

# REINVENTING DISCOVERY

Copyright © 2012 by Michael Nielsen  
Requests for permission to reproduce material from this work should be sent to  
Permissions, Princeton University Press

Published by Princeton University Press, 41 William Street, Princeton, New Jersey 08540  
In the United Kingdom: Princeton University Press, 6 Oxford Street, Woodstock,  
Oxfordshire OX20 1TW  
press.princeton.edu

All Rights Reserved

Fourth printing, and first paperback printing, 2014  
Paperback ISBN 978-0-691-16019-1

The Library of Congress has cataloged the cloth edition of this book as follows

Nielsen, Michael A., 1974–  
Reinventing discovery : the new era of networked science / Michael Nielsen.  
p. cm.

Includes bibliographical references and index.

ISBN 978-0-691-14890-8 (hardback)

1. Research—Technological innovations. 2. Discoveries in science. 3. Internet.  
4. Information networks. 5. Information technology. I. Title.

Q180.55.M4N54 2011

509—dc23 2011020248

British Library Cataloging-in-Publication Data is available

This book has been composed in Minion  
Printed on acid-free paper ∞

Typeset by S R Nova Pvt Ltd, Bangalore, India  
Printed in the United States of America

10 9 8 7 6 5 4 3 2 1

## Contents

<a href="#">1. Reinventing Discovery</a>	1
<b><a href="#">PART 1. AMPLIFYING COLLECTIVE INTELLIGENCE</a></b>	<b>13</b>
<a href="#">2. Online Tools Make Us Smarter</a>	15
<a href="#">3. Restructuring Expert Attention</a>	22
<a href="#">4. Patterns of Online Collaboration</a>	44
<a href="#">5. The Limits and the Potential of Collective Intelligence</a>	69
<b><a href="#">PART 2. NETWORKED SCIENCE</a></b>	<b>89</b>
<a href="#">6. All the World's Knowledge</a>	91
<a href="#">7. Democratizing Science</a>	129
<a href="#">8. The Challenge of Doing Science in the Open</a>	172
<a href="#">9. The Open Science Imperative</a>	187
<a href="#"><i>Appendix: The Problem Solved by the Polymath Project</i></a>	209
<a href="#"><i>Acknowledgments</i></a>	215
<a href="#"><i>Selected Sources and Suggestions for Further Reading</i></a>	217
<a href="#"><i>Notes</i></a>	221
<a href="#"><i>References</i></a>	239
<a href="#"><i>Index</i></a>	255



# REINVENTING DISCOVERY



## CHAPTER 1

### Reinventing Discovery

Tim Gowers is not your typical blogger. A mathematician at Cambridge University, Gowers is a recipient of the highest honor in mathematics, the Fields Medal, often called the Nobel Prize of mathematics. His blog radiates mathematical ideas and insight.

In January 2009, Gowers decided to use his blog to run a very unusual social experiment. He picked out an important and difficult unsolved mathematical problem, a problem he said he'd "love to solve." But instead of attacking the problem on his own, or with a few close colleagues, he decided to attack the problem completely in the open, using his blog to post ideas and partial progress. What's more, he issued an open invitation asking other people to help out. Anyone could follow along and, if they had an idea, explain it in the comments section of the blog. Gowers hoped that many minds would be more powerful than one, that they would stimulate each other with different expertise and perspectives, and collectively make easy work of his hard mathematical problem. He dubbed the experiment the Polymath Project.

The Polymath Project got off to a slow start. Seven hours after Gowers opened up his blog for mathematical discussion, not a single person had commented. Then a mathematician named Jozsef Solymosi from the University of British Columbia posted a comment suggesting a variation on Gowers's problem, a variation which was easier, but which Solymosi thought might throw light on the original problem. Fifteen minutes later, an Arizona high-school teacher named Jason Dyer chimed in with a thought of his own. And just three minutes after that, UCLA mathematician Terence Tao—like Gowers, a Fields medalist—added a comment. The comments

erupted: over the next 37 days, 27 people wrote 800 mathematical comments, containing more than 170,000 words. Reading through the comments you see ideas proposed, refined, and discarded, all with incredible speed. You see top mathematicians making mistakes, going down wrong paths, getting their hands dirty following up the most mundane of details, relentlessly pursuing a solution. And through all the false starts and wrong turns, you see a gradual dawning of insight. Gowers described the Polymath process as being “to normal research as driving is to pushing a car.” Just 37 days after the project began Gowers announced that he was confident the polymaths had solved not just his original problem, but a harder problem that included the original as a special case. He described it as “one of the most exciting six weeks of my mathematical life.” Months’ more cleanup work remained to be done, but the core mathematical problem had been solved. (If you’d like to know the details of Gowers’s problem, they’re described in the appendix. If you just want to get on with reading this book, you can safely skip those details.)

