Chapter Three: How We Think works (and not)

Crowds are not automatically wise and mobs are not necessarily smart. It all depends on how they are made up and come together. Why do some collaborations turn into We Think - seemingly generating a momentum and intelligence of their own – while other do not? Why do any of them work? Viewed through the lens of traditional industrial era organisations, it seems improbably that skilled people, with busy lives, would give their time for free to mass collaborative efforts, only to give way the fruits of their labour. Or that their many contributions could be brought together into a coherent whole, rather than descending into chaos, fragmenting or getting bogged down in committees and debates about what to do next. So how do creative communities avoid becoming inward looking cliques that ignore new ideas brought by outsiders?

Successful attempts at We Think answer these basic organisational questions effectively. The participants do not have to buy into altruistic ethics or hippy ideals of community. We Think works, it delivers the goods, whether that is a new encyclopaedia, a software programme, a news service or even different ways of betting against other people. It does this when a cocktail of ingredients is brought together. When they are lacking it will not take off and no one should waste their time trying. A couple more examples will help explain how it works.

* * *

The worm *C elegans* is a simple organism: it has a front end, where the food comes in, a rear where the waste exits, a bottom and a top, a left and right. That is pretty much it, except that even the simplest worm achieves a mind-bogglingly complex task: it generates itself from a set of genetic instructions. The puzzle of how the worm achieves this task was unravelled by a collaborative research effort that in turn provided the basis for the global, public effort to map the human genome three decades later. Our map of the genome is the product of an elaborate shared authorship. Scientific collaborations like the one behind the unravelling of the C Elegans genome are a powerful working model for We Think culture that the web is helping to spread.

When Sydney Brenner set out to unravel the worm's genome in 1965, just eight years after Frances Crick and James Watson had uncovered the double-helix structure of DNA, little was known about how genes worked. Brenner set out to find out how the worm's genes directed the organism's growth with only a small team of novice researchers and some crude tools: the scientists lifted worms into petri dishes with sharpened tooth picks. It was as if someone had seen the Wright brothers' first flight and decided to start work on a jumbo jet.

Brenner's Laboratory of Molecular Biology provided the community's core. He had the resources to get started and just enough momentum to attract other

laboratories to collaborate on a project that intrigued others because it was so ambitious. The working practices of Brenner's small lab set the tone for what eventually became a global project involving thousands of researchers. Brenner's laboratory was hard working and meritocratic, egalitarian and conversational. People often discussed ideas in the coffee room. They were exploring new territory, devising the process as they went along, so there were no fiefdoms to defend. Sharing ideas quickly became normal. As the community grew, researchers communicate their progress through the relentlessly practical Worm Breeder's Gazette. (The Gazette was like a cross between the Lean Reporter, which organised innovation in the Cornish tin mines and Stewart Brand's Whole Earth Catalogue, which listed useful technologies.) Brenner's openness set off a virtuous cycle of knowledge sharing which was the only way to get the work done. He had identified a task so complex that no single laboratory could complete it. Knowledge about what a particular gene did, was worthless unless it could be combined with information about other genes. The jig-saw puzzle had so many pieces it could only be completed through collaboration on a massive scale. Bob Waterston, one of the US leaders of the explained:

"The more we put out there the less of a problem it was to get other people to contribute. The more we restricted the flow of knowledge, the more people felt they had to bargain with us before they would release their results. If you just put the data out there then everyone was on the same footing and they were all free to talk about it."

The community grew along with the common store of knowledge it created. In 1975, ten years after Brenner launched the project, the first international meeting of worm genome researchers attracted 24 participants. A decade later there was enough information to fit into a sizeable textbook. When the complete gene sequence was announced in 1998 US vice-president Al Gore greeted it as the equivalent of the moon landings. By 2002 the worm researchers' meeting attracted 1,600 participants. One thesis listed all 5,000 connections between neurons in the worm's brain. The worm was the most completely understood organism on earth.

