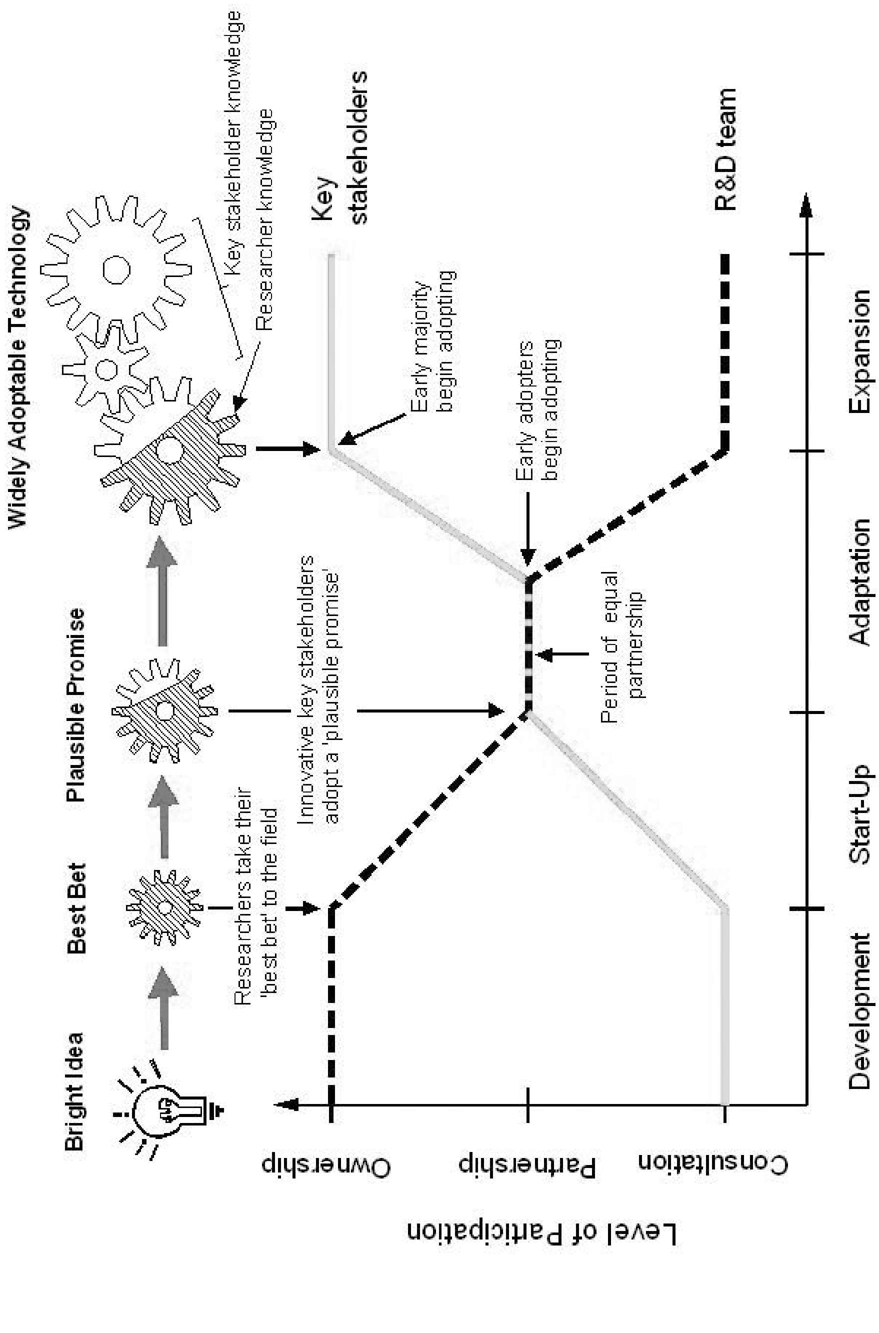

Figure 2. Stages and Participation in a Learning Selection Innovation Process.

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.

7. 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 largely 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. Once the product champion becomes defensive, he or she will be far less useful at sorting out problems.

8. 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.

9. 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.

10. 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.

In the following section I describe how ideas from learning selection can help trigger systemic change.

