

The Principles of Distributed Innovation

“No matter who you are, most of the smartest people work for someone else” is known as Joy’s Law in the high-tech industry. Attributed to Sun Microsystems co-founder Bill Joy, this “law” emphasizes the essential knowledge problem that faces many enterprises today, that is, that in any given sphere of activity most of the pertinent knowledge will reside outside the boundaries of any one organization, and the central challenge for those charged with the innovation mission is to find ways to access that knowledge.

The causal explanation of Joy’s Law is provided in the seminal work of economists Friedrich Hayek and Eric von Hippel on the distributed and sticky nature of knowledge and innovation. Hayek¹, in 1945, arguing for the importance of the market economy, emphasized that at the macro level knowledge is unevenly distributed in society, and that centralized models for economic planning and coordination are prone to failure due to an inability to aggregate this distributed knowledge. Thirty years later, micro-level studies by von Hippel² began to suggest that in many industries users were the originators of most novel innovations. Users’ dominant role in originating innovations reflects the fact that knowledge is not only distributed but also “sticky,” that is, relatively difficult and extremely costly to move between locations, thus shifting the locus of innovation to where it is the stickiest.³ Users generate functionally novel innovations because they experience novel needs well ahead of manufacturers, and manufacturers develop dimension of merit innovations (that improve the performance of existing features) because they specialize in producing products for the mass market.⁴

Joy’s Law is exacerbated by the explosion of knowledge in most scientific and technological fields. In the online database of the US National Library of Medicine (Medline), for example, between 1955 and 2005, the number of academic papers published in the life sciences increased approximately six-and-one-half-fold, from 105,000 to 686,000.⁵ Even in relatively narrow and obscure fields, tissue engineering for instance, 6,131 academic publications were authored by 17,044 individuals between 2004 and 2006.⁶ In the face of this explosion of knowledge, most organizations will have difficulty keeping up with significant trends and identifying and locking up key sources of knowledge for competitive gain. Joy’s Law is thus not so

Karim R. Lakhani is an assistant professor in the Technology and Operations Management Unit at the Harvard Business School.

Jill Panetta consults to organizations implementing distributed innovation strategies in R&D. She is a co-founder and former Chief Science Officer of InnoCentive.

much a statement about the declining IQs of workers or poor hiring practices as an acknowledgement that the traditionally closed models of proprietary innovation will have difficulty completing knowledge intensive tasks when most of the needed knowledge resides outside the organization.

The successful development of the Linux operating system and numerous other open source software (OSS) projects provides an alternative model for organizing for innovation. Many practitioners and scholars of innovation did not anticipate the emergence of a distributed and open model for innovation that can aggressively compete with traditionally closed and proprietary models. That complex software systems running mission critical applications can be designed, developed, maintained, and improved for “free” by a virtual “community” of mostly volunteer computer programmers has come as a great surprise to them. Perhaps even more surprising is that some of the largest software companies and the biggest holders of intellectual property (e.g., IBM, Sun, Apple, and Oracle) have embraced OSS communities by encouraging the participation of their own personnel in, and donating copyrighted software and patents to, these communities, and integrating OSS software into their strategic product and service offerings.

OSS communities are the most fully developed example of the appearance of distributed innovation systems characterized by decentralized problem solving, self-selected participation, self-organizing coordination and collaboration, “free” revealing of knowledge, and hybrid organizational models that blend community with commercial success. The achievements of OSS communities have brought the distributed innovation model to general attention, but it is rapidly taking hold in industries as diverse as apparel and clothing, encyclopedias, biotechnology and pharmaceuticals, and music and entertainment.

In this article, we first provide an overview of distributed innovation systems that are achieving success in three different industries with three different organizational models. We then consider in the context of these three examples questions and concerns related to why people participate, the organizing principles of production, and the implications for intellectual property. We close our discussion with a review of potential extensions and limitations of this alternative model of innovation.

MODELS FOR DISTRIBUTED INNOVATION

The Self-Organizing Community

Linux invariably comes to mind when OSS communities are mentioned. Its organizational and commercial success has stunned most observers. Linux’s growth from just over 10 thousand lines of code at its inception to about four million lines of code as of the latest version reflects the contributions of thousands of individuals.⁷ Just last year, 1, 961 developers added 754,000 lines of code.⁸ The commercial ecosystem that surrounds Linux is expected to reach about \$35 billion in 2008 with installations in more than 43 million computing devices ranging from PCs and servers to cell phones, routers, and super computer clusters.⁹

