

CHAPTER ONE

PEER-TO-PATENT: A MODEST PROPOSAL

You must do the things you think you cannot do.

—ELEANOR ROOSEVELT

PATENT LAW IS THE students' least favorite part of the semester-long class, Introduction to Intellectual Property, that I teach at New York Law School. In this survey course they learn about trademarking brands and copyrighting songs. But they also suffer through five jargon-filled weeks on how inventors apply to the United States Patent and Trademark Office (USPTO) to secure a twenty-year grant of monopoly rights. Despite the fact that patents signal innovation to the financial markets and investors and drive economic growth in certain industries, many dread this segment of the course.¹ Patent applications are written in a special language; patentese is a member of the legalese language family that only the high priesthood of patent professionals understands. Even applications for the most mundane inventions are written in dense jargon. The patent application for the sealed crustless sandwich (aka the peanut butter and jelly sandwich patent), which sought to give Smuckers a monopoly on a process to crimp crusts, reads as follows:

Claim: 1. A sealed crustless sandwich, comprising: a first bread layer having a first perimeter surface coplanar to a contact surface; at least one filling of an edible food juxtaposed to said contact surface; a second bread layer juxtaposed to said at least one filling opposite of said

first bread layer, wherein said second bread layer includes a second perimeter surface similar to said first perimeter surface; a crimped edge directly between said first perimeter surface and said second perimeter surface for sealing said at least one filling between said first bread layer and said second bread layer; wherein a crust portion of said first bread layer and said second bread layer has been removed.²

To help my students understand how patents further Congress's constitutional mandate to "promote the progress of science and useful arts," I start by teaching the process by which the government decides whether to grant a patent.³ While this process has its special rules, the decision to award or withhold a patent is not unlike a thousand other decisions made by government every day, decisions that depend upon access to adequate information and sound science. Just as an official of the U.S. Environmental Protection Agency (EPA) must consult epidemiological studies to determine acceptable levels of asbestos or mercury in air and water, the patent examiner must obtain the relevant technological antecedents—known as prior art—to judge if an invention is enough of an advance over what preceded it to warrant a patent. The patent examiner effectively decides who will control the next BlackBerry or the next life-saving cancer drug.

The Patent Office employs 5,500 patent examiners.⁴ While the examiner might have an undergraduate degree in computer science, she does not necessarily know much about cutting-edge, object-oriented programming languages. She's not up on the latest advances coming out of Asia. She may not have seen anything like the patent application for bioinformatic modeling of the human genome or the application for a patent on poetry-writing software!⁵ She has not necessarily been to law school (you don't need a law degree to take the patent bar exam).⁶ She does not necessarily have a Ph.D. in science, and there is little opportunity on the job for continuing education. As an expert in patent examination, she is not and is not expected to be a master of all areas of innovation.

To make things worse, the inventor is not legally required to give her any help—say, by providing background research.⁷ Indeed, the inventor has an incentive not to supply the Patent Office with prior art, since the examiner might use it to determine that the invention lacks sufficient novelty and thus to reject the application.⁸ Sometimes inventors deluge

an examiner with background research, hoping the overworked official will be daunted by the task of sorting the wheat from the chaff. It is no wonder that even Thomas Jefferson, the first patent examiner, in 1791 sought outside help, consulting with University of Pennsylvania chemistry professor Joseph Hutchinson before issuing a patent on an alchemical process for rendering seawater potable.⁹

Today the modern patent examiner works alone (or at most with a supervisor). Her primary resource is USPTO databases (known as East and West) of old and foreign patents, patent applications, and the prior art citations they reference.¹⁰ On average, she has just fifteen to twenty hours to research the patent application and write up her findings.¹¹ Worse yet, her supervisor (with Congress in the background) is breathing down her neck to move on to the next application in the backlog of a million pending applications.¹² Applicants wait upward of three years (and in certain fields closer to five years) to receive their first notice from the Patent Office, and that's usually just the beginning of a series of communications that will be exchanged before the patent is finally granted or rejected.

Even with more time, patent offices around the world still would not have access to the information they need. To know if a particular inventor is the progenitor of a chemical compound or software program, the examiner has to scour the literature. Government patent offices naturally have access to the historical corpus of patents, and they have access to excellent and up-to-date journals, but the information needed is not always found in traditional government or academic sources. Inventors in cutting-edge fields may discuss their work on the web rather than in print. John Doll, the U.S. commissioner of patents, complains of the dispersed databases and inconsistent search protocols that impede examiners' efforts to decide whether an invention is new, useful, and nonobvious—in a word, patentable.¹³ The result is an inefficient, inaccurate process: of the 2 million patents in force in the United States, many would not survive closer scrutiny.¹⁴

All this got me to thinking. What if the patent examiner worked with the broader community? What if the public augmented the official's research with its own know-how? What if the scientific and technical expertise of the graduate student, industry researcher, university professor, and hobbyist could be linked to the legal expertise of the patent examiner to

produce a better decision? What if, instead of traditional peer review, a process of open review were instituted, wherein participants self-select on the basis of their expertise and enthusiasm? What if, instead of a social network like Facebook, a scientific and technical expert network were built? I nicknamed this “peer-to-patent.” The online tools available today could be employed to connect the government institution and the increasingly networked public to collaborate on an ongoing basis.

