

The Ignorance of Crowds

by Nicholas G. Carr

05/31/2007

a strategy+business exclusive

The Ignorance of Crowds

The open source model can play an important role in innovation, but know its limitations.

by **Nicholas G. Carr**

Ten years ago, on May 27, 1997, a little-known software programmer from Pennsylvania named Eric Raymond presented a paper at a technology conference in Würzburg, Germany.

Titled “The Cathedral and the Bazaar,” the paper caused an immediate stir, and its renown has only grown in the years since. It is now widely considered one of the seminal documents in the history of the software industry.

Raymond’s subject was the open source software movement, as exemplified by what was then — and still is — its most famous product, the Linux operating system. Open source projects, he pointed out, represented a radically new method of software development. Traditionally, sophisticated programs had always been “built like cathedrals, carefully crafted by individual wizards or small bands of mages working in splendid isolation.” An open source project, in contrast, was the product of a large and informal community of volunteers who in aggregate “seemed to resemble a great babbling bazaar of

differing agendas and approaches.” What was amazing, Raymond wrote, was that “the Linux world not only didn’t fly apart in confusion but seemed to go from strength to strength at a speed barely imaginable to cathedral-builders.”

The bazaar model of “peer production” was unthinkable before the Internet came along. It was only when software programmers around the world gained access to a cheap, high-speed communication network that they could start sharing their code in a speedy and efficient manner. As Raymond observed, it was probably not a coincidence “that the gestation period of Linux coincided with the birth of the World Wide Web, and that Linux left its infancy during the same period in 1993–1994 that saw [an] explosion of mainstream interest in the Internet.” The Net formed the thoroughfare of the bazaar.

Of course, that thoroughfare wasn’t open only to software engineers. It was open to every person and to every company. The Net brought the bazaar, and its peer production model, right up to the doors of every business in the world. It’s hardly a surprise, then, that



Raymond's metaphor soon came to be applied far more broadly than he originally intended. Connected to the global masses through the Internet, companies no longer had to pursue innovation in splendid isolation. They had the option of replacing the traditional, closed cathedral model with the new, open bazaar model. Michael Schrage noted the importance of this phenomenon in the pages of this magazine back in 2000 (See "Open for Business," *s+b*, Fourth Quarter 2000). Open source, he wrote, is "transforming how organizations of all kinds seek to create and manage value. [It] will be central to capturing more profits from innovation." In their recent book *Wikinomics* (Portfolio, 2006), Don Tapscott and Anthony Williams similarly argued that peer production can help businesses "take innovation and wealth creation to new levels."

But even as the corporate world has begun to embrace the idea of the bazaar as a forum for innovation,

programs from scratch. That fact has led some to conclude that peer production is best viewed as a means for refining the old rather than inventing the new; that it's an optimization model more than an invention model.

Now that we've arrived at the 10th anniversary of the first appearance of "The Cathedral and the Bazaar," it seems like an opportune moment to take a closer look at both the benefits and the limitations of peer production as a means of business innovation. What's the bazaar good for, and what isn't it good for?

Limits of Linus's Law

In his paper, Raymond highlighted the greatest advantage of the open source model for software development: It dramatically increases the speed with which problems, or bugs, are uncovered and fixed. When only a relatively small number of programmers work on a complex program, debugging consumes

Nicholas G. Carr

(ncarr@mac.com), a contributing editor to *strategy+business*, is the author of *Does IT Matter? Information Technology and the Corrosion of Competitive Advantage* (Harvard Business School Press, 2004). His next book, *The Big Switch*, is forthcoming from W.W. Norton.

Some have concluded that peer production is a means for refining the old rather than inventing the new.

software programmers have continued to debate the strengths and weaknesses of peer production. The open source model has proven to be an extraordinarily powerful way to refine programs that already exist — Linux, for instance, is an elaboration of the venerable Unix operating system, and the open source Firefox browser builds on Netscape's old Navigator — but it has proven less successful at creating exciting new

huge amounts of time and causes lots of delays — and many bugs still manage to sneak through. When you mobilize hundreds or thousands of people, however, they find and fix bugs much more quickly and thoroughly. It's not all that different from, say, an Easter egg hunt. If you hide 100 eggs and have two children search for them, you'll wait a long time for them to finish, and a lot of the eggs will likely remain

undiscovered. If you put two dozen kids to work, however, they'll locate all the eggs in no time. Raymond condensed this simple idea into an aphorism that would become his paper's most famous line: "Given enough eyeballs, all bugs are shallow." He called it Linus's Law, after Linus Torvalds, Linux's founder and presiding genius.