The Principles of Distributed Innovation

Its beginnings belie the contemporary scope and value of this global movement. Linux's genesis in 1991, in the pre-Web era of the Internet, was a series of announcements and requests for help posted by then 22-year-old Linus Torvalds on a message board for computer operating systems.¹⁰

Hello netlanders,

Do you pine for the nice days of minix-1.1, when men were men and wrote their own device drivers? Are you without a nice project and just dying to cut your teeth on an OS you can try to modify for your needs?... :-)

I'm doing a (free) operating system, just a hobby, won't be big and professional

I'd like any feedback on things people like/dislike... This is a program for hackers by a hacker. I've enjoyed doing it, and somebody might enjoy looking at it and even modifying it for their own needs.... Drop me a line if you are willing to let me use your code.

Linus (torvalds@kruuna.helsinki.fi)

These announcements set in motion a loose, informal collaboration that led to the establishment of a framework for interaction among the global community of software developers that created the Linux kernel (the core of a computer operating system). As Torvalds' messages make clear, the initial drivers of participation in Linux were user need and fun. The promise of Linux was of a powerful Unix-like operating system, previously available only on high-end computers that could run on commodity Intel hardware. Because it was available on the Internet, users could download the source code to their own computers and modify it to suit their needs and interests. Modifications made to the source code were then sent back to Torvalds in the hope that they would be included in the next release of the kernel. The growing community established its own Internet-based discussion forums and began to work collaboratively to resolve technical issues related to Linux development.

Although over the past 16 years the number of people and firms interested in Linux has continued to grow, the basic model of participation on the basis of user need or curiosity and having fun has not changed. To participate one need only sign up for the Linux kernel mailing list (LKML) and be competent to modify source code. LKML is the rendezvous point for technical discussions of the features being developed by contributors. LKML participants report and fix bugs, contribute and modify code, and discuss the technical evolution of the kernel. Although Torvalds has final say as to what goes into the kernel, much of the development is organic, determined by the actions of community members and not by any measure of explicit project management within the community. Torvalds' philosophy regarding management of the community runs counter to most expecta-

tions of how complex technological projects should be run. A recent debate within LKML highlights the issues.¹¹

Rik van Riel (an active contributor to Linux):

“It seems like Linux really isn’t going anywhere in particular and seems to make progress through sheer luck”

Linus Torvalds (in several emails in a longer thread):

“Hey, that’s not a bug, that’s a FEATURE

Well, sheer luck, AND:

- free availability and `_crosspollination_` through sharing of ‘source code,’ although biologists call it DNA.

- a rather unforgiving user environment, that happily replaces bad versions of us with better working versions and thus culls the herd (biologists often call this ‘survival of the fittest’)

- massive undirected parallel development (‘trial and error’)

Do I direct some stuff? Yes. But, quite frankly, so do many others. Alan, Al, David, even you. And a lot of companies are part of the evolution whether they realize it or not. And all the users end up being part of the ‘fitness testing’...

“A strong vision and a sure hand sound good on paper. It’s just that I have never met a technical person (including me) whom I would trust to know what is really the right thing to do in the long run....

“Too strong a strong vision can kill you— you’ll walk right over the edge firm in the knowledge of the path in front of you...

“I’d much rather have ‘brownian motion,’ where a lot of microscopic directed improvements end up pushing the system slowly in a direction that none of the individual developers really had the vision to see on their own.

And I’m a firm believer that in order for this to work `_well_`, you have to have a development group that is fairly strange and random.”

Most managers would balk at developing complex technological artifacts with the help of “fairly strange and random” individuals. But Torvalds and Linux show the potential benefit of organizing work such that many individuals can self-select and lead elements of development without much ex-ante guidance and control.

Blending Community and Commerce

Threadless.com, an online t-shirt company, foreshadows the commercial enterprise of the future that is built to leverage community-based distributed innova-

tion. Firms in the apparel and fashion business face two critical challenges, (1) to attract the right designer talent at the right time to create recurring fashion hits, and (2) to forecast sales so as to be better able to match production cycles with demand cycles. Threadless solves these problems by letting its international community of customers take over such core functions as innovation, new product development, sales forecasting, and marketing.

Threadless was started in 2000 by friends Jake Nickell and Jacob DeHart who were active participants in a Chicago-based online community of designers called Dreamless. The experience of winning a t-shirt design competition sponsored by the Dreamless community exposed Nickell, and by association his friend DeHart, to the idea that co-creation with a community was a relatively untapped market. Both were amazed by the variety and high quality of submissions received by the Dreamless community organizers. As budding designers, they realized early on that the fashion industry was fickle and they had no monopoly on good design ideas. But a platform that solicited design ideas from anyone and provided for community-based selection of submitted designs might overcome their own limitations. Hence Threadless was formed.