Such a process is already happening outside of government. Some business and nonprofit organizations recognize that processes that were once the purview of an individual might usefully be opened up to participation from a larger group. Cancer patients, for example, provide medical information to each other via the Association of Online Cancer Resources website and its 159 associated electronic mailing lists. The website Patients Like Me allows patients to share information about their symptoms and the progress of their diseases. Patients Like Me also has data-sharing partnerships with doctors, pharmaceutical and medical device companies, research organizations, and nonprofits to encourage patients to supply information to those who are working to develop cures.

Other examples abound. Amazon’s web-based Mechanical Turk project outsources the work of answering simple questions, such as tagging people and places in pictures, measuring the size of molecules in a microscopic image, identifying land mines from photographs, and creating links to or from a Google map. YouTube depends on amateurs to post video content. Volunteers populate the Internet Movie Database (IMDb), which offers information about close to one million movie titles and more than two million entertainment professionals.¹⁵ Almost 30,000 Korean-speaking citizen-journalists report on stories for OhMyNews.com, where “every citizen is a reporter.”¹⁶ Korean speakers also answer each other’s search queries via the Naver search engine, which far outpaces the popularity of such algorithmic search engines as Google and Yahoo!¹⁷ The Mozilla Corporation, maker of the Firefox browser, enlists the help of several thousand of its 180 million users to work on marketing campaigns, respond to queries on Mozilla message boards, write or edit documentation for developers, and even create the software code for the browser.¹⁸

More than 9,000 companies participate in technology giant SAP’s global partner networks, and 1.2 million individuals participate in its online discussion communities, which are designed to generate innovation for the firm while making individuals more successful at their jobs.

Inspired by these examples, once the spring 2005 term ended, I wrote up a posting for my blog entitled “Peer-to-Patent: A Modest Proposal.”¹⁹ I proposed that the Patent Office transform its closed, centralized process and construct an architecture for open participation that unleashes the “cognitive surplus” of the scientific and technical community. I called on the Patent Office to solicit information from the public to assist in patent examination and, eventually, to enlist the help of smaller, collaborating groups of dedicated volunteers to help decide whether a particular patent should be granted. Through this sort of online collaboration, the agency could augment its intelligence and improve the quality of issued patents. “This modest proposal harnesses social reputation and collaborative filtering technology to create a peer review system of scientific experts ruling on innovation,” I wrote. “The idea of blue ribbon panels or advisory committees is not new. But the suggestion to use social reputation software—think Friendster, LinkedIn, eBay reputation points—to make such panels big enough, diverse enough, and democratic enough to replace the patent examiner is.”

Just as I posted my thought experiment the phone rang. Daniel Terdman, a reporter for *Wired News*, was trolling for stories. “Heard anything interesting?” he asked. I reeled off three or four initiatives of various colleagues. “That’s all well and good, but what are *you* up to?” Daniel probed, hoping I might have something to report. “Catching up on my blog and making improbable proposals to revolutionize the Patent Office, improve government decisionmaking, and rethink the nature of democracy,” I modestly replied.

On July 14, 2005, *Wired News* ran an article titled, “Web Could Unclog Patent Backlog.”²⁰ As a reporter who wrote about videogames, not government, Daniel was uninhibited about calling the patent commissioner for a quote. Commissioner John Doll responded: “It’s an interesting idea, and an interesting perspective.” Peer review, he added, “is something that could be done right now, and I’m a little surprised that somebody hasn’t started a blog” for that purpose.

THE MODEST PROPOSAL TAKES OFF

The day the article appeared, Manny Schecter, the associate general counsel and managing attorney for intellectual property at IBM, sent me an e-mail: “I saw the story on Peer-to-Patent. We should talk.” Manny

Schechter, Marian Underweiser, and Marc Ehrlich are known as the 3Ms of the intellectual property law department at IBM. Responsible for the company's 42,000 patents (28,000 in the United States alone), these three senior attorneys and their staff ensure that IBM continues its unbroken fifteen-year streak as the holder of the largest patent portfolio in the world. The firm now receives between 3,000 and 4,000 U.S. patents each year. In addition to strengthening the competitive position of IBM's products, these patents generate \$1 billion annually in licensing fees from other businesses wishing to incorporate IBM's scientific inventions into their products and services. The size of IBM's patent portfolio signals to the market that the firm is an innovator, which may be responsible for its rising share price and increased shareholder value.²¹

As the USPTO's biggest client, IBM is one of the companies with the most to gain from an efficient patent system. It also stands to lose if the patenting process breaks down. With the pace of patent examination out of sync with the pace of innovation, firms like IBM are forced to wait ever longer for patents. And these innovations, on which their licensing strategies depend, may even turn out to be invalid. In addition, critics charge that the granting of undeserved patents, in combination with growing uncertainty over patent quality, has led to an increase in costly litigation. Patents provide a license to sue others for damages for using a patented invention. Companies with deep pockets, such as IBM, are more likely to be sued for patent infringement than smaller firms. Software patents, which represent the bulk of IBM's portfolio, are more than twice as likely as other patents to be litigated.²² The cost of defending such a suit, even for the victorious, makes the game not always worth the candle, especially when the alternative is to pay the plaintiff a five- or six-figure fee.