BEYOND LEARNING SELECTION: RESEARCH TO TRIGGER A "BLUE REVOLUTION" IN AGRICULTURE

I work for one of the 15 international agricultural research centers based in Africa, Asia, Latin America, Europe and the United States of America that constitute the Consultative Group on International Agricultural Research (CGIAR). In the 1960 and 1970s research in CGIAR Centers helped spark the Green Revolution by breeding improved higher yielding crop varieties that were then disseminated to farmers by national agricultural research and extension systems. This "pipeline" approach is effective for developing certain types of technology, like seed and vaccines—technologies that have some characteristics of magic bullets. In the case of seed and the Green Revolution, farmers planted improved seed, harvested more grain, sold it for more, decided to plant it again and gave some seed to their neighbors. Adoption of new varieties and improved yields spread like a virus. The role of research was clear—keep breeding improved varieties to replace those in the field when inherent resistance to pests and diseases breaks down. But the pipeline approach does not work for more complex technologies than seed, as the Burma story at the beginning of the article showed.

Today the CGIAR is confronting a new challenge: catalyzing a "blue revolution" to content with the global challenge of water scarcity. Unless water use patterns change substantially, within 25 years agriculture can be expected to be using an additional 500 km^3 of water if the world is to feed the additional billions who will live on the planet by then. This is more water than flows down the Mekong River in one year. This water would have to come from the world's major rivers, aquifers, wetlands and lakes that are already under pressure. Already many large rivers now run dry or clog up before they reach the sea, and an estimated half of the world's five million lakes are endangered. Unless there is "more crop per drop," many aquatic ecosystems will collapse and conflicts over water will increase. Climate change only makes the challenge greater.

In contrast to the Green Revolution, there is no magic bullet for a Blue Revolution. As a consequence, a pipeline approach to research will not be effective. What is more, the feedback between innovation and a successful outcome is far less direct in the case of water conservation than it is in the case of increased agricultural yields. Water flows between groups of people. Most of the time, when farmers save water it is people downstream who benefit, often people they don't know. Improving how water is managed often requires technical innovation, but is at least as much about linking people together, improving social processes of negotiation and changing norms.

Realizing this, the CGIAR researchers focused on the task of improving global water management have sought not simply to develop new technologies to push into target regions, but rather to build the capacity of networks of water and agricultural scientists to develop local solutions with the people who would use them. The result is the Challenge Program on Water and Food (CPWF), an international, multi-institutional research initiative with a strong emphasis on "north-south" and "south-south" partnerships. The initiative brings together research scientists, development specialists, and river basin communities in Africa, Asia and Latin America.

With a budget of $15 million/year, the CPWF is extended over nine of the world's largest and most important river basins including the Nile, Mekong, Limpopo, Volta and Yellow River. One way these scarce resources can effect change at the needed scale is if they motivate large-scale learning-selection-type innovation processes. This is possible as the case study in Appendix A shows—a story of adoption of a no-till technology saved 1.16 km^3 of irrigation water in India and Pakistan.

The learning selection model can also help guide top-down system-level interventions. The model and complexity theory[10] suggests that rate of innovation in a system can be changed by three sets of interventions: 1) changes in the variation or novelty of system components (i.e., types of stakeholders involved and the strategies and technologies they use); 2) changes in interaction patterns between stakeholders, in particular through changes to social networks; and, 3) changes in the way selection decisions are made.The research that helped trigger the Green Revolution focused mainly on developing and introducing novelty. Research for a Blue Revolution needs to be more balanced, as we now explore.

Many of the solutions to water management problems already exist. Consequently, part of the CPWF research should be to identify promising local solutions. Once identified these "best bets" have the potential to spark similar improvements in similar systems.

Although many solutions may exist, too much novelty and variation creates uncertainty in people and can prevent them from adopting and innovating. Research is needed to reduce this uncertainty by describing the variation that already exists. People have been managing water for millennia and do not need more variation so much as better understanding of what already exists.

Much more research is also needed on understanding how changing interaction patterns and selection processes can help people manage their water properly. As we've already discussed, in some ways the Green Revolution was easy. Farmers clearly saw the benefit in growing improved seed varieties in terms of higher yields and incomes. Those making the changes directly benefited. Feedback was strong and clear and adoption increased quickly. Politicians too understood and adoption was supported by massive state endorsement, subsidies on fertilizers, the building of irrigation schemes and numerous policy changes.

Improving how water is managed is a different story. When farmers improve the way they manage water the benefits may not accrue to them but to other water users downstream. Conversely, if farmers use excessive amounts of water the effects are not felt by them but by people downstream for whom water becomes unavailable. Feedback is weak or non existent so there is almost no incentive to select and promulgate better water management practices. Learning selection cannot work. Hence it is research on changing interaction patterns, promulgation pathways and selection mechanisms to improve information flow, feedback, negotiation and decision making which probably offers the greatest potential for triggering a Blue Revolution.

Research to trigger a Blue Revolution should further develop a twin strategy of fostering local "pilot site" changes while looking for opportunities to catalyze much larger scale changes. This requires mechanisms, such as innovation funds, that can support potential winners. It also means using research to improve systemic understanding to become better at spotting early winners, and knowing in which systems they are likely to first emerge.