Threadless.com's business model revolves around an ongoing competition to which anyone, professional graphic designers and amateurs alike, can submit designs for new t-shirts. The community is polled on both the designs (which are rated using a scale of zero to five) and willingness to buy. Threadless uses this information to select for production six to ten new designs each week. Winning designs' creators receive cash and prizes worth \$2,500, are recognized for their accomplishment on the company's Web site, and have their screen name printed on the t-shirt label. Community members also critique submitted designs and provide feedback to help designers improve their ideas going forward. Threadless also populates its online catalog with photographs of community members wearing t-shirts bearing winning designs.

Threadless has become both a commercial and community success story. In 2006, it sold more than 1.5 million t-shirts to customers around the world, and its active community exceeds 600,000 members. Threadless receives more than 800 new design submissions per week, each of which is typically rated and assigned a demand signal indicating intent to purchase by more than 500 community members. More than half of purchasers have, at one time or another, voted on t-shirt designs. The online community is vibrant, logging in excess of 1,000 new member posts per day discussing design and art and submitting music and video inspired by designs. All of this has been accomplished with no reliance on traditional forms of advertising and customer recruitment. Such has been its success that the company has on many occasions declined overtures by large-scale retailers to sell Threadless t-shirts in stores around the world.

Getting Outsiders to Innovate

InnoCentive.com is changing how the toughest science-based R&D problems get

solved by traditionally closed enterprises in the pharmaceutical, biotechnology, consumer goods, and high-technology industries. The business offers firms that encounter difficult science problems an alternative to devoting laboratory time and resources to the search for a solution. Firms can now post such problems, together with a designated cash prize (typically ranging from \$5,000 to \$100,000) for an acceptable solution, on InnoCentive.com. Problem posters and prospective solvers, who self select to the attempt to devise or formulate a solution, remain anonymous to one another throughout the process. InnoCentive's role is that of knowledge broker, providing the seeker firms that post problems with solutions solvers have been motivated to submit. The seeker firm chooses the most appropriate solution, if any, and receives from the solver, in return for the prize money, all rights to the intellectual property related to the solution.

InnoCentive was spun off from Eli Lilly and Company's Internet incubator in 2000. The driver of InnoCentive was then-VP of R&D Alph Bingham, who recalled from his doctoral student days that most science problems were potentially amenable to multiple approaches and diverse solutions, and that often top students in one problem area were not necessarily at the top in another. Yet, within firms, he realized, science problems were typically assigned to a single scientist or a small team possibly either not at the top in the problem domain or unaware of alternative approaches. Having seen his share of projects and problems being tackled internally that had gotten stuck or lain fallow because solutions were not forthcoming, Bingham reasoned that a dedicated company that could connect diverse "outside" scientists with "inside" problems could be the answer to some of the scientific productivity challenges in pharmaceuticals and other industries.

InnoCentive's solver network includes more than 120,000 scientists from around the world. More than 400 problems that could not be solved by the R & D laboratories of some 50 firms have been posted. Each problem piques the interest of more than 200 scientists, about ten of whom submit solutions. About one-third of the problems posted by seeker firms have been solved and the associated prizes awarded. Solutions arrive from unexpected sources and are typically not what the originating problem holder scientists had envisioned as possible.¹² For example, one firm had experienced difficulty transferring a chemical powder to a specialty container. Whereas most of the unsuccessful solutions proposed within the firm had attempted to enable transport by modifying the material's chemical properties, the winning solution submitted by an InnoCentive solver employed instead the use of electrostatic charge, in essence, applying a physics solution to a chemistry problem. In another case, an aerospace physicist, a small agribusiness owner, a transdermal drug delivery specialist, and an industrial scientist all submitted diverse winning solutions to the same scientific problem: identification of a polymer delivery system.

MOTIVATIONS TO PARTICIPATE

“Why do people work and participate for ‘free?’” is one of the first questions asked when distributed innovation systems are encountered. The emergence of OSS communities and sheer numbers of participants, in the hundreds of thousands, first raised the question of motivation in distributed innovation. Certainly, the common view of “homo economicus,” of purely self-interested participants, is not the answer when many participate with no promise of a direct financial reward for their efforts. The answer lies rather in a more expansive view that acknowledges, as well as the role of economic motivations, notions of enjoyment and having fun together with identity and the social benefits of community.