The 3Ms, therefore, had been contemplating ideas for patent reform that were similar to Peer-to-Patent. The company had been experimenting internally with technology for distributed collaboration for a long time, and senior executives credit IBM's rescue from the brink (it is one of the 16 percent of large companies tracked from 1962 to 1998 to have survived) to the digitally aided development of a culture of collaboration.²³


IBM's lawyers were intrigued by the simplicity and promise of the Peer-to-Patent proposal, particularly since it could be implemented, at least as a pilot, without legislative or Supreme Court action. By spring 2006 they were ready to help the idea become reality. The 3Ms at IBM

offered a research grant to New York Law School to allow me to (ironically) take a break from teaching Introduction to Intellectual Property and to flesh out the blog posting into a design for a practical prototype. Little did I know that by yielding to the temptation of a semester off to write a research paper I would end up launching an experiment to improve the flow of information to the Patent Office and running the government's first open social networking project.

In short order, corporate patent counsel at the major technology firms began to hear about Peer-to-Patent, and Microsoft joined the project with a commitment to submit patents for public review and to contribute much-needed additional sponsorship. After all, it would smack of regulatory capture and delegitimize the work if the largest customer of the Patent Office were to be the sole supporter, designer, and funder of a plan to reform it. Then came Hewlett-Packard, followed by Red Hat, General Electric, CA (Computer Associates), and finally Intellectual Ventures, the invention company founded by former Microsoft chief technology officer Nathan Myhrvold. These companies not only offered to submit their patent applications through this process but also contributed money to the development of the legal and technical infrastructure. In addition, New York Law School received support from the MacArthur Foundation and the Omidyar Network, the organization that channels the philanthropic activities of eBay founder Pierre Omidyar.

Dozens of lawyers, technologists, and designers gave their time and expertise to refining the design of the project. The result was a series of workshops at Harvard, Yale, Stanford, the University of Michigan, and New York Law School during 2006–07. The planning of Peer-to-Patent created educational opportunities for New York Law School students, who practiced law reform and acquired professional skills by running the project at every stage. They produced educational videos about patent law and prior art (think Schoolhouse Rock for the patent system). They wrote the directions for each page of the website, explaining to new users how to find and upload prior art in connection with a patent application or how to comment on prior art submitted by others. Students also drafted privacy and copyright policies, terms of use, and solicitations to inventors to invite them to submit their applications. Above all, they learned how to work as a team, using technical, legal, and communication tools to implement a solution to a complex problem in the real world.

FIGURE 1-1. PEER-TO-PATENT HOME PAGE AT WWW.PEERTOPATENT.ORG



in cooperation with USPTO

Hello, Joanne | Logout | Contact Us

COMMUNITY PATENT REVIEW

ABOUT P-TO-P


APPLICATION LIST | ARCHIVED APPLICATIONS | US PATENT CLASSIFICATIONS

Home

Search

MY PROFILE

TUTORIALS



00:00 | 00:00

PEER TO PATENT VIDEOS
PJ (Pamela Jones)
Founder and Editor, Groklaw

[More Videos](#)

PEER TO PATENT ACTIVITY

<p>Discuss Patent Applications 399 comments posted</p> <p>Size of Community: 2017</p>	<p>Upload + Explain Prior Art 175 submitted</p>	<p>Annotate and Evaluate Prior Art 188 prior art ratings 175 citations</p>
---	--	---

WELCOME TO PEER TO PATENT

Peer-to-Patent opens the patent examination process to public participation for the first time. Become part of this historic pilot program. Help the USPTO find the information relevant to assessing the claims of pending patent applications. Become a community reviewer and improve the quality of patents.

- Click here to see a list of all applications.
- Click here to be notified of any new applications via RSS
- Click here to be notified about any new applications via email (requires login)

New Applications

- 10 New! Methods of de identifying an object data
- 12 New! Recommending contacts in a social network
- 10 New! Method and system for internet search
- 11 Detecting missing elements in a storage area ne...
- 14 Systems and methods for integrating a patient k...
- 11 Method and apparatus for an inductive doubling ...
- 10 Method and apparatus for computer network security
- 14 Method and apparatus for xml parsing using para...
- 12 System and method for ontology-based translatio...