At the level of project management, a practical approach to making the most of complexity is to facilitate a collaborative process in which project staff and stakeholders come to a common understanding of how they see a project achieving outcomes and impacts. Doing so can help the project achieve impact by mapping out promising "impact pathways." Monitoring and evaluation of projects' progress along their impact pathways enables early identification of opportunities and challenges, which if acted on also makes impact more likely.

Finally, to effect change, research findings further need to be packaged into plausible promises. Without being packaged as plausible promises, key stakeholders will not engage with the novelty. Without engagement there will be no behavior change and no impact. Packaging of plausible promises is needed as much for research outputs such as policy briefs, models and methodologies as it is for rice harvesters and wind turbines. Packaging of plausible promises usually involves a learning selection process with the key stakeholders.

CONCLUSION

The world faces huge challenges in the 21st Century, of which triggering a Blue Revolution to improve water use in agriculture is one. Much of the response to these challenges will come through innovation. Research can and does enable innovation, but the way that research and innovation processes are conceptualized and managed makes a huge difference to the ability of engineers and researchers to foster change. The paper describes the learning selection model that can guide setting up and managing grassroots innovation processes that capitalize on complexity by building on peoples' ingenuity, motivations, and their implicit theories of how change occurs. Enabling innovation requires fostering change at different scales. The learning selection model can also help guide "top-down" changes by identifying three sets of interventions that alter innovation rate in a system: 1) changes to novelty and variation of actors and technologies in the system; 2) changes to interaction patterns between actors; and 3) changes to the way selection decisions are made. Traditionally agricultural research has attempted to leverage change by changing system novelty, through, for example, breeding new crop varieties. The learning selection model helps us see that bringing about a Blue Revolution is more about changing how people interact and make decisions, and less about developing new technology.

We invite reader comments. Email <editors@innovationsjournal.net>.

Acknowledgements

This paper is based on the author's book *Enabling Innovation: A Practical Guide to Understanding and Fostering Technological Change* (London: Zed Books), 2002. It reworks and adds to an article first published in 2002 in the CIGR Journal of Scientific Research and Development.[11] The author would like to thank Larry Harrington for writing the no-till case study in the Appendix and to Simon Cook for his contribution to the second part of the paper on the Blue Revolution.

APPENDIX. A GRASSROOTS INNOVATION THAT SAVES WATER[12]

No-till wheat after rice is currently being planted on over 3 million ha. of the Indo-Gangetic plains, producing net benefits of around USD 239 million per season, along with a 1.18 km^3 reduction in the extraction of irrigation water.[13]

Work on no-till began in the Indian Punjab in about 1970. It was restricted to hand-sown, small-plot, on-station trials, with little or no scientist-farmer interaction. Trials as designed masked the true advantages of zero-till: earlier sowing, higher yields, reduced costs, and improved weed control. In the mid-1980s, researchers in the Pakistan Punjab began a separate program. A key factor was access to a "best bet"—a no-till implement from New Zealand. These were, however, expensive, heavy and few in number. Although appropriate for research, they were not very well suited for farmer use. Researchers sponsored the development of a local prototype but this, also being too heavy, was not well received. It did not make a plausible promise of being useful.

From here, the initiative passed to India. An international scientist, familiar with the no-till work in Pakistan, donated four imported implements to colleagues in India. On-farm testing with one of these began in 1990 at Pantnagar University. Zero till wheat performed well, with good crop establishment, higher yields and lower costs. Nonetheless, progress was slow. The implement was again too heavy, and openers were prone to breakage and unable slice through rice stubble. In the following season, one of the scientists took a simple "recombination" step, attaching "inverted-t" cross-slot openers from the imported drill to his own frame design. This was the original "Pantnagar drill" and the first plausible promise.

As it happened, there was near Pantnagar a dealership for National Agro Industries, a Ludhiana-based farm implement company. The local dealer became aware of the Pantnagar drill and in 1992 introduced the researchers to one of the company's owners. Subsequently, National learned to forge its own inverted-t openers, which then were installed on the frame of a National conventional-till drill. This was the "Pantnagar drill" Mark II. National was soon joined by another company, Amritsar-based ASS Foundries. Several dozen design changes were progressively introduced, largely inspired by farmer testing. By 1995, a well-adapted design was ready for commercial production. And just at this time, an emergency occurred which sent researchers looking for just such a drill.