The polymaths aren’t standing still. Since Gowers’s original project, nearly a dozen Polymath and Polymath-like projects have been launched, some attacking problems even more ambitious than Gowers’s original problem. More than 100 mathematicians and other scientists have participated; mass collaboration is starting to spread through mathematics. Like the first Polymath Project, several of these projects have been great successes, really driving our understanding of mathematics forward. Others have been more qualified successes, falling short of achieving their (sometimes extremely ambitious) goals. Regardless, massively collaborative mathematics is a powerful new way of attacking hard mathematical problems.

Why is mass online collaboration useful in solving mathematical problems? Part of the answer is that even the best mathematicians can learn a great deal from people with complementary knowledge, and be stimulated to consider ideas in directions they wouldn’t have considered on their own. Online tools create a shared space where this can happen, a short-term collective working memory where ideas can be rapidly improved by many minds. These tools enable us to scale up creative conversation, so connections that would ordinarily require fortuitous serendipity instead happen as a



giant edifice. That edifice is too vast to be comprehended by any individual working alone. But new computerized tools can help us find meaning hidden in all that knowledge.

If the Polymath Project illustrates a shift in how scientists collaborate to create knowledge, and GenBank and the genetic studies illustrate a shift in how scientists find meaning in knowledge, a third big shift is a change in the relationship between science and society. An example of this shift is the website Galaxy Zoo, which has recruited more than 200,000 online volunteers to help astronomers classify galaxy images. Those volunteers are shown photographs of galaxies, and asked to answer questions such as “Is this a spiral or an elliptical galaxy?” and “If this is a spiral, do the arms rotate clockwise or anticlockwise?” These are photographs that have been taken automatically by a robotic telescope, and have never before been seen by any human eye. You can think of Galaxy Zoo as a cosmological census, the largest ever undertaken, a census that has so far produced more than 150 million galaxy classifications.

The volunteer astronomers who participate in Galaxy Zoo are making astonishing discoveries. They have, for example, recently discovered an entirely new class of galaxy, the “green pea galaxies”—so named because the galaxies do, indeed, look like small green peas—where stars are forming faster than almost anywhere else in the universe. They’ve also discovered what is believed to be the first ever example of a quasar mirror, an enormous cloud of gas tens of thousands of light-years in diameter, which is glowing brightly as the gas is heated by light from a nearby quasar. In just three years, the work of the Galaxy Zoo volunteers has resulted in 22 scientific papers, and many more are in the works.

Galaxy Zoo is just one of many online citizen science projects that are recruiting volunteers, most of them without scientific training, to help solve scientific research problems. We’ll see examples ranging across science, from volunteers who are using computer games to predict the shape of protein molecules, to volunteers who are helping understand how dinosaurs evolved. These are serious scientific projects, projects where large groups of volunteers with little scientific training can attack scientific problems beyond the reach of small groups of professionals. There’s no way a team of professionals could do what Galaxy Zoo does—even working full

time, the pros don't have the time to classify hundreds of thousands (or more) of galaxies. You might suppose they'd use computers to classify the galaxy images, but in fact the human volunteers classify the galaxies more accurately than even the best computer programs. So the volunteers at projects such as Galaxy Zoo are expanding the boundary of what scientific problems can be solved, and in so doing, changing both who can be a scientist and what it means to be a scientist. How far can the boundary between professional and amateur scientist be blurred? Will we one day see Nobel Prizes won by huge collaborations dominated by amateurs?

Citizen science is part of a larger shift in the relationship between science and society. Galaxy Zoo and similar projects are examples of institutions that are bridging the scientific community and the rest of society in new ways. We'll see that online tools enable many other new bridging institutions, including open access publishing, which gives the public direct access to the results of science, and science blogging, which is helping create a more open and more transparent scientific community. What other new ways can we find to build bridges between science and the rest of society? And what will be the long-run impact of these new bridging institutions?

The story so far is an optimistic story of possibility, of new tools that are changing the world. But there's a problem with this story, some major obstacles that prevent scientists from taking full advantage of online tools. To understand the obstacles, consider the studies linking genes to disease that we discussed earlier. There's a crucial part of that story which I glossed over, but which is actually quite puzzling: *why* is it that biologists share genetic data in GenBank in the first place? When you think about it, it's a peculiar choice: if you're a professional biologist it's to your advantage to keep data secret as long as possible. Why share your data online before you get a chance to publish a paper or take out a patent on your work? In the scientific world it's papers and, in some fields, patents that are rewarded by jobs and promotions. Publicly releasing data typically does nothing for your career, and might even damage it, by helping your scientific competitors.