Most Active Teams

Applications in Need

News

Most important, despite the first-of-its-kind nature of the plan, the USPTO was on board by the end of 2006. Sold first on the idea of channeling more information to overworked patent examiners from the website and second by a promise we made to forward only the ten most relevant public submissions, the Patent Office agreed to conduct this pilot in “open” patent examination. New York Law School hired Eric Hestenes, former vice president of technology for a large financial services firm, to lead a team of programmers in developing the million-dollar software platform that would help create teams of self-selecting scientific and technical experts to contribute information to the Patent Office online. (Not only have the cost of these tools come down, but options are available that would obviate designing from scratch. Strategies like cloud computing and shared services could enable many collaboration pilots to run off the same infrastructure.) A steering committee of corporate patent law experts and an advisory board comprising academics, journalists, and independent patent experts spent the next six months devising policies and designing the processes that would be enabled by the Peer-to-Patent website.

The USPTO then convened a team of eight executives to manage the project for the agency. Headed by the deputy commissioner for patent examination, Jay Lucas (later succeeded by John Love), and run by Jack Harvey, the director responsible for computer technology patents, the group collaborated with us on drafting the legal consent requirements. We agreed the pilot would include a maximum of 250 applications (enough to generate useful data without overwhelming the agency) relating to software and computer hardware (so-called Technology Center 2100 inventions), with a maximum of 15 applications from any one firm.²⁴ The USPTO chose and trained the examiners who would take part in the pilot, crafted the workflow processes for managing the receipt of public information and the distribution of feedback to the public, and circulated surveys to participating officials.

By 2007 press reports began to mount in anticipation of Peer-to-Patent’s rollout. *Fortune* profiled the project in August 2006, and the *Washington Post* devoted a front-page story to Peer-to-Patent in the spring.²⁵ Finally, on June 15, 2007—twenty-three months after my initial blog post—New York Law School, in cooperation with the USPTO, launched the Peer-to-Patent website.²⁶ (A screenshot of the home page is shown in figure 1-1 and can be found at www.peertopatent.org.) The

first five patent applications in the pilot came from HP, IBM, Intel, and Red Hat. The applications involved a wide range of computer technology relating to wind farming, virtual collaboration, and social networking. As an incentive to participate in this peer review process, the USPTO offered to examine Peer-to-Patent applications first, allowing companies to jump the million-application queue.

IMPLICATIONS FOR THE FUTURE: FROM WIKIPEDIA TO WIKILAW

In brief, here's how the collaboration works. As part of the process by which the patent examiner determines whether a patent application meets the legal standards set forth by the Patent Act, the Peer-to-Patent website solicits the public to submit information—namely prior art—relevant to evaluating a pending application. Because participating in this process requires enthusiasm and expertise, those who respond to the Peer-to-Patent invitation are self-selecting volunteers. Anyone can join but only an expert would. Participation requires working on an application in collaborative teams. Several team members might research the application, uploading relevant publications and suggestions for further research for use by the patent examiner. Others might comment on the relevance of submitted pieces of prior art. Following online discussion, each team vets the submissions made by its members. The group votes on which ten submissions are most relevant. Those are then forwarded to the Patent Office.

In the pilot's first year, inventors submitted eighty-four applications through Peer-to-Patent, and over 2,000 volunteers signed up to offer their expertise through the website. The numbers were small but the results demonstrated that the public is indeed in possession of information not readily accessible to the patent examiner and that, despite the complexity of the patent examination process, the public will take the time to contribute to it. At the end of the pilot's first year, 89 percent of participating patent examiners reported that the materials they had received from the public had been useful; 92 percent indicated they would welcome the opportunity to examine another application with public participation; and 73 percent wanted the Peer-to-Patent program implemented as regular office practice.²⁷ (This last number is quite high

considering that examiners had no way of knowing if institutionalized public participation would eventually put them out of a job!)

In June 2008 the USPTO extended the pilot for a second year and expanded the subject matter of Peer-to-Patent from computer software to include so-called business methods, or patent applications pertaining broadly to methods and processes for doing business (such as the one-click shopping cart). Support for this experiment in collaborative governance also came from outside the Peer-to-Patent community. Among many media mentions, the head of the Ewing Marion Kauffman Foundation, Carl Shramm, and its vice president for research and policy, Bob Litan, also wrote in *The American* that, “assuming this experiment proves to be as promising as it sounds, the next president should urge the PTO to adopt and Congress to accept this new way of assessing patents much more broadly.”²⁸ Jonathan Schwartz, CEO of Sun Microsystems, named Peer-to-Patent one of the “leading institutions promoting . . . patent reform.”²⁹ In his campaign’s technology platform, President Barack Obama called for incorporating Peer-to-Patent into USPTO’s regular procedure.³⁰ The U.S. Chamber of Commerce endorsed the adoption of Peer-to-Patent.³¹ In addition, Peer-to-Patent was nominated for the Prix Ars Electronica cyber arts prize for digital communities and the Silicon Valley Tech Museum Award for technology benefiting humanity.