Technology played its part in the project's success. By the end, researchers who had started out by using toothpicks were working with automated gene sequencing machines that could do in minutes tasks that had taken months in the mid 1960s. Yet the worm project – and the human genome project that followed in its footsteps - were as much a triumph of social organisation as technology. Eric Raymond, the guru of the open source software movement, famously described mass collaborative innovation as a bazaar – open, cacophonous, with no one in control – rather than a cathedral, where craftsmen implement a master plan. The worm project could never have been like building a cathedral. There was no master plan because no one knew for sure what they would find next and how it would fit together. Researchers were obliged to fan out and explore, to share their

ideas to piece together the map. Yet neither was the worm project like a bazaar. It only got going and sustained itself thanks to Brenner's leadership and a core of contributors. His ambition animated the community and led by example and his laboratory set the egalitarian, open yet challenging, style of working that characterised the whole project. He set the norms, releasing information early to encourage others to do likewise. Brenner made sure all the pieces fitted together to create something of lasting value.

The worm project is a recipe for We Think. Brenner found a way to mobilise a vast community of researchers and to combine their different skills and interests with little hierarchy and bureaucracy. His laboratory provided the community's core; the Worm Breeder's Gazette and frequent meetings provided a way for researcher to connect and combine their ideas; the open sharing of information allowed thousands of people to collaborate. More popular, non-scientific efforts at We Think are now replicating this recipe using the web. The most famous example is the open source software community Linux.

A single grain added to a pile of sand can cause it to collapse. Spotting which grain of sand that will be is virtually impossible. That is perhaps why most of the computer world did not notice in September 1991 when Linus Torvalds, then a wispy computer science student in Helsinki, released onto the Internet the first version of a computer programme he had written: Linux. Torvalds did not just put the programme online but also its source code, its basic recipe, leaving it for software enthusiasts to take away and tamper with, to criticise and propose improvements. Open source is software that nobody owns, everyone can use and anyone can improve and open source licensing is a way to hold ideas and information in common in a way that under the right conditions can encourage mass collaborative innovation. That is what Torvalds eventually set off an experiment in geek democracy, people thinking and creating together, that has created a programme that is complex, robust and now widely used. The Linux community is also the most impressive example of sustained We Think, with ideas shared among a large community over more than 15 years to develop a highly sophisticated but reliable product. Linux shows sharing makes sense, especially when ideas are at stake and the web is in play.

Like many radical innovations, Linux is not as new as it sounds. It draws on more than four decades of innovation in sharing software and computers but it arrived at just the right time. Linux drives Internet applications and as the Internet spread, so demand for it grew. In the five years from 1999 Linux installations grew at 28% a year and by 2006 Linux accounted for about 80% of software on computer servers worldwide. The 29m registered Linux users in the world by 2006 vastly underestimates of the true user base for open source software. Everytime anyone runs a search on Google they are a Linux user because Google's servers run on Linux.

Linux's market share expanded as the software became more complex. The version distributed in March 2000 by Red Hat, which specialises in installing Linux for corporations, had 17m lines its source code. One estimate suggests it would have taken 4,500 person years of work for professional software coders to develop this at a cost of about $600m. The version released by Debian in June 2005 had 229m lines in its source code and would have taken 60,000 man years of software coding at a cost of perhaps $8bn. To put that in context the version of Microsoft Windows XP released in 2002 had an estimated 40m lines in its source code. IBM is just one software company that has recently started investing in shared, open source platforms such as Linux. IBM executives estimate it would have cost the company ten times as much to develop Linux itself as it does to participate in the community that had already developed it. (quote from von Hippel conference..)

The Linux community has sustained this phenomenal growth by adopting the organisational features of We Think. At the core is a small band of trusted programmers working closely with Torvalds, a quietly inspirational, leader guiding by example. Membership of this core group, which looks after the programme's kernel and its future development, is earned by putting in long hours of high quality volunteer programming. In 1994 there were just 80 people in the core; by 2001 there were 400. A much larger community of users and contributors has formed beyond this. By 2007 there were 655 Linux user groups in 91 countries sharing ideas ideas through websites and bulletin boards as well as face to face at conferences.