In part for these reasons, GenBank took off slowly after it was launched in 1982. While many biologists were happy to access others' data in GenBank, they had little interest in contributing

their own data. But that has changed over time. Part of the reason for the change was a historic conference held in Bermuda in 1996, and attended by many of the world's leading biologists, including several of the leaders of the government-sponsored Human Genome Project. Also present was Craig Venter, who would later lead a private effort to sequence the human genome. Although many attendees weren't willing to unilaterally make the first move to share all their genetic data in advance of publication, everyone could see that science as a whole would benefit enormously if open sharing of data became common practice. So they sat and talked the issue over for days, eventually coming to a joint agreement—now known as the Bermuda Agreement—that all human genetic data should be immediately shared online. The agreement wasn't just empty rhetoric. The biologists in the room had enough clout that they convinced several major scientific grant agencies to make immediate data sharing a mandatory requirement of working on the human genome. Scientists who refused to share data would get no grant money to do research. This changed the game, and immediate sharing of human genetic data became the norm. The Bermuda agreement eventually made its way to the highest levels of government: on March 14, 2000, US President Bill Clinton and UK Prime Minister Tony Blair issued a joint statement praising the principles described in the Bermuda Agreement, and urging scientists in every country to adopt similar principles. It's because of the Bermuda Agreement and similar subsequent agreements that the human genome and the HapMap are publicly available.

This is a happy story, but it has an unhappy coda. The Bermuda Agreement originally only applied to human genetic data. There have since been many attempts to extend the spirit of the agreement, so that more genetic data is shared. But despite these attempts, there are still many forms of life for which genetic data remains secret. For example, as of 2010 there is no worldwide agreement to share data about the influenza virus. Steps toward such an agreement remain bogged down in wrangling among the leading parties. To give you the flavor of how many scientists think about sharing non-human genetic data, one scientist recently told me that he'd been "sitting on a genome" for an entire species (!) for more than a year. Without any incentive to share, and with many reasons

not to, scientists hoard their data. As a result, there's an emerging data divide between our understanding of life-forms such as human beings, where nearly all genetic data are available online, and life-forms such as influenza, where important data remain locked up.

This story makes it sound as though the scientists involved are greedy and destructive. After all, this research is typically paid for using public funds. Shouldn't scientists make their results available as soon as possible? There's truth to these ideas, but the situation is complex. To understand what's going on, you need to understand the incredible competitive pressures on ambitious young scientists. On the rare occasion a good long-term job at a major university opens up, there are often hundreds of superbly-qualified applicants. Competition for jobs is so fierce that eighty-hour-plus workweeks are common among young scientists. As much of that time as possible is spent working on the one thing that will get such a job: amassing an impressive record of scientific papers. Those papers will bring in the research grants and letters of recommendation necessary to find long-term employment. The pace relaxes after tenure, but continued grant support still requires a strong work ethic. The result is that while many scientists agree in principle that they'd love to share their data in advance of publication, they worry that doing so will give their competitors an unfair advantage. Those competitors could exploit that knowledge to rush their results into print first, or, worse, even steal the data outright and present the results as their own. It's only practical to share data if everyone is protected by a collective agreement such as the Bermuda agreement.

A similar pattern has seen scientists resist contributing to many other online projects. Consider Wikipedia, the online encyclopedia. Wikipedia has a vision statement to warm a scientist's heart: "*Imagine a world in which every single human being can freely share in the sum of all knowledge. That's our commitment.*" You might think Wikipedia was started by scientists eager to share all the world's knowledge, but you'd be wrong. In fact, it was started by Jimmy "Jimbo" Wales, who at the time was cofounder of an online company mostly specializing in adult content, and Larry Sanger, a philosopher who left academia to work with Wales on online encyclopedias. In the early days of Wikipedia there was little involvement from

scientists. This was despite the fact that anyone in the world can edit Wikipedia, and, in fact, it's written entirely by its users. So here's this incredibly exciting project, which anyone can get involved in, which is taking off rapidly, and which expresses core scientific values. Why weren't scientists lining up to be involved? The problem is the same as with the genetic data: why would scientists take the time to contribute to Wikipedia when they could be doing something more respectable among their peers, like writing a paper? That's the kind of activity that leads to jobs, grants, and promotions. It doesn't matter that contributing to Wikipedia might be more intrinsically valuable. In the early days work on Wikipedia was seen by scientists as frivolous, a waste of time, as not being serious science. I'm happy to say that this has changed over the years, and today Wikipedia's success has to some extent legitimized work on it by scientists. But isn't it strange that the modern-day Library of Alexandria came from outside academia?

There's a puzzle here. Scientists helped create the internet and the world wide web. They've taken enthusiastically to online tools such as email, and pioneered striking projects such as the Polymath Project and Galaxy Zoo. Why is it that they've only reluctantly adopted tools such as GenBank and Wikipedia? The reason is that, despite their radical appearance, the Polymath Project, Galaxy Zoo, and similar undertakings have an inherent underlying conservatism: they're ultimately projects in service of the conventional goal of writing scientific papers. That conservatism helps them attract contributors who are willing to use unconventional means such as blogs to more effectively achieve a conventional end (writing a scientific paper). But when the goal isn't simply to produce a scientific paper—as with GenBank, Wikipedia, and many other tools—there's no direct motivation for scientists to contribute. And that's a problem, because some of the best ideas for improving the way scientists work involve a break away from the scientific paper as the ultimate goal of scientific research. There are opportunities being missed that dwarf GenBank and Wikipedia in their potential impact. In this book, we'll delve into the history and culture of science, and see how this situation arose, in which scientists are often reluctant to share their ideas and data in ways that speed up the advancement of science. The good news is that we'll find leverage



## **PART 1**

---

### **Amplifying Collective Intelligence**





## CHAPTER 2

### Online Tools Make Us Smarter

In 1999, world chess champion Garry Kasparov played a game of chess against “the World.” In this event, organized by Microsoft, the idea was that anyone in the world could go to the game website, and vote on what move should be taken next. On a typical move more than 5,000 people voted, and over the entire game 50,000 people from 75 countries voted. The World Team decided on a new move every 24 hours, and on any given turn the move taken was whichever got the most votes. The game was billed as “Kasparov versus the World.”