To build on this incipient success, New York Law School established the Center for Patent Innovations to promote and facilitate public participation in the patenting process around the world. The patent offices of the United Kingdom and Japan were the first to follow the USPTO’s example, adapting Peer-to-Patent for their own national patent systems and launching similar websites in 2008 and 2009. The patent offices of Australia, Canada, and Europe are exploring online public participation for their own examination processes. The Trilateral Cooperation (the European Patent Office, the Japan Patent Office, and the USPTO) has begun to discuss a multioffice pilot to network the global scientific community to the national patent offices. Pending versions of patent reform legislation left over from the 2007–08 legislative calendar all include provisions to allow for third-party commentary on applications. These bills would also expand the regulatory authority of the USPTO to enable the agency unambiguously to adopt innovations like Peer-to-Patent.

But the implications of Peer-to-Patent reach far beyond the field of intellectual property. Far from being unique to the patent system, the lessons we learned about soliciting far-flung, self-identifying expertise to improve government decisionmaking can be applied to a broad range of environmental, educational, and other policy domains. The technology and social processes that drive Peer-to-Patent can be used to solicit participation in governance on the basis of professional expertise, or local context and experience, or willingness to do research and hard work. For example, the web could be used to structure participation by local communities in EPA decisionmaking about clean air and water. Technology could connect experts in every level of government to one another to solve problems more effectively and more efficiently. An online network of independent university experts—an online brain trust—could be created to advise. Citizen juries could be appointed to oversee the work of every cabinet official or agency head and generate greater accountability. Local groups could even be empowered to spend agency money, report back on how they addressed specific problems, and thereby become eligible for more funding. So much innovation is still possible.

Public conversation about the power of networks is already proliferating. Books such as the *Starfish and the Spider, Here Comes Everyone, Crowdsourcing*, and *Momentum* describe ordinary people coming together into caucuses mediated by technology to promote change. But while the new literature includes inspirational stories about the power of social networking tools, there is still a need for deep and serious thinking about how to apply what is learned about technology to the betterment of public policymaking—or how, in other words, to enhance political institutions with the power of networks.

Connecting the power of the many to the work of the few in government has little precedent, making it difficult to visualize its potential. As the NYU media scholar and critic Jay Rosen comments, “Crowdsourcing will not create any genuinely new things unless people know what is being asked of them.”³² Users of Wikipedia know what to do because they understand what it means to write an entry for an encyclopedia. People share a common image of that collective goal. But despite the growing popularity of online collaboration, experience is fairly limited when it comes to participating in government decisionmaking.

There are new networking tools available to go from Wikipedia to “Wikilaw.” The first government authority to start a blog was the Transportation Safety Authority, and that wasn’t until 2008!³³ While Silicon Valley and Route 128 develop increasingly powerful tools to connect people, policymakers downplay the role of technology in governance. They have not come to grips with the disruption created by this new way of working. This should not come as a surprise. Few institutions readily invite their own obsolescence. The *Encyclopaedia Britannica* did not create Wikipedia. The *New York Times* did not create Craigslist classifieds. Record companies did not create the MySpace social networking and music-sharing site. Existing institutions lack clear incentives to change their own business plans. More important, they lack a blueprint for doing so. Were it only a matter of more technology and a faster Internet, collaborative governance would have come to government long ago.

THE CORE IDEA

In a speech at the New America Foundation, Google CEO Eric Schmidt said about Peer-to-Patent:

At the Patent and Trademark Office, which is as overloaded as it has ever been, they’re running a very significant experiment where they publish the patent applications early for public comment. And guess what? All the players who cared deeply about this bizarre and nerdy patent really go after it because there’s no way where their patent examiners can fundamentally get all the insight that the wisdom of crowds can do. Why is that not true of every branch of government? It makes perfect sense, use all those people who care so passionately, and who have a lot of free time, to help you.³⁴

The presidential campaign generated unprecedented public engagement. The American public turned out in record numbers to vote and also to participate in getting out the vote. Thousands of experts joined policy committees to advise the Obama campaign via closed listservs, and tens of thousands of “ordinary” people participated in online policy discussions via the open transition website (change.gov). President Obama championed volunteerism with the launch of a national service initiative in honor of Martin Luther King’s birthday. The campaign

drew—or just drew attention to—a groundswell of enthusiasm for involvement, giving rise to the question, What next?

The Peace Corps and Americorps already offer opportunities for full-time engagement. Now VolunteerMatch and other web-based services help to hook up and reduce the coordination costs for the vast majority who prefer to do part-time community service. While the Internet may have increased participation in mass campaigns and enabled individual participation in civic life, both are divorced from the work of governing.

There is too little diversity of participation in the work of managing society, both participation in traditional government practices and innovative technological strategies that might connect government to the public to solve problems in new ways. While people can take full-time jobs in government, there is no equivalent of VolunteerMatch to connect a network of doctors to the Department of Health and Human Services or to allow a team of scientists to assist with evaluating climate change data for the EPA or the economist and the physicist to collaborate on modeling economic forecasts for Treasury. More to the point, government is not articulating priorities that enable the venture capitalist and the entrepreneur to build new businesses. There are too few projects where government articulates a problem and then the public coordinates the solution, such as NetDay did in the 1990s, when volunteers collaborated to connect local California schools to the Internet, in response to and with the encouragement of the federal and state governments.³⁵ While there are myriad public-private partnerships, these singular events do not address the opportunity for sustained collaboration and institutional redesign.