Were Linux a one off it would be interesting and exotic. But it is not. Most of the Internet relies on open source software, much of its created through collaboration: most websites rely on servers running the open source Apache programme; MySQL is an open source data base programme; Perl and Python are open programming languages. There are many more in their wake with names like Dupral, Evolt, Tomcat and JBoss. In 2007 the directory Sourceforge.net which listed more than 90,000 open source initiatives.

Linux has succeeded as a product only because the community that supports it has organised itself systematically to create, share, test, reject and develop ideas in a way that flouts conventional wisdom. Successful We Think projects are based on five key principles that were all present in Linux and the worm project.

Core

Everything has to start somewhere. Somebody has to be willing to work harder than everyone else or nothing ends up getting done. Innovative communities invariably start with a gift of knowledge provided by someone just as Linux started with the kernel that Linus Torvalds slaved over and which he posted on the Internet.

A good core attracts a community of capable contributors and developers around it. The kernel has to be solid but unfinished, so open to improvement. If it were complete there would be few opportunities to add to it. Jane McGonigal says the core to a successful game like *I Love Bees* depends on the starting point being ambiguous and open to interpretation. Both the worm project and *I Love Bees* began with a puzzle that could only be solved with the collaborative efforts of people with different skills. Steven Weber, a political scientist at Berkeley University in California, found that successful open source software projects tended to be "multi-dimensional" and complex, thus inviting the involvement people with different skills. (refs Success of Open source) Thomas Khun summed up the ambiguous character of the core to a new intellectual community in his history of scientific revolutions. Kuhn argued the possibility of a new scientific paradigm emerged when a small group of pioneers made a breakthrough that was:

"sufficiently unprecedented to attract an enduring group of adherents away from competing modes of scientific activity. Simultaneously, it was sufficiently open ended to leave all sorts of problems for the redefined group of practitioners to resolve."

A core will, however, only develop if its creators give away the material on which others can work, to which they can add and which they can refine. Successful innovation comes from a creative conversation between people who combine their different skills, insights and knowledge to explore a problem. We Think is creating a new way for these conversations to emerge. A good core starts a creative conversation, and invites people to contribute.

Contribute

A successful creative community has to attract the right mix of people, who have different ideas and outlooks and access to tools that enable them to contribute. We Think only takes off by getting the right answer to each of the following questions: who contributes; what do they contribute; why do they do so; and how do they do it?

Creative communities have a social structure. As we have seen a relatively small, committed core group tends to do most of the heavy lifting: the discussion moderators in Slashdot; the original inhabitants of Second Life. These are the Web 2.0 aristocracy: people who because they have been around longer and done more work, tend to get listened to more. There is nothing unusual in this. Most innovative projects, whether inside a company a theatre group or a laboratory, start with intense collaboration among a small group who share a particular passion or want to address a common problem, as did the worm researchers who gathered around Sydney Brenner at Cambridge. Often, however, such communities can become closed and inward looking. To be dynamic, they have to open out to a wider world of more diverse contributors who add their knowledge or challenge conventional wisdom.

We Think projects take off when they attract a much larger crowd, who are less intensely engaged with the project. Their occasional, smaller contributions may in aggregate be as the significant as the work initially done by the core. As well as 400 key programmers at the core, Linux for example, has close on 150,000 registered users – akin to members - who may only report the occasional bug in the programme. Yet such a report may provide the starting point for a much more significant effort at innovation. The make up of the crowd is as important as brainpower of the highly committed core. Crowds are intelligent only when their members have a range of views and enough self-confidence and independence to voice their opinions. Scott Page, a professor of complex systems at University of Michigan has used sophisticated computer models to find that groups with diverse skills and outlooks came up with smart solutions more often than groups of very clever people who shared the same outlook and skills. Groups made up of many people who think in different ways can trump groups of people who are very bright but very alike, Page argues, so long as they are organised in the right way.

