

Wiki-Government

How open-source technology can make government decision-making more expert and more democratic.

George Bernard Shaw once wrote, “All professions are conspiracies against the laity,” and nowhere is this more the case than in a democracy. Although political legitimacy demands accountability to an electoral process, those living in a democracy readily submit to what sociologist Michael Schudson calls the “permanent embarrassment” of expertise. We believe that administrative governance by a professional elite is the best way to organize decision-making in the public interest. Experts decide on acceptable levels of mercury emissions in the air, anti-discrimination rules in education and the workplace, and the standards for cross-ownership of newspapers and broadcasting stations.

The justification for this professional decision-making, articulated by theorists ranging from Max Weber to Walter Lippmann, is that while citizens can

BETH SIMONE NOVECK is professor of Law and Director of the Institute for Information Law & Policy at New York Law School and the McClatchy Visiting Associate Professor in the Department of Communication at Stanford University.

express personal opinions based on values, they are incapable of making fact-based decisions on matters of policy. For Weber, the complexities of modern governance call for “the personally detached and strictly objective expert.” Only institutionalized and governmental professionals possess the expertise, resources, discipline, and time to make public-policy decisions. And citizen participation is hard to organize and administer, and even harder to scale. It is one thing for 10 bureaucrats to debate a policy and come to an informed consensus; try getting the same result with 10,000 people—or 10 million.

Now, however, new technology may be changing the relationship between democracy and expertise, affording an opportunity to improve competence by making good information available for better governance. Large-scale knowledge-sharing projects, such as the Wikipedia online encyclopedia, and volunteer software-programming initiatives, such as the Apache Webserver (which runs two-thirds of the websites in the world), demonstrate the inadequacy of our assumptions about expertise in the twenty-first century. Ordinary people, regardless of institutional affiliation or professional status, possess information—serious, expert, fact-based, scientific information—to enhance decision-making, information not otherwise available to isolated bureaucrats. Partly as a result of the simple tools now available for collaboration and partly as a result of a highly mobile labor market of “knowledge workers,” people are ready and willing to share that information across geographic, disciplinary, and institutional boundaries.

Consider how communal pooling of knowledge is already at work. An individual may submit the first “stub” about the history of the Ming Dynasty or the biography of Winston Churchill in Wikipedia, and a wider community of thousands collaborates on writing, editing, and refining every article. Wikipedia is open enough to allow expertise to emerge, but it is also structured enough, with outlines and to-do lists, to set the rules for a certain kind of group collaboration—and that collaboration is producing high-quality results.

Or take sites that utilize self-reinforcing “reputation” systems to improve quality and reliability. These so-called social-networking sites—like Dopplr for travelers, LinkedIn for business professionals, or Facebook and MySpace—use community rating and friend-of-a-friend (FOAF) accreditation mechanisms to build the reputation and trust necessary to form knowledge groups and communities. Making expertise relevant for the complex processes of policy-making also requires forming communities that can collaborate, but it goes beyond that. It demands “civic networking,” tools designed for groups to transform data into knowledge useful to decision-makers, as well as the concomitant institutional practices designed to make use of that knowledge.

Political philosophers from Aristotle to Rousseau to Rawls have suggested that when groups engage in the public exchange of reason, they produce better ideas. In practice, however, more talk usually slows decision-making and comes with the attendant problem of groupthink. Increasingly, however, we are discovering how to use computers to enable deliberation without endless talk and without having to be in the same room. And those structures—enforced through software—are what transform the subjective, free-wheeling, dynamic expertise of amateurs into effective communities of experts.

For example, the Omidyar Network, the philanthropy launched by eBay founder Pierre Omidyar, asks the public to participate in awarding its grants. Rather than invite submissions from thousands of individuals—which would have strained resources to review—the Omidyar Network created an online framework for the interested community to deliberate and winnow the proposals first. Or consider New Assignment, which was launched to demonstrate that “open collaboration over the Internet among reporters, editors and large groups of users can produce high-quality work that serves the public interest, holds up under scrutiny, and builds trust.”