After Election Day, those who participated in the 2008 political campaign had the opportunity to engage in government directly. Traditional public participation practices, like peer review or federal advisory committees, select participants by means of complex vetting processes. But only a handful can ever serve. Yet outside of government people are coming together every day coordinated by Internet technology to strive toward common outcomes. Schmidt was therefore right to ask why there are not more opportunities for people to participate in governance.

It is overdue to rethink the legitimacy of attenuated participation in a small number of representative institutions. Instead, democratic theory and the design of governing institutions must be rethought for the age of networks. The opportunity now is to move toward collaborative

democracy (of which Peer-to-Patent is an exemplar), in which institutions afford the public the opportunity to select themselves to participate actively in diverse ways.

Collaborative democracy is a new approach for using technology to improve outcomes by soliciting *expertise* (in which expertise is defined broadly to include both scientific knowledge and popular experience) from self-selected peers working together in groups in open networks. By lending their expertise and enthusiasm, volunteer experts can augment the know-how of full-time professionals and coordinate their own strategies. By taking advantage of technology's cost savings, hierarchies can be transformed into collaborative knowledge ecosystems and radically change the culture of government from one of centralized expertise to one in which the public and private sector—organizations and individuals—solve social problems collectively.

The private sector has been quicker than government to recognize that making better decisions requires looking beyond institutionalized centers of expertise. Don Tapscott and Anthony Williams have chronicled this phenomenon in the private sector in *Wikinomics*.³⁶ In this IBM 2006 global study that asked chief executive officers where they looked for fresh ideas, they cited clients, business partners, and employees far more than their research and development labs.³⁷ IBM conducts digital brainstorming sessions known as *World Jams*, which allow IBM employees across the globe to make and refine proposals collaboratively for the improvement of the company. Far from being gimmicky online happenings, World Jams are taken so seriously by the blue chip company that the CEO of IBM established a \$100 million fund to implement the ten best resulting ideas.³⁸

A handful of employees in an institution—any institution—cannot possess as much information as the many dispersed individuals who make up a field. This is why Eli Lilly set up Innocentive Inc. to farm out problems from life sciences companies to a network of 160,000 “solvers.” One company recently paid a \$1 million bounty for the solution to a complex chemistry problem. The solver was not even a scientist but a lawyer with a knack for chemistry. He answered the intractable question in fewer than four hours! In technology, this insight has been popularized as Joy's law: “No matter who you are, most of the smartest people work for someone else.”³⁹ This quip, attributed to Bill Joy, cofounder of Sun Microsystems, pinpoints the core problem faced by all

organizations in an exploding information ecosystem, including government: most knowledge lies outside the boundaries of the institution.

Collaboration is distinct from the concept of crowdsourcing. Jeff Howe, an editor at *Wired* magazine, coined the term *crowdsourcing* to describe the burgeoning phenomenon of “taking a job traditionally performed by a designated agent (usually an employee) and outsourcing it to an undefined, generally large group of people in the form of an open call.”⁴⁰ (He does use Peer-to-Patent as his one public sector example.)⁴¹ But whereas crowdsourcing generally refers to aggregating the responses of individuals across a network, collaborative democracy aspires to the kind of intentional peer production and shared group effort of Wikipedia, in which volunteers sign up to write encyclopedia entries as a group. While crowdsourcing activities like prediction markets aggregate individual preferences, collaboration implies more robust and diverse coordinating structures that enable people to divvy up tasks and roles. Collaboration does not so much imply throwing people at a problem as coordinating the right people in different roles. Role differentiation not only helps to structure work done across a distance, it also conveys the sense of working as a team. Unlike peer production, which includes purely civic, bottom-up activities, collaborative democracy emphasizes shared work by a government institution and a network of participants. Collaborative participation is the “smoke-filled aquarium”—to borrow an overheard coinage—that combines open-source volunteer participation with government’s central coordination, issue framing, and bully pulpit.

In *Wiki Government*, the case for a collaborative vision of democratic theory is bolstered by three arguments woven through the book: collaboration as a distinct form of democratic participation, visual deliberation, and egalitarian self-selection.

First, collaboration is a crucial but not well understood claim of democratic practice. There is a belief that the public does not possess as much expertise as people in government. Furthermore, the technology has not previously existed to make collaboration possible on a large scale. These spurious assumptions have produced an anemic conception of participatory democracy. Participation has generally referred to once-a-year voting or to community deliberation, in which neighbors engage in civil dialogue and public opinion formation on a small scale. New social and visual technologies (sometimes referred to as web 2.0) are demonstrating

that people are knowledgeable about everything from cancer to software and that, when given the opportunity to come together on a network and in groups, they can be effective at solving problems (not only deliberating about them). We must therefore distinguish between deliberation and collaboration as forms of participatory practice (which we'll do more of in chapter 2). *Wiki Government* explores many examples of ordinary people joining together to do extraordinary things coordinated via the Internet. Peer-to-Patent is a paradigmatic case of database programmers and wind-farming experts working with patent examining professionals to make a better decision.