Page's explanation is that the more vantage points there are from which a complex problem is seen the easier it becomes easier to solve. A group of experts who think in the same way are probably no better at devising a solution than just one of them, so adding more people who think in the same way is unlikely to improve a group's ability to come up with different solutions. Groups who think the same way can often find themselves stuck at the same point – akin to the peak of a foothill in a mountain range - unable to climb to the higher peaks that lie beyond. A group that thinks in diverse ways, in contrast, is more likely to address a problem from many angles, less likely to get stuck and more likely to find a way out if it does get stuck. Diverse viewpoints are likely to generate more possible solutions and evaluate them in a wider range of ways. The right perspective can make a difficult problem seem easy. Innovation often involves trying out many vantage points on a problem before finding the one that makes the problem look simple. As Thomas Edison put it:

"We have found 1,000 ways not to make a light bulb."

Bugs in a software programme often become apparent only when the programme is tested in many different settings. Better 1,000 people making different tests at the same time, than a single person making 1,000 tests one after the other. That explains why open source programmes are often more robust that proprietary software: it has been tested much earlier by a much wider group of users. Bart Nooteboom, a professor at Rotterdam University, argues that distributed testing of this kind is vital to most innovation. He examines the development of 17th century Dutch sailing ships and found the designs mutated when the community of sailors, tested and then adapted them to meet different conditions: first canals, then lakes, larger inland waterways, offshore sailing, the North Sea, the Atlantic and so on. We Think allows ideas to be tested from a larger, more diverse set of

vantage points more quickly and when ideas continually pass between the tight knit core who develop them and the crowd who test them out.

This testing only becomes possible when people can make the kind of contributions they feel happy with, which requires tools to allow them to get involved. Mass computer games thrive by making it easy for player-developers to pick up tools to create content. Blogging depends on easy to use software for writing and publishing online. The camera phone is now a ubiquitous tool for citizen journalism. Such tools are taking to mass scale the self-help ethic of the original computer hackers. Lone programmers created the first versions of the Unix operating system on which Linux is based and could not afford to provide tech support to their users. So when they sent their programmes to people, usually on a stack of floppy discs, they included a set of tools that allowed users to sort out problems themselves. When people can get hold of tools that allow them to produce aspects of a service, they start becoming players, participants and developers: newspaper readers become writers, publishers and distributors; bystanders become photographers; the audience can become reviewers and critics.

Perhaps the most perplexing question is not how people contribute, but why they do so, especially when they are not being paid and their work is given away? In open source software projects, a few are inspired by a hatred of proprietary software providers, especially Microsoft. A minority are driven by altruistic motives. Some see their involvement as a way to get a better job: by showing off their skills in the open source community they can enhance their chances of being employed. For the majority the main motivation is recognition: they want the acknowledgement of their peers for doing good work that they enjoy, which gives them a sense of achievement and in the process solves a problem for which other people are seeking a solution. Many of the most striking Web 2.0 success stories started when users created tools to solve a problem they faced – keeping track of all the blogs being created, sharing video and photographs online – and which quickly got taken up by others who faced similar problems.

Open source gives away intellectual property so other people can freely use it. We Think requires more than that: it is also an invitation to participate and collaborate in creating something. Open source ownership of a project becomes powerful when it enables mass collaborative approaches to innovation. For that to be possible many ideas have to be combined; contributors have to meet and connect with one another.

Connect

At the St Louis world fair in 1904 an ice cream stand ran out of cups. The owner of the waffle stand next door started rolling his waffles to form cones. There was nothing new in either ingredient, but the combination of ice cream and the waffle created something entirely new. The more combinations a community can create, the more innovation there will be. Cities are creative when they make these

combinations possible. The same is true of We Think.

Diversity counts for little unless the different ideas that are floating around can be brought together to cross-pollinate. A community that is diverse but balkanised will not be creative. People with different ideas must find a way to connect and communicate with one another. When they do, in the right way, the results can be explosive. James Watson and Francis Crick unravelled the double helix structure of DNA because they found a way to combine their very different outlooks. Crick's training spanned physics, biology and chemistry. Watson had trained as a zoologist but became fascinated by DNA after studying viruses. They combined their ideas through constant, intense conversation the kind of which their rivals were incapable. Watson and Crick's collaboration was a case of one plus one equals twelve.