Research on OSS communities has shown motivation to participate to break out broadly into extrinsic, that is, direct or indirect rewards for performing a task, and intrinsic, that is, valuing a task for its own sake.¹³ In the context of OSS, participation is driven primarily by user need. Programmers observe that they contribute time, effort, and intellect because they have a direct need for a particular software functionality that is not available from commercial sources. That individuals participate because they can use the software and features to which they have directly contributed for work or non-work purposes is consistent with the fact that approximately 40 percent of the participants in OSS communities are paid to participate. Employees are encouraged by their employers to contribute code to OSS communities because software that addresses needs of the organization might result through community development.

In the case of output that is not needed by the contributor, a cash reward might be tendered for substantial contributions. Such efforts are not undertaken with any ex-ante guarantee that they will be rewarded. Rather, payment is at the discretion of the sponsoring organization, made after the work has been completed and evaluated as meeting a certain criteria, and is usually attended by the formal transfer of intellectual property between contributor and sponsor.

Beyond pecuniary benefits, extrinsic reasons for participation include job market signaling and skill and reputation building. Distributed innovation communities provide a relatively open and transparent platform for exhibiting skills and talents to prospective employers. Participants don't need high-level credentials to directly demonstrate their abilities in highly specialized domains, and employers can screen and hire talent by directly observing or soliciting third-party verification of skills. Peer review benefits both members and products. The experience of the community can be leveraged both to improve the quality of contributions and to provide skill-enhancing feedback to contributors. Participants' open activities also accrue reputation among peers. Consistently contributing top-notch code and helping to bring along other members of the community earn status that often translates into privileges within and reputation outside the community, with the attendant possibility of future rewards.

If writing code, designing graphic elements, and solving tough science problems are construed by outsiders to be unremunerated blood, sweat, and tears, the

contributors themselves are more likely to insist that the work is a source of significant satisfaction that derives from the pure joy of engagement in the work, or with the group or community, or both. Research has shown that members of OSS communities quite simply enjoy the programming task sufficiently to want to devote their incremental free time to it.

Programming has been observed to put some in a “flow state” whereby enjoyment of the task is maximized and intense and focused concentration is achieved.¹⁴ A flow state results when a person’s skill matches the challenge of a task. A task beyond an individual’s skill provokes anxiety; a task beneath an individual’s skill induces boredom. Achieving a flow state also correlates with a personal sense of creativity. A majority of respondents in OSS communities report their involvement in various software projects to be among the most creative work they’ve undertaken in their lives. Challenge, enjoyment, and creativity are hallmarks of “voluntary” participation in distributed innovation. Similar findings regarding the importance of the intellectual challenge and enjoyment of the task at hand have been reported on the part of InnoCentive solvers and Threadless designers.

A strong sense of identity and community belonging also motivates participation. Individuals who strive to be active players in the community are inclined to act in a manner consistent with its norms. Contributors are socialized by their participation into acting in a way that advances the collective. For example, because the norms of free revealing and code sharing are universally accepted and expected in OSS communities, many participate openly. Members who have benefited personally from using source code developed by many other members, moreover, feel obliged to give back to their community. Finally, OSS communities can be a source of a sense of self-identity that can lead members to undertake tasks that benefit the community generally. This sense of belonging, which has been observed as well in communities formed by firms, has been found to be quite powerful.

The intrinsic and extrinsic motivations to participate in distributed innovation systems are not intuitively obvious to new observers of the phenomenon. Most, in fact, find to be counterintuitive the association of fun, enjoyment, and a personal sense of identity with the accomplishment of complex technical tasks. But the research findings strongly suggest that the functioning of these systems is driven by mixed and heterogeneous motivations. Consequently, optimizing on only one dimension might have the effect of limiting participation.

ORGANIZING PRINCIPLES

“Brownian motion-based management” is not yet taught in any business schools. But the participation of commercial enterprises in OSS communities and other distributed innovation systems suggest that organizing principles for participation, collaboration, and self-organization can be distilled. Importantly, these systems are not “managed” in the traditional sense of the word, that is, “smart” managers are not recruiting staff, offering incentives for hard work, dividing tasks, inte-

grating activities, and developing career paths. Rather, the locus of control and management lies with the individual participants who decide themselves the terms of interaction with each other.

Key to participation is contributors' self-selection to tasks. In the case of OSS, contributors self-select to tasks that will generate functionality that they need or eliminate a bug that is hindering their use of functionality that is otherwise available to them. Self-selection to tasks can also be influenced by what other community members are indicating to be potential issues and opportunities. But whatever the driver, the matching of individuals' skills and tasks at hand is entirely at the discretion of the individual contributors. Rarely do community "managers" allocate tasks or attempt to perform this matching. InnoCentive contributors similarly self-select to science problems for which they perceive a match between their knowledge base and abilities and the requirements of the solution, and the designs submitted to Threadless reflect contributors' individual interests, inspirations, and graphic design skills.