Second, the medium matters. To enable collaboration at scale requires designing the practices to make participation manageable and useful and then enabling those practices by means of technology. While the forms of participation will differ when information gathering or priority setting or data analysis is required, the technology should always be designed to reflect the work of the group back to itself so that people know which role they can assume and which tasks to accomplish. This second insight is what I term *visual deliberation*. In traditional deliberative exercises, strict procedures for who can talk govern the public conversation. But collaboration depends, instead, on having tools that convey the structure and rules of any given collaborative practice. This kind of social mirroring can be communicated through software. Peer-to-Patent uses visualizations to communicate the work flow by which information goes from the government institution to the public and back again. The website helps to convey what it means to review a patent application. It exploits rating and reputation techniques that help each group work together as a group, even across a distance. Hence, designing new democratic institutions also depends on designing the appropriate collaborative practices and embedding that design in software.

Third, collaboration is a form of democratic participation that is egalitarian—but egalitarian in a different way than the traditional understanding of the term. Typically, mass participation like voting is thought of as being quite democratic because everyone can participate in the same way. By contrast, Peer-to-Patent is not mass participation. It demands highly technical expertise. Successful participation depends upon the participant's interest in and knowledge of patents. If Peer-to-Patent were the only example of collaborative participation, it would not be egalitarian. But Peer-to-Patent multiplied by a thousand would be

more institutionally diverse and complex. If the patent expert and the doctor and the teacher each has a vehicle for engagement, contexts would be created in which they each uniquely possess expertise and derive meaning.

In other words, people do not have to participate in the same exercise. One person may want to work on Peer-to-Patent, another may want to get involved in health care debates. One person may want to work on energy policy; another may want to organize a corps of energy “scouts” to go door-to-door and help neighbors evaluate their energy usage. The ability to self-select to participate in the arena of one’s choosing is what makes collaborative democracy egalitarian. A person may be an expert on wetlands because she possesses professional credentialing. Another person may be an expert on wetlands because she lives near one. Perhaps it is a level of know-how or the enthusiasm to commit more time that generates status in other domains. For every project, there is a different kind of expertise, which could be sought. Experts will flock to those opportunities that exploit their intelligence. In this choice lies the equality of opportunity.

What does collaborative democracy look like in practice? In the old way of working, the bureaucrat might decide to repair a bridge in response to an opinion poll or vote that randomly obtains feedback. Or the bureaucrat might publish a fully developed plan to repair the bridge, ostensibly soliciting comment in response to a notice of proposed regulation, attracting participation by formal interest groups and lobbyists but not ordinary citizens, who can never hope to match the power and influence of corporate interests. Community groups might use the web to lobby for bridge repair but with no greater opportunity to get involved in detailed decisions. The government or a nongovernmental organization (NGO) might organize a face-to-face deliberative discussion about the bridge and hope to use the event to trigger a newspaper article that will influence the decision. A similar online discussion may or may not attract attention.

Under a collaborative strategy, the bureaucrat establishes the process, then frames and asks the questions that will get targeted information from bridge users (the truck driver, the commuter), from an engineer, and from the informed enthusiast. The public can contribute evidence and data to help inform specific decisions, analyze data once gathered, and share in the work of editing, drafting, and implementing policies. Alternatively, if officials articulate the priority of bridge safety, they might

spur private sector businesses, nonprofits, and individuals to develop their own strategies, such as organizing a volunteer corps of bridge safety inspectors who log their work on a shared website. Citizens are no longer talking about the process: they are the process.

The future of public institutions demands that we create a collaborative ecosystem with numerous opportunities for those with expertise about a problem) to engage. There is a Plum Book, which lists government jobs, and there is a Prune Book, which lists the toughest management positions. The pluot is supposed to be the sweetest variety of plum (or plum plus apricot). Yet there is no Pluot Book cataloging opportunities for part-time participation in government! When participatory democracy is defined to include diverse strategies for collaboration, when these thousands of opportunities to self-select come to light, a Pluot Book may well be needed.

OVERVIEW OF THE BOOK

This book offers a rethinking of the meaning of participatory democracy in the digital age. At the same time, it is a how-to guide for bringing about collaborative democracy and the practices of collaborative governance using the tools of law, policy, and technology. Practical experience with the Peer-to-Patent program enhances understanding of the core problem: a failure to grasp the changing nature of expertise in the digital age and the resulting misconception of both effective institutional practices and legitimate democratic theory.