The larger the group and the more diverse perspectives there are involved the greater the benefits from combining them. Take five people, each with a different skill. You have 10 possible pairings of skills. Add a sixth person with a different skill. That does not create 12 pairs but another 5 possible pairings. A group with twenty different tools at its disposal has 190 possible pairs of tools and more than 1,000 combinations of three tools. A group with 13 tools has almost as many tools – 87% - of a group with 15 tools. Not much of a gap. But if a task requires combining four tools it is a different story. The group with 15 tools has 1,365 possible combinations of four tools. The group with 13 tools has 715, or about 52%. Groups with larger sets of diverse tools and skills are at an advantage if they can combine effectively to take on complex tasks.

Markets are not the best way for people with diverse skills can connect and combine. A market might provide a way for someone with a problem to find someone else who might have a solution: if you have a leaking tap you look for a plumber. That is the model of Innocentive, the scientific problem solving community that was spun off from the drugs company Eli Lilly. Companies can post their scientific problems on Innocentive's web site to see if they can be solved by one of the more than 100,000 scientists signed up to the market. But markets of this kind have inherent limitations: they work for specific problems that need exactly the right individual to solve them. They do not provide the basis for sustained creativity and innovation to explore difficult complex puzzles. That is a kind of problem solving that only comes from intense collaboration. In the worm project, the researchers started by meeting in the coffee room at Brenner's laboratory. In We Think crowds need meeting places, neutral spaces for creative conversation, moderated to allow the free flow of ideas. That is why, at their heart, these projects they have open discussion forums and wikis, bulletin boards and community councils or simple journals like the *Lean Engine Reporter* and the *Worm Breeder's Gazette*, so that people can come together in a way that allows one plus one to equal twelve many times over.

In We Think projects, the task of combining ideas is made easier because the products usually fit together like Lego bricks: they are made from many interconnecting modules. Modularity is not new; it has been a feature of computer development since at least the 1960s, when IBM was developing its system 360 computer. Fred Brooks, the person responsible, wanted everyone involved to be kept abreast of what everyone else was doing. Daily notes of changes to the programme were shared with everyone. Quite soon people were starting work each day by sifting through a two-inch wad of notes on design changes. The costs of communication and coordination spiralled out of control. Miscommunication and misunderstandings grew. Adding people to the project did not solve the problem: more work got done, but more misunderstandings were created and with them more bugs. When the wad was five feet thick Brooks decided to break the S360 into discrete modules which could be worked on separately. A core team set some design rules that which specified what modules were needed and how they should to click together. That meant module makers could concentrate on their patch while the core team looked after the architecture of the system as a whole. New and better modules could be fitted into the system without having to redesign it from scratch.

Modularity really pays dividends when it is combined with open ways of working. When it enables a mass of experiments to proceed in parallel, with different teams working on the same modules, each proposing different solutions. This combination is how open source gets the Holy Grail: a mass of decentralised innovation that all fits together. Just as Lego bricks come in a dizzying array of colours, shapes and even sizes, but they all have the same system of connectors, We Think projects have rules for making connections that usually come from the core team. This is what allows a mass of independent but interconnected innovation. Mass computer games, collaborative blogs, open source programmes, the human genome project all share this feature: they click together masses of modules.

A Lego brick structure is not however enough to make We Think work. Groups also need to make decisions. Diverse contributors can combine their ideas only if they can agree how to collaborate. Any commons will fall into disrepair if it is not effectively self-regulated. That is far easier said than done.

Collaborate

A mass of contributions does not amount to anything unless together they create something ordered and complex. An encyclopaedia is not a mass of random, individual contributions; it’s structured account of knowledge. People playing a game or building a community need to agree rules to govern themselves or chaos ensues. How do We Think communities govern themselves without an obvious hierarchy being in charge, enforcing the law? This challenge is not technical but political. We Think works only when it has responsible self-governance, and that is a particularly thing to achieve in highly diverse communities.