The game exceeded all expectations. After 62 moves of innovative chess, in which the balance of the game changed several times, the World Team finally resigned. Kasparov called it “the greatest game in the history of chess,” and revealed that during the game he often couldn’t tell who was winning and who was losing; it wasn’t until the 51st move that the balance swung decisively in his favor. After the game, Kasparov wrote a book about it, and in that book he commented that he expended more energy on this one game than on any other in his career, including world championship games.

Although the World Team had input from some strong players, none was as strong as Kasparov himself, and the average quality of player was far below Kasparov. Yet collectively the World Team played a game far stronger than any of the individuals contributing would ordinarily have played—indeed, one of the strongest games in the history of chess. Not only did they play Kasparov at his best, but much of their deliberation about strategy and tactics was carried out in public, an advantage Kasparov used extensively. Imagine not

only playing Garry Kasparov, but also having to explain to him the thinking behind your moves!

How was this possible? How could thousands of chess players, most of them amateurs, compete in a chess game with Kasparov at his peak? The World Team contained people at all levels of chess ability, from beginners to grandmasters. Moves regarded by experts as obviously mistaken sometimes obtained up to 10 percent of the vote, suggesting that many beginners were participating. On one move, 2.4 percent of the votes were cast for moves that weren't merely bad, but actually violated the rules of chess!

The World Team coordinated their play in several ways. Microsoft set up a game forum where people could discuss the game, and also appointed four official advisors to the World Team. These were outstanding teenage chess players, among the best of their age in the world, although none were in Kasparov's class. On each move, the advisors published their recommendations on the Microsoft website, and, if they wanted, a commentary explaining the recommendation. This was done well before World Team voting closed, so the recommendations could influence the voting. As the game progressed several other strong chess players also offered advice. Particularly influential, although not always heeded, was the GM School, a Russian chess club containing several grandmasters.

Most of the advisors and other strong players ignored the discussion on the game forum, making no attempt to engage with the bulk of people making up the World Team, and so distancing themselves from the people whose votes decided the World Team's moves. But one of the advisors did actively engage with the World Team. This was an extraordinary young chess player named Irina Krush. Fifteen years old, Krush had recently become the US Women's chess champion. Although not as highly rated as two of the other World Team advisors, Krush was certainly in the international elite of junior chess players.

Unlike her expert peers, Krush devoted considerable time and attention to the World Team's game forum. Shrugging off abuse and insults, she extracted many of the best ideas and analyses from the forum, wrote extensive commentary describing the thinking behind her recommended moves, and gradually built up a network of strong

brain's origin and hardware are completely different from those of the internet, and there's no compelling reason to suppose the brain is an accurate model of how collective intelligence works, or of how it can best be amplified. Whatever our collective brain is doing, it seems likely to work according to very different principles than the brain inside our heads. Furthermore, we don't yet have a good understanding of how the human brain works, so the metaphor is in any case of limited use at best. If we're going to understand how to amplify collective intelligence, we need to look beyond the metaphor of the collective brain.

Many books and magazine articles have been written about collective intelligence. Perhaps the best-known example of this work is James Surowiecki's 2004 book *The Wisdom of Crowds*, which explains how large groups of people can sometimes perform surprisingly well at problem solving. Surowiecki opens his book with a striking story about the scientist Francis Galton. In 1906, Galton was attending an English country fair, and among the fair's attractions was a weight-judging contest, where people competed to guess the weight of an ox. Galton expected that most of the competitors would be far off in their estimates, and was surprised to learn that the average of all the competitors' guesses (1,197 pounds) was just one pound short of the correct weight of 1,198 pounds. In other words, collectively, if one averaged the guesses, the crowd at the fair guessed the weight almost perfectly. Surowiecki's book goes on to discuss many other ways we can combine our collective wisdom to surprisingly good effect.

This book goes beyond *The Wisdom of Crowds* and similar works in two ways. First, our goal is to understand how online tools can actively *amplify* collective intelligence. That is, we're not just interested in collective intelligence, per se, but in how to design tools that dramatically increase collective intelligence. Second, we're not just discussing everyday problems like estimating the weight of an ox. Instead, our focus is on problems at the limit of human problem-solving ability, problems like competing with Garry Kasparov at the peak of his chess-playing might, or smashing mathematical problems that challenge the world's best mathematicians. Our main interest will be in scientific problem-solving, and of course it's problems at the limit of human problem-solving ability that scientists

most dearly want to solve, and whose solution will bring the greatest benefit.