Furthermore, Carliss Baldwin and Kim Clark have shown how "The architecture of participation" in distributed innovation systems is driven by the granularity and diversity of the tasks available in a given context.¹⁵ The more granular and diverse the available tasks, the larger the potential pool of participants. Participation in Threadless, for example, is not limited to individuals with bona fide graphic design skills. Non-designers can provide feedback, suggest changes in color and graphical elements, for example, and also indicate their preferences for and willingness to purchase particular designs, contributions that are as essential to the success of the business as the submission of designs. There is also an important role in the grass roots marketing of the business and community that is played by those who contribute digital photographs and videos of themselves wearing Threadless t-shirts, and post to the site's lively blogs.

The task granularity and diversity observed in OSS communities is reflected in the range of opportunities open to contributors, who not only might update existing as well as write new code, but also report or fix bugs, request new features, engage in discussions of approaches to coding, write documentation, create and improve graphical user interfaces, translate interfaces into different languages, and provide user-to-user technical support. Tasks generally remain latent until they pique the interest of a contributor with the requisite skills. Core members of OSS communities, often possessing use experience in narrow domains, rely on other participants to help identify and then fulfill the missing elements.

Even in the case of the problems posted on InnoCentive, the solutions to most of which might be expected to be formulated by individuals with advanced scientific training, granularity is important. The problems posted are not of infinite scope, "finding a cure for cancer or discovering anti-gravity," for example, but rather are sufficiently decomposed to accommodate attempts by individual scientists to solve them using locally available materials, methods, and tools. Granularity also helps to assure that solvers will not devote inordinate amounts of time to attempts to devise solutions that ultimately prove unsuccessful.

Task granularity and diversity in distributed innovation systems are further enhanced by the information processing nature of most tasks and widespread availability of low cost tools that support innovation.¹⁶ OSS communities emerged as an ideal type of distributed innovation system because the functionality required for software invention, innovation, production, and distribution are digital and information-based. The tools required for software production—text editors, compilers, debuggers, and source code repositories—are widely available and cheap (in most instances, even free). The high degree of fidelity with respect to error detection and correction provided by these tools enables contributors to share and evaluate each others' as well as newcomers' contributions.¹⁷ The efficacy of assertions about software designs and operations is quickly proved by the requirement that they be converted into code that is then run on contributors' computers. Or not. If code doesn't run, or generates errors, the problems can be readily identified and the code either rejected or repaired by the contributor or other interested members of the community. The advent of e-mail, by making all members with Internet access universally available and connectable, reduced the cost of coordination and collaboration within OSS communities. Going forward, any individual with sufficient background, training, or experience in computer science could participate at little to no cost to themselves or to the community.

Threadless leverages the ubiquitousness of computers, graphic design software, and the Internet to transform a material good (i.e., a t-shirt) into an information good. Submissions created by contributors on their own computers using their preferred graphics packages are uploaded to Threadless's Web site, critiqued and evaluated, and possibly re-worked. Community participation takes the form of representing ideas, inspiration, and tastes as information signals that can be easily aggregated and evaluated. Threadless and the members of its community effectively share the material costs of near universal, virtual participation. Threadless developed and hosts the Web site that provides the information infrastructure and platform that support and facilitate participation, and community members use their own computers and design software to participate and interact with one another.

In the case of InnoCentive, the tools required for participation vary with the type of problem posted. Problems for which a "paper solution," that is, a research proposal solution, is required are essentially information problems for which recourse to local knowledge stores and scientific journals is often sufficient. Problems that require a "reduction to practice" solution, that is, submission of a chemical or biological agent, are likely to attract solvers who have easy and ready access to the necessary materials and equipment. One winning solver, for example, was a retiree with a fully equipped, home-based organic chemistry laboratory that he used for his hobbies.

More generally, two trends are making the tools needed for scientific problem solving more accessible. First, much scientific work can now be done in-silico, that is, the information component of material objects is extracted and modeled and further developed on computers. Computer simulations that provide good first

approximations of the viability of proposed solutions, by shrinking the solution space, reduce or eliminate the costly trial-and-error phase of bench-based problem-solving. Second, the cost of physical tools needed to generate solutions is also declining rapidly and dramatically. For example, polymerase chain reaction machines used for genomic amplification and sequencing can now be purchased on e-Bay for less than \$1,000. InnoCentive's distributed problem solving leverages the sunk costs across the tools owned by solvers.