People often think in different ways because they have very different values; what matters to them differs. Someone who sees the world through art and images will acquire skills – drawing and painting – which make it easier for them to work. Someone who sees the world in numbers and money is more likely to become an accountant, to use a calculator rather than a paintbrush. A large tool box that includes both calculators and paintbrushes, artists and accountants is good for innovation.

The trouble is that people with fundamentally different values often find it difficult to agree on what they should do and why. Diverse ways of thinking are essential for innovation; diverse values, based on differences about what matters to us, often lead to squabbles. That is why diverse communities often find it more difficult agree on how to provide public goods, such as health care, welfare benefits and social housing. Diverse groups can become very unproductive when their differences overwhelm them, provoking conflicts over resources or goals. Elinor Ostrom found that shared fisheries, forest and irrigation systems required effective self governance and local monitoring by participants to make sure no one was over-using resources. When local self governance fails, the commons collapses and innovation becomes impossible.

We Think succeeds by creating self-governing communities that make the most of their diverse knowledge without being overwhelmed by the differences between people. That is possible only if these communities are joined around a simple animating goal; they develop legitimate ways to review and sort ideas and they have the right kind of leadership. What they are not, ever, is egalitarian self governing democracies.

As an example consider the open source community which produces Ubuntu, a user friendly version of Linux. Mark Shuttleworth, Ubuntu's founder is like a benevolent dictator and reserves some decisions to himself, such as the design of the Ubuntu website. The heart of the community, the Technical Board, meets online to set technical standards and to define what should be included in the different versions of the programme. The Board's decision making is transparent and open: anyone can propose additions to these policies through the Ubuntu wiki, the board's agenda is available as a wiki every two weeks and anyone can attend the online meetings as an observer. The decisions are taken, however, by Shuttleworth and four other board members, who he appoints, albeit subject to a vote among the community's lead programmers. A separate Ubuntu community council, meanwhile, supervises the social structure, creating new projects and appointing leaders for teams that support different releases and features of the programme, like those for laptop users for example. Then there are the LoCo teams around the world which promote the use of Ubuntu in their country. Someone can become an Ubuntu member (an Ubuntero) by making by coding software, documenting changes, contributing artwork or acting as an advocate for

Ubuntu. As of mid 2007 the community had 283 core members. Those with most power and responsibility – dubbed Masters of the Universe – are the core developers and they have their own council to determine who should be allowed into their guild.

The lesson of Ubuntu – which is still far from a proven success – is that effective governance of creative communities is like a lattice work. Decision making is very open: anyone can see what is decided and how, anyone can make suggestions about what should be done. But the way decisions are made is rarely democratic. Ubuntu the product may be open source; the community that sustains it is far from open ended. These are not like the utopian communes of the 1960s: which is why they might be more successful than cooperatives of the past.

Create

We Think enables a mass social creativity which thrives when many players, with differing points of view and skills, the capacity to think independently and tools to contribute are brought together in a common cause. If the players are distributed they must have a way to share, combine and cohere around a common goal. However for much of the time, contributors may work independently and in parallel, often reworking elements of a core central product – whether that's an epic poem in Ancient Greece, a piece of genetic code, a latter day software programme or an encyclopaedia. The product grows through accretion and a reciprocal process of observation, criticism, support and imitation. Most people take part because they get an intrinsic pleasure from the activity and seek recognition from their peers for the work they have done. These communities must have places – forums, web sites, festivals, gazettes and magazines – where people can publish and share ideas. Social creativity is not a free for all; it is highly structured. Although the lines between expert and amateur, audience and performer, user and producer may be blurred, those with more standing in the community, based on the history and quality of their contribution, from something like a tightly networked craft aristocracy. Social creativity collapses without effective self governance: decisions have to be made about what should be included in the source code, published on the site, pushed to the top of the news list. Participants who do not abide by the community's rules have to be excluded somehow. They must respect the judgements of their peers.