Superficially, the idea that online tools can make us collectively smarter contradicts the idea, currently fashionable in some circles, that the internet is reducing our intelligence. For example, in 2010 the author Nicholas Carr published a book entitled *The Shallows: What the Internet Is Doing to Our Brains*, arguing that the internet is reducing our ability to concentrate and contemplate. Carr's book and other similar works make many good points, and have been widely discussed. But new technologies seldom have just a single impact, and there's no contradiction in believing that online tools can both enhance and reduce intelligence. You can use a hammer to build a house; you can also use it to break your thumb. Complex technologies, especially, often require considerable skill to use well. Automobiles are amazing tools, but we all know how learner drivers can terrorize the road. Looking at the internet and concluding that the main impact is to make us stupid is like looking at the automobile and concluding that it's a tool for learner drivers to wipe out terrified pedestrians. Online, we're all still learner drivers, and it's not surprising that online tools are sometimes used poorly, amplifying our individual and collective stupidity. But as we've already seen, there are also examples showing that online tools can be used to increase our collective intelligence. Our concern will therefore be with understanding how those tools can be used to make us collectively smarter, and what that change will mean.

We're still in the early days of understanding how to amplify collective intelligence. It's telling that many of the best tools we have—tools such as blogs, wikis, and online forums—weren't invented by the people we might suppose are the experts on group behavior and intelligence, experts from fields such as group psychology, sociology, and economics. Instead, they were invented by amateurs, people such as Matt Mullenweg, who was a 19-year-old student when he created Wordpress, one of the most popular types of blogging software, and Linus Torvalds, who was a 21-year-old student when he created the open source Linux operating system. That tells us we should be wary of current theory: while we can learn a great deal from existing academic studies, the picture of collective intelligence that emerges is also incomplete. For this reason, we'll ground our

discussion in concrete examples in the mold of the Polymath Project and Kasparov versus the World. In part 1 of this book we'll use these concrete examples to distill a set of principles that explain how online tools can amplify collective intelligence.

I have deliberately focused the discussion in part 1 on a relatively small number of examples, with the idea being that as we develop a conceptual framework for understanding collective intelligence, we'll revisit each of these examples several times, and come to understand them more deeply. Furthermore, the examples come not just from science, but also from areas such as chess and computer programming. The reason is that some of the most striking examples of amplifying collective intelligence—examples such as Kasparov versus the World—come from outside science, and we can learn a great deal by studying them.

As our understanding deepens, we'll see that scientific problems are especially well suited for attack by collective intelligence, and in part 2 we'll narrow our focus to how collective intelligence is changing science.

## CHAPTER 3

### Restructuring Expert Attention

In 2003, a young woman named Nita Umashankar, from Tucson, Arizona, went to live for a year in India, where she worked with a not-for-profit organization to help sex workers escape the sex trade. What she found in India frustrated her. Many of the sex workers had so few skills that it was almost impossible to help them find jobs outside prostitution. Returning to the United States, Umashankar decided she would start a not-for-profit organization that addressed the core problem, by training at-risk Indian girls in technology, and then helping them find jobs with technology companies.

Eight years later, the nonprofit she founded, ASSET India, has opened technology training centers in five Indian cities. They've helped hundreds of young people escape the sex trade, and have plans to expand. Unfortunately, many of the smaller towns they'd like to expand into don't have the reliable electricity needed to power crucial technologies such as the wireless routers used to access the internet. ASSET has experimented with using solar-powered wireless routers, but found that the devices already on the market won't run reliably over the long hours their training centers are open.

To solve their problem with wireless routers, ASSET tried something unconventional, searching for help using an online marketplace for scientific problems called InnoCentive. InnoCentive is like eBay or Craigslist, but aimed at scientific problems. The idea is that participating organizations can post online "Challenges"—scientific problems they want solved—with prizes for solution, often tens of thousands of dollars. Anyone in the world can download a detailed description of a Challenge, try to solve the problem, and win the prize.

Using \$20,000 in prize money put up by the Rockefeller Foundation, ASSET posted an InnoCentive Challenge to design a reliable solar-powered wireless router, using low-cost, easily available hardware and software. In the two months the Challenge was posted at InnoCentive it was downloaded 400 times, and 27 solutions were submitted. The \$20,000 prize was awarded to a 31-year-old Texan software engineer named Zacary Brown, and a prototype is being built by engineering students at the University of Arizona.

Zacary Brown wasn't just any software engineer. An enthusiastic amateur wireless radio operator, he was working toward a goal of making radio contact with every country in the world. While growing up, he was enchanted by his parents' explanation of how the solar panels Jimmy Carter installed at the White House made electricity from sunlight, and as an adult he was experimenting with using solar panels to power his wireless radio equipment. Over the long run, he hoped to power his entire home office using solar power. In short, Zacary Brown was exactly the right person for ASSET to be talking to. InnoCentive simply provided a way of making the connection.