Distributed innovation systems are organized so as to lower the cost of participation for contributors. Reducing or eliminating barriers to entry expands the population that can self-select into the community. There typically being no screening, joining a community tends to be easy, even trivial. Neither is there any *a priori* guarantee of acceptance. With task granularity, too, the degree to which an innovation outcome depends on the contribution of any given individual is reduced. Granules of activity can be parceled out to individuals working independently, and then aggregated by those same individuals, working collaboratively, into a larger outcome. Co-creation not only limits the cost to individuals, but also, because a broader base of knowledge and perspectives is brought to bear in the creation process, tends to produce more robust innovations.

OPENNESS AND INTELLECTUAL PROPERTY

Inviting "outsiders" to participate in the innovation process naturally implies a different orientation towards issues of openness and intellectual property (IP). The traditional road to innovation, which is paved in secrecy and walled off to competitors and customers and emphasizes the accumulation of a large IP portfolio, would clearly not be effective in a distributed innovation setting in which participation is invited from many individuals. In such settings, individuals and especially organizations must be comfortable with the requisite degree of openness in the innovation process and adjust IP policies to encourage greater sharing and reuse of knowledge and expertise. There is no standard approach to evolving towards open innovation. The degree to which openness and creativity can be accommodated by IP is dependent on the context of operation and norms of the community industrial setting and business model.

OSS communities represent the most radical edge of openness and sharing observed to date in complex technology development. OSS communities are open in the sense that their outputs can be used by anyone (within the limits of the license), and anyone can join by subscribing to the development e-mail list. Openness in joining leads, in turn, to transparency in the development process, since the bulk of communication about projects and their direction generally occurs in public. This means that project leadership is accountable to the wider community for growth and future direction, and also that everyone will be aware of shortfalls and issues. Transparency also affords individuals self-determination with respect to the level of effort they choose to commit, and awareness of others' efforts that they might be able to fold into their own contributions.

Traditional means of IP protection (e.g., copyrights and patents) are not available within the context of OSS projects. Submitted code, although nominally copyrighted to the contributors or their employers, is, according to the terms of most OSS licensing arrangements, immediately available for use and further modification by others. This ethos of sharing and modification extends beyond code to the actual technology development process in the sense that community members engage in joint problem solving via open sharing of ideas and ongoing interaction.

OSS communities license code to ensure that all contributed software is available to all users both within and outside the community. Formal licensing arrangements vary greatly by project, but mostly follow the example of the Free Software Foundation in using the General Public License (GPL), which stipulates that all modifications to the source code must be made public if the modified code has been redistributed. Linux, for example, is licensed by Torvalds under the GPL, and all modifications to it by commercial entities, such as TiVo's modification to the Linux kernel for its consumer DVR, are public. Projects that do not impose this stipulation often use a Berkeley Software Distribution (BSD)-like license, which permits redistribution without release of modifications. Communities that use a BSD license, the Apache Software Foundation for example, rely on the speed of the development process and frequency of updates to ensure that all modifications come back to the community, as it would be very costly to keep private modification in sync with public relations.

Designers for Threadless need to feel comfortable revealing their designs to community members and accepting feedback, negative as well as positive. There is also the risk of submitted ideas inspiring others to create related but different designs that might outperform the designs that inspired them. Most designers will also need to be psychologically prepared to lose in public, as very few designs are converted into t-shirts. Finally, whereas non-winning designers retain all rights to their work, winning designers, in exchange for the cash prize, assign the copyright for their work, and exclusive use of the design on t-shirts, to Threadless. That Threadless management has embraced openness and transparency is reflected in most decisions about the interaction platform, voting and selection rules, and even manufacturing strategy being taken openly and in consultation with the community, which has on many occasions provided direction and guidance unanticipated by Threadless management. Recently, for example, issues of copyright infringement over a winning design were identified and rectified through direct involvement of the community. Threadless retains sufficient rights to IP to ensure the viability of its business model, but leaves rights to all other uses to the designers. The company's claims on IP are thus not based on the use of its platform.

The sine qua non of InnoCentive is the willing broadcast by seeker firms of their R&D laboratories' current difficult to solve in-house science problems, traditionally regarded as highly proprietary and often jealously guarded. InnoCentive works with firms to generalize their problems such that no company specific information is revealed. That seekers and solvers remain anonymous to one another throughout the process further mitigates the risk of releasing proprietary knowl-

edge about internal scientific programs. A firm that finds a suitable solution to its broadcasted problem acquires the IP from the solver in exchange for the agreed prize. Most IP transfer clauses grant the seeker rights to internal use and the solver rights to use in applications not required by the seeker. Through contractual arrangements that provide for R&D unit laboratory output audits, InnoCentive ensures that solutions viewed but not acquired by seeker firms do not somehow show up in the firm's IP portfolio, thereby protecting non-winning solvers.