The raw material of these collaborations is creative talent, which is highly variable. People are good at different things and in different ways. It is difficult to tell from the outside, for example by time and motions studies, who is the more effective creative worker. It is impossible to write a detailed description for a creative job specifying what new ideas need to be created by whom and by when. Open source communities resolve the difficulties of managing creative work by decentralising decision making down to small groups which decide what to work on, depending on what needs to be done and the nature of their skills. It is very difficult for someone to pull the wool over the eyes of their peers: they will soon

be found out. When it works peer review excels at the sharing ideas and maintaining quality at low cost.

Conclusions

We Think will not work where there is no core around which a community can form; where experimentation is costly and time consuming, and so feedback is slow; where decision making becomes cumbersome or opaque, beset by complex rules; where the project fails to attract a large and diverse enough community. It will not take off if tools to add content are difficult to use; if contributors cannot connect to one another, if communities cannot govern themselves effectively and either fracture or ossify. For many important activities, We Think will make no sense at all: medical operations, cooking meals, running nuclear reactors, railways or steel mills. It is not well suited to tasks where only professional expertise will do. In late 2006 I had a minor operation and was very glad to find the surgeon was not assisted by a group of Pro Am butchers, bakers and candlestick makers who were taking their lead from the Wikipedia entry on the procedure they were about to perform on me.

We Think only works under certain conditions. Usually, a small group creates a kernel that invites further contributions. Their project must be regarded as exciting, intriguing and challenging by enough people with the time, means and motivation to contribute. Tools should be distributed, experimentation cheap and feedback fast, enabling a constant process of trialling, testing and refinement. The product should benefit from extensive peer review, to correct errors and verify good ideas. Tasks should be broken down into modules around which small, close-knit teams can form allowing a range of experiments to run in parallel. There should be clear rules for how the modules for fitting the modules together and separating good ideas from bad. Ownership of the project must have a public component, otherwise the sharing of ideas will not make sense.

It is not all or nothing but a matter of degree: from No We Think at one end of the spectrum, where traditional, closed and hierarchical models of organisation still make sense, to Full We Think at the other end, with the likes of Linux and Wikipedia. In the middle, there will be lots of opportunities to blend some of these ingredients in different ways.

Blogging is a prime example: it allows a mass of people to contribute their views, but only rarely do they find a core to build around. Mostly, bloggers communicate into the ether. They have no desire to build something with others, merely to leave their mark on their little patch of digital space. Blogging is high on participation, low on collaboration. Flickr, the photo sharing site, and YouTube, the video site, fit in this Low We Think category: they allow a mass of participants to connect with an audience and one another. Yet there is relatively little collaborative creativity. When YouTube becomes a platform for people to collaborate in making films together it may acquire some of the features of We Think.

Social networking is Medium We Think. Sites such as MySpace, CyWorld and Bebo have not yet encouraged much deliberate collaborative creativity, although some participants have started to use them for example to support political candidates or rally around causes they care about. Collaborative filtering and the book reviews and ratings on Amazon, and social tagging tools like Technorati and del.i.cious, through which people help one another find interesting material on the web fit in this category.

Only when all our five conditions come together at scale to provide a deliberate, conscious form of social creativity in which many people contribute and collaborate does Full We Think emerge. Oh My News, the south Korean citizen journalist news service fits in here, as do mass computer games like World of Warcraft and scientific collaborations like the project to unravel the worm's genome. Full We Think is the deliberate and organised combination of contributions from a mass of distributed and independent participants.

It would be silly to suggest We Think can work in every situation and that it is always the best organisational recipe. The challenge is to produce more We Think when it is appropriate which is when we are collectively trying to solve a complex problem, create something that no individual could produce and critical to develop ideas. We Think will not touch all organisations but some will be transformed, and many will find some aspects of what they do changed, possibly quite fundamentally by this new organisational recipe. How We Think will change the way we run, lead and own organisations is the subject of the next chapter.