Underlying InnoCentive is the premise that there is enormous untapped potential for scientific discovery in the world, potential that can be released by connecting the right people. This premise has been confirmed, with more than 160,000 people from 175 countries signing up to InnoCentive, and prizes for more than 200 Challenges awarded. The Challenges range across many areas of science and technology. Examples include finding more cost-effective methods of manufacturing drugs for tuberculosis, designing a solar-powered mosquito repellent (I'm not making this up!) to combat malaria, and finding better ways of identifying people at risk of developing motor neuron disease. Many of the successful solvers report, as Zacary Brown did, that the Challenges they solve closely match their skills and interests. Furthermore, as in the ASSET story, connections are usually made between parties who otherwise would only have met accidentally. InnoCentive makes such connections systematically, not as lucky one-offs, but at scale.

The reason the connections made by InnoCentive are so valuable is, of course, the big gap between the skills of the people posing the Challenges and those solving the Challenges. While designing

The game was widely publicized within the chess community, and hundreds of experienced chess players were following the game. Chess is so rich with possible variations that many of those players had their own individual areas of microexpertise where they too equaled or even surpassed Kasparov. The key to the World Team's play was to ensure that all this ordinarily latent microexpertise was uncovered and acted upon in response to the contingencies of the game. So although it was a lucky chance that Krush *in particular* was the person whose microexpertise was decisive at move 10, given the number of experienced chess players involved, it was highly likely that latent microexpertise from those players would come to light at critical points during the game, and so help the World Team match Kasparov.

This is, in fact, exactly what happened. As an example, after the game ended Krush singled out move number 26 as one of her three favorite World Team moves. Move number 26 wasn't Krush's idea, or the idea of one of the established chess experts following the game. Instead, move 26 was proposed by one of the posters on the game forum, using the name Yasha, later revealed to be Yaaqov Vaingarten, a reasonably serious but not elite junior player. This was part of a pattern, as during the game Krush drew extensively on the thinking of many unknown or even anonymous contributors to the game forum, people using pseudonyms such as Agent Scully, Solnushka, and Alekhine via Oujii. At the same time, she also consulted with established chess players such as international masters Ken Regan and Antti Pihlajasalo, and grandmaster Alexander Khalifman, of the GM School. The World Team wasn't lucky at all. Rather, the World Team had such a diverse collection of talent available that each time a problem arose, a member of the team rose to the occasion; someone with just the right microexpertise would leap in to fill the gap.

### Designed Serendipity

We've seen how collaborative projects such as Kasparov versus the World and InnoCentive harness latent microexpertise to overcome challenges that would stymie most members of the collaboration.



In the most successful online collaborations this use of micro-expertise approaches an ideal in which the collaboration *routinely* locates people such as Yasha and Zacary Brown, people with just the right microexpertise for the occasion. In particular, as creative collaboration is scaled up, problems can be exposed to people with a greater and greater range of expertise, greatly increasing the chance that someone will see what seems to most participants like a hard problem and think, “Hey, that’s easy to solve.” Instead of being an occasional fortuitous coincidence, serendipity becomes commonplace. The collaboration achieves a kind of *designed serendipity*, a term I’ve adapted from the author Jon Udell.

To understand the value of such serendipity in creative work, it helps to have a concrete historical example. Let’s take Einstein’s work on his greatest contribution to science, his theory of gravity, often called the general theory of relativity. He worked on and off developing general relativity between 1907 and 1915, often running into great difficulties. By 1912, his work had led him to the astonishing conclusion that our ordinary conception of the geometry of space, in which the angles of a triangle add up to 180 degrees, is only approximately correct, and a new kind of geometry is needed to describe space and time. Now, in case you’re wondering what the geometry of space and time has to do with gravity, you’re in good company: it came as a surprise to Einstein, too. When setting out to understand gravity, Einstein had no idea that he’d end up thinking of it as a geometric problem. Nonetheless, there he stood in 1912 with the idea that gravity was somehow connected to a nonstandard type of geometry. And he was stuck, because such geometric ideas were outside his expertise. He talked his problems over with a long-time mathematician friend, Marcel Grossmann, telling him, “Grossmann, you must help me or else I’ll go crazy!” Fortunately, for Einstein, Grossmann was just the right person to be talking to. He told Einstein that the geometric ideas Einstein needed had already been worked out in full, decades earlier, by the mathematician Bernhard Riemann. Einstein quickly dove into Riemannian geometry, and realized that Grossmann was right. Riemannian geometry became the mathematical language of general relativity.