Currently, due to seeker firms' concerns about "clean" IP transfer and the allocation of prize money, InnoCentive solvers cannot work together or discover through other scientists on the platform complementary approaches to solving a particular problem. This naturally limits the innovative capability of the distributed network. As has been observed in OSS settings, the free flow of ideas facilitates creative solutions to often intractable problems. The limitations are not necessarily within the InnoCentive web platform; rather, it is that seeker firm's IP lawyers and managers need to become comfortable with acquiring IP created by multiple and disparate collaborators. In general, distributed innovation systems thrive when organizers embrace openness, transparency, and IP regimes that sustain continued collaborative participation.

LIMITATIONS AND EXTENSIONS

Although we have presented a fairly positive view of the success and possibilities of distributed innovation systems, they are not without limitations. First, we observe a relatively high failure rate at various levels of analysis. The Linux operating system, Apache web server, and Firefox web browser are spectacularly successful and thriving, but many OSS projects are stillborn. In Sourceforge.net, an online repository of OSS projects, there are in excess of 100,000 software projects engaging more than 1.4 million users, but few are making any meaningful progress towards code shipment and the cultivation of active development communities.¹⁸ At the micro-level, data from Threadless and InnoCentive reveal that most attempts to create designs and solve science problems fail. The six to ten new t-shirt designs Threadless releases each week are selected from a base of approximately 800 new design submitted weekly. Similarly, InnoCentive's problem resolution rate of 30% is on the base of about one successful solution out of ten submitted for each problem.

Second, there is a non-deterministic element of distributed innovation systems in that they cannot be used in the manner of traditional R&D organizations to deliver innovations "on demand" and according to annual plan. Anyone who expects strict planning guidelines and milestone-based innovation development in OSS communities will be disappointed. Contributors to these communities are not employees and they cannot be expected, nor do they care, to be in tune with the pressures and methods of sponsoring organizations. One who desires that a particular feature be added to OSS software is well advised to develop it oneself.

Finally, there is within traditional organizations a great deal of internal resist-

ance to embracing distributed innovation systems. Many organizations are quite good at absorbing external knowledge for internal consumption, but many fewer are comfortable being transparent about internal issues and problems that need to be resolved. Concerns about trade secrecy and IP protection are immediately raised when the distributed innovation alternative is considered. Staff often believe that revealing knowledge about internal developments will put the organization at risk and might tip off competitors about future plans. Many insiders also believe that they have a monopoly on relevant knowledge and are already in contact with knowledgeable external experts, rendering a waste of time interaction with random individuals outside the organization. A stronger but generally unstated reservation is fear of loss of employment. Some internal staff view the creation of a distributed innovation system as a first step towards outsourcing their jobs. Internal staff who are cynical about management's motives in embracing distributed innovation often resist cooperating with such efforts.

These limitations notwithstanding, many individuals and organizations have been inspired by the success of distributed innovation systems and are applying the principles to other domains. Perhaps the most successful and widely known is Wikipedia, the free online encyclopedia established in 2001. Open in this case means that virtually anyone can contribute a new article or edit an existing one. By the end of June 2006, Wikipedia had accumulated 4.2 million articles totaling 1.4 billion words in 250 languages, 2.3 million photographs and illustrations, more than five million links to other websites, and 85.4 million between-article cross-reference links, and occupied approximately 12 gigabytes. Wikipedia's explosive growth has been fueled by more than 300,000 volunteer contributors each of whom has made at least 10 changes to the encyclopedia, and two full-time system administrator employees of the non-profit Wikimedia foundation. Analysis of the quality of articles has found error rates to be only slightly higher for Wikipedia than for *Encyclopedia Britannica*.¹⁹

Distributed innovation systems will take hold first in areas of endeavor dominated by information and knowledge, but not necessarily limited to the purveyance of pure information goods. As more and more work of all types is done in-silico, more areas of economic activity will become amenable to the distributed innovation model. Separation of the information and physical components of goods will likely give rise to new organizational forms that specialize in either the information or physical foundry side of production. This has occurred already in the area of application specific integrated circuits, with the design component distributed to users around the world via specialized toolkits for innovation and the manufacturing component limited to specialized silicon fabs.²⁰ A similar trend has been observed in the sports equipment industry in which manufacturers have become conduits for transforming innovations conceived by a distributed base of users into mass-market products.²¹