In the Indo-Gangetic Plains, continuous rice-wheat rotations favor a weed called Phalaris minor. For many years this weed caused little damage—farmers had learned to control it with isoproturon herbicides. But with millions of farmers using isoproturon for Phalaris control over many years, a herbicide-resistant Phalaris evolved. It was during the 1992-93 wheat season that a scientist working at Haryana Agricultural University (HAU) first reported such strains. By the 1995-96 season, the weed problem had been become a crisis. The affected area in Haryana continued to expand and began to move into neighboring states. Farmers grew desperate for a solution.

Some scientists felt that "desperate times call for desperate measures", with zero tillage being one of the "desperate measures". In order to test the effect of zero tillage and new herbicides on Phalaris populations, researchers needed zero till drills. The new Pantnagar drill had just become available. The newly-formed Rice-Wheat Consortium donated to HAU four new National no-till drills. These were delivered in October of 1996, as the Phalaris crisis peaked. With wheat sowing just weeks away, researchers moved to organize a research program, but encountered an unexpected obstacle: farmers placed a high value on tillage and wished to have nothing to do with "zero tillage". No-till was an invention that challenged a constraint long taken for granted: the need to till the soil to prepare for the next crop. In the end, however, a combination of new herbicides plus zero tillage worked well. Phalaris populations fell dramatically, and yields were excellent.

Although the very first zero till trials in Haryana were established by HAU scientists, farmer experimentation with the zero tillage soon followed. This happened for one simple reason: instead of returning the drills to the university campus after sowing the experiments, scientists decided to leave them in those villages where the trials were located. Farmers were encouraged to try out the drills on their own, and were provided with training and technical support. Learning selection began. Farmer experimenters found that zero till helped control Phalaris—but also very substantially reduced production costs. With this, adoption of no-till began to spread through farmers' networks.

This was spurred by the purchase (by ACIAR and the Rice Wheat Consortium) of additional drills for farmer testing. These still-scarce zero till drills were shifted from one village to another each wheat season. Farmers interested in no-till were invited to purchase their own drills—which they began to do in large numbers. By 1997, 150 drills had been sold to universities, ICAR institutions, and individual farmers. The Rice-Wheat Consortium, who became one of the product champions, began to organize study tours, whereby farmers from different districts—and even different States—could see for themselves the progress being made in Haryana on zero till wheat. There were even study tour participants from Pakistan!

Finally, adoption of zero till wheat in Haryana was further accelerated by an unexpected event. An agricultural department official was testing zero tillage on his own farm, with excellent results. In 1998, his daughter got married. The wedding happened during the wheat season and was celebrated in his own home. The no-till plots were located near the path leading up from the road. As a consequence, hundreds of the most influential farmers and state officials saw for themselves the extraordinary performance of zero till. This led to further adoption.

---

1. Mokyr J. (1990). *The Lever of Riches: Technological Creativity and Economic Progress* (Oxford and New York: Oxford University Press), p. 155.
2. von Hippel, Eric (1988). *The Sources of Innovation* (Oxford and New York: Oxford University Press), p. 3.
3. Douthwaite, Boru (2002). *Enabling Innovation: A Practical Guide to Understanding and Fostering Technological Change* (London: Zed Books)
4. Kolb, D.A. (1984). *Experiential Learning: Experiences as the Source of Learning and Development* (Englewood Cliffs, NJ: Prentice-Hall).
5. Raymond, Eric. (2001). *Cathedral and the Bazaar: Musings on Linux and Open Source by an Accidental Revolutionary* (O'Reilly Media; Revised Edition).
5. For an elaboration, see Elliot Maxwell, "Open Standards, Open Source, and Open Innovation: Harnessing the Benefits of Openness," in the previous issue of this journal.
6. Mokyr (1990), p. 9.
7. Rogers (2003) identifies five types of adopter (i) innovators, (ii) early adopters, (iii) early majority, (iv) late majority, and (vi) laggards. See Everett M. Rogers (2003), *Diffusion of Innovations* (New York: Free Press).
8. John Gall (1978). *Systemantics: How Systems Really Work and How They Fail* (New York: Pocket Books).
9. In particular, Axelrod, Robert and Michael D. Cohen (1999). *Harnessing Complexity: Organizational Implications of a Scientific Frontier* (New York: The Free Press).
10. Douthwaite, Boru (2002). "How to Enable Innovation: An Invited Overview Paper." *Agricultural Engineering International: the CIGR Journal of Scientific Research and Development*. Vol. IV. October.
11. This story is based on an unpublished manuscript by Larry Harrington titled "A brief history of zero and reduced tillage revolution in the Indo-Gangetic Plains."
12. Gupta, R. (2006). "RCT Induced Impacts in Indo-Gangetic Plains". *RWC Research Highlights 2006* (New Delhi: Rice Wheat Consortium for the Indo-Gangetic Plains). <www.rwc.cgiar.org/pubs/180/Highlights06.pdf>.