Serendipitous connections like this are crucial in creative work. In science, especially, every active scientist carries around in their

head a host of unsolved problems. Some of those problems are big (“Figure out how the universe began”), some of them are small (“Where’d that damned minus sign disappear in my calculation?”), but all of them are grist for future progress. If you’re a scientist, it’s mostly up to you to solve those problems by yourself. If you’re lucky, you might have a few supportive colleagues who can help you out. Very occasionally, though, you’ll solve a problem in a completely different way. You’ll be talking with an acquaintance, when one of your problems comes up. You’re chatting away when BANG, all of a sudden you realize that this is exactly the right person to be talking to. Sometimes, they can just outright solve your problem. Or sometimes they give you some crucial insight or idea that provides the momentum needed to vanquish the problem. This kind of fortuitous connection is one of the most exciting and important moments in science. The problem is, such chance connections occur too rarely. The reason designed serendipity is important is because in creative work, most of us—even Einstein!—spend much of our time blocked by problems that would be routine, if only we could find the right expert to help us. As recently as 20 years ago, finding that right expert was likely to be difficult. But, as examples such as InnoCentive and Kasparov versus the World show, we can now design systems that make it routine. Designed serendipity enables us to rapidly and routinely solve many of those previously insoluble problems, and so expands the range of our problem-solving ability.

### Conversational Critical Mass

It’s challenging to convey the experience of designed serendipity. It’s one thing to describe examples, but it’s quite another thing to be part of a collaboration where designed serendipity is actually going on. All of a sudden, you feel as though your mind has grown wings. You’re liberated from much of the burden of niggling problems, problems that would be routine if only you had access to an expert with just the right skills. It’s profoundly enjoyable to instead spend your time concentrating on the problems where you have a special insight and advantage. Designed serendipity is something that must be experienced to be fully understood. But with that said, there

is a simple model that can help explain why designed serendipity is important, and how it can qualitatively change the nature of collaboration. That model is a nuclear chain reaction. By reminding ourselves of what happens during a chain reaction we will gain insight into why designed serendipity is important.

The way a chain reaction works is simple. Imagine you have somehow come into possession of a small piece of uranium—uranium-235, the type of uranium that goes into nuclear bombs. (There are several types of uranium, but they don't all undergo nuclear chain reactions. From now on, when I say "uranium" I mean uranium-235.) Uranium atoms, it turns out, aren't very stable. Every once in a while, the nucleus of a uranium atom will disintegrate, spitting out one or more neutrons. That neutron then flies off through the piece of uranium. Uranium, like all solids, only looks solid to the human eye. In fact, at the atomic level it's mostly empty space, and the neutron can travel a long way before it encounters the nucleus of another uranium atom. In a small piece of uranium—say, half a kilogram (about a pound)—the chances are pretty good that the neutron will never encounter another nucleus, and will instead fly all the way out of the piece of uranium, and just keep going. But if the piece of uranium is just a little bit bigger—say, a kilogram—the chances are a fair bit higher that the neutron will smash into the nucleus of another uranium atom. That nucleus then disintegrates, and, it turns out, releases three more neutrons. Now there are four neutrons whizzing through the uranium—it's four because we need to include in our count the original neutron that started the process, which continues to move, even after smashing into the nucleus. Each of those four neutrons is, in turn, likely to smash into four other nuclei, with the result that 16 neutrons are now on the loose. They are likely to crash into still more nuclei, and things rapidly cascade out of control: after 40 collisions like this, we have a trillion trillion neutrons whizzing around. It's because of this incredibly rapid rate of growth that the process is called a chain reaction. Below a certain mass, called the critical mass, a piece of uranium is simply an inert lump of rock. Atoms inside are occasionally decaying and releasing neutrons, but for each such neutron the average number of so-called daughter neutrons caused by further collisions is less than one, and any possible chain reaction quickly dies out. But with just

a slightly larger piece of uranium, larger than the critical mass, the average number of daughter neutrons is slightly more than one. And if the average number of daughter neutrons is even a tiny bit larger than one then the chain reaction will take off, and cascade out of control. If the average number of daughter neutrons is 1.1, then after just 200 collisions the uranium will have more than 100 million neutrons flying around inside, causing still more collisions. This is why two apparently similar pieces of uranium will behave in completely different ways. One will lie inert, while another just slightly larger piece will explode with the force of thousands of tons of dynamite. A small increase in size can cause a complete qualitative change in behavior.

Something similar goes on in a good creative collaboration. When we attempt to solve a hard creative problem on our own, most of our ideas go nowhere. But in a good creative collaboration, some of our ideas—ideas we couldn't have taken any further on our own—stimulate other people to come up with daughter ideas of their own. Those, in turn, stimulate other people to come up with still more ideas. And so on. Ideally, we achieve a kind of conversational critical mass, where the collaboration becomes self-stimulating, and we get the mutual benefit of serendipitous connection over and over again. It's that transition that is enabled by designed serendipity, and which is why the experience of designed serendipity feels so different from ordinary collaboration. It occurs when collaboration is scaled up, increasing the number and diversity of participants, and so increasing the chance that one idea will stimulate another new idea. In the Polymath Project, for example, Tim Gowers commented that the main thing that sped up the process was that he and other participants often "found [themselves] having thoughts that [they] would not have had without some chance remark of another contributor." In Kasparov versus the World the same thing happened, with an idea from one team member often sparking ideas from others, enabling the World Team to explore many different directions.