Organizers of distributed innovation systems will be concerned not only with providing access to tools that enable information-based innovation, but also with the "architecture of participation" for contributors, which of necessity includes

consideration of the intellectual property regime that underlies their efforts. The architecture of participation is concerned with designing modularity and granularity into the task structure so as to minimize the cost of, and motivate with intrinsic or extrinsic (or both) rewards, participation by contributors. The task structure should facilitate accretion of effort in a way that improves the overall quality of a desired innovation. Attempts that enable many to participate, but limit the benefits of the outcome to a few, will likely fail. Successful, sustainable efforts will be built on IP regimes that reward participation with perpetual free use for contributors, attribution of effort, direct compensation, business models that do not restrict community engagement and development, or most likely some combination of these incentives.

CONCLUSION

Joy's Law applies to most organizations that are responsible for continually delivering innovations to stakeholders. Distributed innovation systems are an alternative approach to organizing for innovation that seem to meet the challenge of accessing distributed knowledge. They demonstrate the effectiveness of new methods and organizational structures for improving innovation outcomes by engaging a broader base of outside knowledge holders. Traditional organizations should not, however, seize on distributed innovation systems as some silver bullet that will solve their internal innovation problems. Rather, these systems are an important addition to an organization's portfolio of innovation strategies.

Those who would adopt or create a distributed innovation system, however, must be prepared to acknowledge the locus of innovation to be outside the boundaries of the focal organization. And this will require a fundamental reorientation of views about incentives, task structure, management, and intellectual property.

-
1. See Hayek (1945), "The Use of Knowledge in Society," *The American Economic Review*.
 2. See von Hippel (2005), *Democratizing Innovation*, Cambridge MA: MIT Press, for an excellent overview and analysis of the user-driven and distributed innovation paradigm.
 3. von Hippel, (1994), " 'Sticky information' and the locus of problem solving: Implications for innovation." *Management Science*.
 4. See Riggs and von Hippel (1994), "Incentives to innovate and the sources of innovation: The case of scientific instruments," *Research Policy*, for an analysis showing how novel functionality emerges from users.
 5. Data obtained by searching on this website <http://dan.corlan.net/cgi-bin/medline-trend?Q=>, accessed on September 16, 2007.
 6. Analysis done by Intellectual Property Practice of The Boston Consulting Group.
 7. See Amor-Iglesias et al, 2005, "Measuring Libre Software Using Debian 3.1 (Sarge) as A Case Study: Preliminary Results," Upgrade.
 8. See <http://lwn.net/Articles/222773/>, accessed on September 16, 2007
 9. See <http://www.techweb.com/wire/showArticle.jhtml?articleID=55800522>, accessed on September 16, 2007.
 10. Text shows messages sent to the comp.os.minix Usenet group between July and October 1991.
 11. See <http://kerneltrap.org/node/11>, accessed on October 2, 2007.

12. See Lakhani and Jeppesen (2007), "Getting Unusual Suspects to Solve R&D Puzzles," *Harvard Business Review*.
13. See Lakhani and Wolf (2005), "Why Hackers Do What They Do: Understanding Motivation and Effort in Free/Open Source Software Projects" in *Perspectives on Free and Open Source Software*, Cambridge, MA: MIT Press, for an overview of findings on motivations to participate in open source communities.
14. See Csikszentmihalyi (1990), *Flow: The Psychology of Optimal Experience*, New York: Harper and Row, for an excellent overview of how work can be fun and satisfying in a wide range of professions.
15. See Baldwin and Clark (2006). "The Architecture of Participation: Does Code Architecture Mitigate Free Riding in the Open Source," *Management Science*, for a very novel analysis of how code structure can impact community participation.
16. See von Hippel and Katz (2002), "Shifting innovation to users via toolkits," *Management Science*, for a discussion on how most users may be able to innovate through manufactured supplied tool-kits.
17. See Carlile (2004), "Transferring, Translating and Transforming: An Integrative Framework for Managing Knowledge across Boundaries," *Organization Science*.
18. On an absolute basis, however, even if only ten percent are successful, this represents more than 10,000 projects.
19. See Giles (2005), "Internet encyclopedias go head to head," *Nature*, for a comparison of Science articles in Wikipedia and Britannica.
20. See footnote 2.
21. See Hienerth et al (2006), "How user innovations become commercial products: a theoretical investigation and case study," *Research Policy*, for an economic model and case study of user-driven entrepreneurial effects.