The power that a crowd of contributors has to solve problems derives not just from its sheer size, although that is important, but from its diversity. It's only because the members of the crowd have, as Raymond put it, "differing agendas and approaches" that they're so effective at finding so many bugs (or so many Easter eggs) so quickly. If the participants shared similar outlooks, they'd all end up looking for the same things in the same places. What an unorganized, fairly random group of people provides is not just a lot of eyeballs but a lot of different ways of seeing. As University of Michigan professor Scott Page writes in his new book, *The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools, and Society* (Princeton University Press, 2007), "When solving problems, diversity may matter as much as, or even more than, individual ability."

Raymond also made another, very important observation. What makes the open source model so well suited to finding and fixing software flaws is that debugging is a task that requires little coordination among workers. Debuggers are able to sift through chunks of code in isolation — whether "splendid" or not — without knowing or caring what their fellow bug finders are doing. "Debugging," as Raymond puts it, "is parallelizable." All the

debuggers have to do is communicate their findings and fixes to some central authority, like Linus Torvalds. The central authority takes care of synthesizing the work of the crowd, choosing the best contributions, melding them together into a coherent product, and then redistributing the work to the crowd for the next go-round.

But in Raymond's observation, we also begin to see some of the limitations of the bazaar. First, peer production works best with routine or narrowly defined tasks that can be pursued simultaneously by a big crowd of people. It is not well suited to a job that requires a lot of coordination among the participants. If members of a large, informal group had to coordinate their efforts closely, their work would quickly bog down in complexity. The crowd's size and diversity would turn from a strength to a weakness, and the speed advantage would be lost. Second, because it requires so many "eyeballs," open source works best

filtered through a central authority who holds supreme power as a synthesizer and decision maker. As the Linux project has grown, Torvalds has gathered a hierarchy of talented software programmers around him to help manage the crowd and its contributions. It's not a stretch to say that the Linux bureaucracy forms a cathedral that coordinates the work of the bazaar and molds it into a unified product.

If Raymond made a mistake in his paper, it was in drawing too sharp a distinction between the cathedral and the bazaar. They're not two different and incompatible approaches to innovation. Their relationship is symbiotic. Without the bazaar, the cathedral model moves too slowly. Without the cathedral, the bazaar model lacks focus and discipline.

The People's Encyclopedia

Outside the realm of software, the best-known example of a work created through peer production is

The open source model — when it works effectively — is not as egalitarian or democratic as it is often made out to be.

when the labor is donated or partially subsidized. If Linus Torvalds had had to compensate all his "eyeballs," he would have gone broke long ago.

Third, and most important, the open source model — when it works effectively — is not as egalitarian or democratic as it is often made out to be. Linux has been successful not just because so many people have been involved, but because the crowd's work has been

Wikipedia, the giant online encyclopedia that is being written by many thousands of volunteers. The creation of Wikipedia, and its predecessor, Nupedia, was directly inspired by the open source movement. In fact, Jimmy Wales, a cofounder of the encyclopedia, has said that Eric Raymond's essay "opened my eyes to the possibility of mass collaboration."

Wikipedia has been a roaring

success by many measures. Its original English-language version contains well over a million articles, covering everything from Aachen to ZZ Top, and it has become one of the most visited sites on the Web, attracting more than 150 million people a month. Its achievement underscores the power of peer production to dramatically accelerate narrowly defined tasks that require little coordination — in this case, finding and paraphrasing descriptions of many different subjects — and reinforces the important role that diversity plays in such efforts. In contributing to the encyclopedia, each Wikipedia volunteer naturally focuses on those subjects that interest him or her, and it's the wide range of those interests that has enabled the encyclopedia to get so big so fast.