Of course, the chain reaction model shouldn't be taken too literally as a model of collaboration. Ideas aren't neutrons, and the goal of collaboration isn't simply to go "critical," producing a rapidly ballooning number of ideas. We need to, at least occasionally, have

each participant's attention where it is best suited—that is, where they have maximal comparative advantage. Ideally, the collaboration will achieve designed serendipity, so that a problem that seems hard to the person posing it finds its way to a person with just the right microexpertise to easily solve it (or stimulate further progress). Conversational critical mass is achieved, and the collaboration becomes self-stimulating, with new ideas constantly being explored. In the next chapter, chapter 4, we'll see many collaborative patterns that can help achieve these ends, including:

- Modularizing the collaboration, that is, figuring out ways to split up the overall task into smaller subtasks that can be attacked independently or nearly independently. This reduces barriers to entry by new people, and thus broadens the range of available expertise. Modularity is often difficult to achieve, requiring a conscious, relentless commitment on the part of participants.
- Encouraging small contributions, again to reduce barriers to entry, and to broaden the range of available expertise.
- Developing a rich and well structured information commons, so people can build on earlier work. The easier it is to find and reuse earlier work, the faster the information commons will grow.

In chapter 5 we'll examine the limits to collective intelligence. We'll find that for collective intelligence to be successful, participants must be committed to a shared body of methods for reasoning, so disagreements between participants can be resolved, and do not cause permanent rifts. Such a shared body of methods is available in fields such as chess, programming, and science, but not always in other fields. For example, artists may be fundamentally divided over basic aesthetic principles. Such divisions will prevent collaboration from scaling up, and so prevent designed serendipity and conversational critical mass.

## How Online Collaboration Goes Beyond Conventional Organizations

Using collective intelligence to solve problems is not new. Historically, groups have used three main ways to solve creative problems: (1) large formal organizations, such as the hundreds or thousands of people who may be involved in creating a movie, say, or a new electronic gadget; (2) the market system; and (3) conversation in small informal groups. In the remainder of this chapter we'll investigate how online tools can take us beyond these three existing ways of doing group problem solving.

To understand how online collaboration goes beyond conventional organizations, consider a movie production. A modern blockbuster movie may employ hundreds or even thousands of people—the 2009 movie *Avatar* employed 2,000 people. But unlike the participants in Kasparov versus the World or the Polymath Project, each employee has their own assigned role in the production. An employee in the movie's art department won't usually give advice to a violin player in the orchestra. Yet that's exactly the kind of decision making that happened in Kasparov versus the World. Recall the critical move number 26 suggested by Yasha. In movie terms, it was as though an unknown stranger had wandered on set, made a crucial suggestion to the director, completely changing the course of the movie, and then wandered off.

Of course, there are such stories in the movies. Actor Mel Gibson got his big break when a friend who was auditioning for the movie *Mad Max* asked to be driven to the audition. Gibson wasn't auditioning, but had gotten into a brawl at a party the night before, and had bruises all over his face. The casting agent decided that was the look the movie needed, and Gibson was invited back, completely changing the movie, and launching him on the path to international stardom.

In the world of movies this is an unusual story. But in Kasparov versus the World this kind of occurrence wasn't a lucky one-off, it was the essence of the way the World Team played. There was no preplanned, static division of labor, as in a conventional organization. Instead, there was a dynamic division of labor, in

which every player on the World Team had the opportunity, at least in principle, to be involved in every move.

Let me make more precise what I mean by a dynamic division of labor. It's a division of labor where all participants in a collaboration can respond to the problems at hand, as they arise. Zacary Brown saw ASSET's problem, and realized he could solve it. Yasha followed along the World Team's progress, and realized he had a special insight at move 26. And all participants in the Polymath Project could follow the rapidly evolving conversation, and jump in whenever they had a special insight. In conventional offline organizations, such flexible responses are usually only possible in small groups, if at all. In larger groups different group members focus on their own preassigned areas of responsibility. Online tools change this, making it possible for large groups to harness each participant's special areas of microexpertise, just-in-time as the need for that expertise arises. That's what I mean by a dynamic division of labor. Ideally, as we saw earlier, this will lead to designed serendipity. But even when that doesn't happen, the dynamic division of labor is still strikingly different from the conventional static division of labor.

None of this is to deny the value of a static division of labor. We've achieved enormous improvements in our ability to manufacture goods by improving the static division of labor—think of Henry Ford's assembly line, or even Adam Smith's hypothetical pin factory. But while such a division of labor is well suited to the manufacture of goods, using a predictable and repetitive process, it's been less useful in solving hard creative problems. The reason is that in creative work it's often the unplanned and unexpected insights and connections that matter the most. In many cases, what makes a creative insight important is precisely the fact that it combines ideas that previously were thought to be unrelated. The more unrelated, the more important the connection—recall the astounding connection Einstein and Grossmann made between gravity and Riemannian geometry. Because of this, the greatest creative work can't be planned as part of a conventional static division of labor. No one could have predicted that Kasparov versus the World would unfold the way it did, and so it wasn't possible to anticipate that Krush's special microexpertise would be needed to cope with the position that occurred at move 10. And it certainly wasn't possible