Chapter 2 argues that the “single point of failure” in government can be transformed through new mechanisms for obtaining expertise. Decisionmaking is currently organized around the notion that the government official knows best. In reality, agencies make decisions every day without access to the best information or the time to make sense of the information they have. Citizen participation traditionally focuses on deliberation but, in the Internet age, it will not be as successful as collaboration in remedying the information deficit. The broader mandate is to use technology to upend the outdated theory of institutional expertise and replace it with collaborative practices for gathering and evaluating information and transforming raw data into useful knowledge.

Chapters 3 and 4 tell the story of the Peer-to-Patent pilot. Chapter 3 illustrates the single-point-of-failure problem by showcasing the crisis of

patent quality—the problem to which Peer-to-Patent was designed to respond. Whether or not one knows or cares about patents—though there is plenty of reason to do both—the information deficit faced by the Patent Office is paradigmatic of the practices of centralized decisionmaking in government. The aim in chapter 3 is therefore to provide a detailed account of how the Patent Office gets—or fails to get—the information it needs to make important decisions and to detail the consequences of this failure.

Chapter 4 begins to explain how to move toward a collaborative solution to the governance challenge described in chapter 3. It describes the development of the Peer-to-Patent website—what it is, how it worked, and why it worked—to illustrate the process through which innovative participatory practices can be designed and adopted. The story of Peer-to-Patent begins with an in-depth exploration of the innovative role of technology design in making citizen participation practices manageable. Instead of designing for deliberation—pure talk—I argue for what I term *visual deliberation*, namely, ways of using the computer screen to mirror the work of participating groups back to themselves so that they can organize and function as networked publics. Creative uses of the interface through which people interact with the computer and therefore with each other also make information manageable and intelligible and reduce the problem of information overload. From talking about the design of the collaborative project, the chapter concludes with a discussion of the collaborative design process that led to the creation of the project.

Perhaps the most important chapters of the book are those in part 3, “Thinking in Wiki.” These chapters generalize from the Peer-to-Patent project to online participation in other arenas of governance.

Chapter 5 focuses on the role of information in collaboration, arguing for a government information policy that enables the collection and distribution of information in ways that engender participation. Data can become more useful as a result of group participation. Groups not only can help to visualize information in graphic formats that make it more intelligible but these graphical formats can also focus the work on solving problems. As a baseline condition, information must be transparent—accessible, searchable, and usable—to lend itself to collaboration.

Chapter 6 examines the history of citizen collaboration and its future. This chapter situates Peer-to-Patent against the backdrop of transparency

and participation legislation and regulation. The aim is to uncover why—despite past attempts to introduce innovative and participatory practices into administration, including those that exploit Internet technology—agencies have not always had access to enough information nor have citizens enjoyed meaningful participation in government decisionmaking.

Chapter 7 asks what will produce such innovations in government. Peer-to-Patent was brokered by an outside organization that pushed for this citizen participation effort, building on the momentum of web 2.0 technologies. But to transform the culture of government and create lasting change, there has to be evangelism from within as well as without. This should be the job of the senior leadership, such as the new role of U.S. chief technology officer created by President Obama. Senior government management should use the bully pulpit to exhort public institutions to put collaborative democracy into effect. The CTO can be the champion of participatory innovations to connect institutions to public expertise. I offer examples of such innovations, including the policy wiki and the citizen jury, which might produce more open, and ultimately more legitimate, ways for government to work.

Finally, chapter 8 offers lessons for designing better practices to engage the public in government. These lessons apply both to information-gathering projects like Peer-to-Patent and to policy wikis, citizen juries, online brainstorming, and other innovations in participation. Collectively, these lessons form the basis of a new design science of government. Designing for democracy requires law, technology, and policy to create more effective institutions. Such a design approach has the potential to enhance the legitimacy of government; it also empowers participants. Ordinary citizens have more to offer than voting or talking. They can contribute their expertise and, in so doing, realize the opportunity to be powerful.

This book speaks to three audiences: those interested in the story of Peer-to-Patent as a lesson in patent reform; those aficionados of web 2.0 interested in a specific case study of how to apply collaboration in the government arena; and government reformers interested in improving decisionmaking. The chapters of the book unwind the argument about collaborative democracy and the role of social and visual technology in enabling collaboration. Patent experts may want to skim the patent problem in chapter 3 and focus, instead, on the specifics of Peer-to-Patent in chapter 4 and subsequent chapters that describe the lessons

learned. Web 2.0 enthusiasts who already “get” collaboration but do not know the government context can skim the book’s justification, articulated in chapter 2, and dig right into the story of Peer-to-Patent (chapters 3 and 4) and the challenge of collaboration in government (chapters 5 through 8). Government reformers with no particular patent bent will want to read the opening chapters 1 and 2 carefully to understand the distinction between deliberation and collaboration and then focus on the lessons of Peer-to-Patent in chapters 5 to 8.

Peer-to-Patent is an experiment. But that’s the point: the best strategy is to try something: to see what works to bring about a more engaged citizenry. Peer-to-Patent demonstrates a way to solicit help from those with know-how, passion, and enthusiasm.