But for all its breadth and popularity, Wikipedia is a deeply flawed product. Individual articles are often poorly written and badly organized, and the encyclopedia as a whole is unbalanced, skewed toward popular culture and fads. It's hardly elitist to point out that something's wrong with an encyclopedia when its entry on the Flintstones is twice as long as its entry on Homer. Eric Raymond himself has become one of Wikipedia's harshest critics. "The more you look at what some of the Wikipedia contributors have done, the better [Encyclopaedia] Britannica looks," he told the *New Yorker* in 2006. If Wikipedia weren't free, it is unlikely its readers would be so forgiving of its failings.

The Linux operating system, in contrast, is renowned for its high quality. It routinely runs for months on end without crashing. What explains the difference? Wikipedia's problems seem to stem from the fact

that the encyclopedia lacks the kind of strong central authority that exerts quality control over the work of the Linux crowd. The contributions of Wikipedia's volunteers go directly into the product without passing through any editorial filter. The process is more democratic, but the quality of the product suffers.

Aware of Wikipedia's flaws, Wales and other contributors have been trying hard to improve the

evolve. Nevertheless, its most prominent examples are now mature enough that we can begin to draw from them some practical lessons for executives and entrepreneurs looking to the open source model as a means of strengthening the innovativeness of their organizations.

The bottom line is that peer production has valuable but limited applications. It can be a powerful tool, but it is no panacea. It's a great

The bazaar should be defined by diversity, but the cathedral should be defined by talent.

quality of the site's content. A management team has slowly been taking shape, and it is establishing editorial policies and policing contributions. But even though this nascent hierarchy has already become much more bureaucratic than Linux's lean managerial structure, it hasn't yet been able to substantially improve Wikipedia. The failure appears to stem from the makeup of the supervisory group. Whereas the Linux team is a strict meritocracy, Wikipedia's administrators represent a broader mix of contributors. They're often chosen on the basis of how much they've contributed or how long they've contributed rather than on the quality of their contributions or their editorial skill. It seems fair to say that although the bazaar should be defined by diversity, the cathedral should be defined by talent. When you move from the bazaar to the cathedral, it's best to leave your democratic ideals behind.

Peer Production's Place

Peer production is still a new phenomenon and will continue to

way to find and fix problems, to collect and categorize information, or to perform any other time-consuming task that can be sped up by having lots of people with diverse perspectives working in parallel. It can also have the important added benefit of engaging customers in your innovation process, which not only allows their insights to be harnessed but also may increase their loyalty to your company.

But if peer production is a good way to mine the raw material for innovation, it doesn't seem well suited to shaping that material into a final product. That's a task that is still best done in the closed quarters of a cathedral, where a relatively small and formally organized group of talented professionals can collaborate closely in perfecting the fit and finish of a product. Involving a crowd in this work won't speed it up; it will just bring delays and confusion.

The open source model is also unlikely to produce the original ideas that inspire and guide the greatest innovation efforts. That

remains the realm of the individual. Raymond was clear on that point when, toward the end of his paper, he examined some of the “necessary preconditions” for the bazaar model of production. “It’s fairly clear,” he wrote, “that one cannot code from the ground up in bazaar style. One can test, debug and improve in bazaar style, but it would be very hard to *originate* a project in bazaar mode.” In a recent e-mail to me, he was even blunter. “The individual wizard,” he wrote, “is where successful bazaar projects generally start.”

Matt Asay, a software executive with long experience in the open source movement, agrees. “All open source projects — without exception — are started by one or two people and...have a core development group of fewer than 15 developers,” he says. “The most you can hope for [from the broader set of contributors] is bug fixes.” Asay warns that trying to expand the core decision-making group to include more of the “community” can backfire, as the resulting decision-by-committee approach tends to produce “stale, conservative code.” In other words, keep the bazaar out of the cathedral.

So if you’re looking to bolster your company’s creativity, you should by all means look for opportunities to harness the power of the crowd. Just don’t expect the masses to take the place of the lone wizard or the band of mages. The greatest breakthroughs will always begin, to quote Eric Raymond once more, with “one good idea in one person’s head,” and the greatest products will always reach perfection through the concerted efforts of a highly skilled team. +

strategy+business magazine
is published by Booz & Company Inc.
To subscribe, visit www.strategy-business.com
or call 1-877-829-9108.