The site set forth the social practices to elicit collaborative reporting (instead of collaborative gossip-mongering), resulting in the publication of seven original essays and 80 interviews, as well as a series of stories about collaborative journalism for *Wired* magazine. While this fell short of the number of pieces the organizers had wanted, New Assignment still enabled the “crowd” to produce stories as good as any found in a national magazine and demonstrated how to organize (and how not) the process. Similarly, the Sense.us program at the University of California–Berkeley provides public mechanisms to allow people across disciplinary boundaries to collaborate in making, and thereby making sense of, census data graphs and charts. And the United States Patent and Trademark Office (USPTO) is taking this idea to the next level: connecting experts directly to actual decision-making in the “Peer-to-Patent” project.

In each case, we are beginning to see how, if designed with clear, simple tasks to that help create more open and collaborative yet organized practices, the computer screen can shift power from professional sources of authoritative knowledge to new kinds of knowledge networks. Speaking truth to power is easiest to do—and more accurate—when spoken not as an individual, but as a group.

Ordinary people, regardless of institutional affiliation or professional status, possess information to enhance decision-making.

This has particular application to policymaking: Non-governmental participants have something more to offer than voting once a year—namely, good information. In much the same way that we devise legal procedures to ensure fairness in the courtroom or open deliberation in Congress, we can design technology—and the legal and policy framework to support it—that elicits specific, structured, and manageable input, not from individuals, but from collaborative groups. If we can harness the enthusiasm and knowledge of “netizens” to the legal and political processes generally reserved for citizens, we can produce government decision-making that is both more expert and, at the same time, more democratic.

The Problem with Experts

In his award-winning book *On Political Judgment*, social psychologist Philip Tetlock analyzed the predictions of those professionals who advise government about political and economic trends. Pitting these professional pundits against minimalist performance benchmarks, he found “few signs that expertise translates into greater ability to make either ‘well-calibrated’ or ‘discriminating’ forecasts.” It turns out that professional status has much less bearing on the quality of information than we might assume, and that professionals—whether in politics or other domains—are notoriously unsuccessful at making informed predictions.

Moreover, the traditional reliance on institutionalized expertise is fraught with political controversy. Sometimes these pre-selected scientists and outside experts are simply lobbyists passing by another name. The current administration, for example, regularly replaces experts on agency advisory panels with ideologues and political allies. In a published statement titled *Restoring Scientific Integrity in Policy Making*, over 60 preeminent scientists, including Nobel laureates and National Medal of Science recipients, lambasted the Bush Administration for “manipulation of the process through which science enters into its decisions.” But if Bush is among the more egregious violators of the presumed wall between politics and institutionalized expertise, his actions only go to show how easy it is for any executive to abuse his or her power of appointment to disrupt experts’ advisory function.

The problem of relying only on professionals in our bureaucracy has been compounded by institutionalized practices of confidential decision-making, even in areas where public input could have a decisive impact, such as the review of patent applications. Whereas a first- or second-year civil servant at another agency might be drafting memos, a patent examiner with limited supervision is doling out a patent that could impact the fate of an industry or fundamental scientific research. Yet, until 1999, the USPTO, in order to protect

the confidential trade secrets of inventors, did not publish applications. Even today, most (but not all) applications are only published after 18 months, and even then the examiner cannot communicate with the public. The USPTO allows third parties to submit publications with no commentary, annotations, or markings, and only for a fee, effectively eliminating participation by all but the most die-hard corporate competitors. As a result, despite a backlog of 700,000 applications last year (by now it is 800,000 and counting), the USPTO received only 40–100 such submissions, none of which were necessarily used in decision-making.

But if excessive reliance on governmental or selected professionals is the problem, direct public participation, as we have typically known it, is no panacea. Public participation in regulatory decision-making has been an established part of the federal landscape since the end of World War II, when it was enshrined as a right in the Administrative Procedure Act of 1946. This may sound good in the abstract, but there is such a thing as too much public participation. Whereas democracy is to be practiced at the polling booth, bureaucrats are not supposed to feel the pressures of direct democratic involvement. It can corrupt the rationality of the process with the distortions of private interests. And more participation does not mean better participation. How should administrators deal with individuals who carp but offer little useful information to improve decision-making? Or interest groups that electronically submit tens of thousands of identical “postcard comments”?

Scientific peer review provides an alternative mechanism for oversight and quality control. The practice is widespread in government grant-making: The National Science Foundation relies on a network of more than 50,000 reviewers, while the National Institutes of Health relies on outside review groups and advisory councils from the scientific community to review over 70 percent of its grant applications. And the Environmental Protection Agency’s grant-selection process relies heavily on “science review panels,” peer review groups chosen and managed by an outside scientist.

But peer-review panels, like agency advisory boards and other professional commissions, are empanelled, not self-selected. Membership is closed. Even if their deliberations are posted online or made available after the fact, the public has no say over who participates and no voice in the process. People have

How do we square the expertise emerging via Wikipedia with the professional standards which we expect from government?

no option to self-select on the basis of enthusiasm, rather than being chosen on the basis of profession. Even when not unduly subject to political influence, the decision as to who participates is based on institutional status. Those who may have meaningful contributions to make—graduate students, independent scientists, avid hobbyists, retired executives, and other consultants in the “free agent nation”—fall outside the boundaries of professional institutions and status and will of necessity be excluded, regardless of their intellectual assets.

Peer review is also time-consuming to organize and run. The group has to be selected, vetted, and approved, and because fights can arise over membership, conflicts of interest have to be identified and sorted out. Participants must also be convinced to join. It is, perhaps, in part because of the work that must go into maintaining a peer-review system that review generally happens only in limited contexts, and too late in the process to have the maximum impact on regulatory decision-making.

Nonprofessional Expertise

In contrast to what we see in government, in many other fields there is a move away from the preeminence of institutionalized professionals. Instead, technology aggregates and refines the knowledge of distributed, non-institutional experts. Patients, not doctors, are providing medical information to each other about cancer via the Association of Online Cancer Resources website and its 159 associated electronic mailing lists. Almost 30,000 citizen-journalists report on stories for OhMyNews.com. Amazon’s Mechanical Turk project outsources the work of answering simple questions or doing basic tasks to a distributed network. YouTube depends on amateurs to post video content. The Internet Movie Database (IMDb), which offers information about close to one million movie titles and more than two million entertainment professionals, started as a collection of movie trivia submitted by members of two online newsgroups. In South Korea, the Naver search engine, where Korean speakers answer each other’s queries, far outpaces Yahoo and Google as the most popular.

But how might this open, online collaboration improve governmental decision-making? How do we square the expertise and talent we see emerging via Wikipedia or YouTube with the professional standards of science to which we expect governmental decisions to conform? After all, Wikipedia has been known to contain errors and defamation. For every brilliant art film or newsworthy clip, there are thousands of pieces of video junk on YouTube. To put it bluntly, information quality may not be a matter of national security when it comes to a fifth-grade book report, but it is essential to nuclear regulations or environmental safety standards or the quality of issued patents. While we know that

excessive reliance on professionals is problematic, we do not want to replace one set of abuses with another by eliminating the professionals and replacing them with direct, popular decision-makers.

Rather, we want to design practices for “collaborative governance,” shared processes of responsibility in information-gathering and decision-making that combine the technical expertise of public experts with the legal standards of professional decision-makers. There are plenty of people with expertise to share if their knowledge can successfully be connected to those decision-makers who need it. It is not necessary to pre-select authenticated and known professionals when structures can be put in place to ensure that informational inputs are discernable, specific, well-labeled, and easy to search, sort, and use. An online system will not be without its own problems and abuses. But the assumption is that greater public participation, not only in setting values but also in supplying information or making sense of information provided by others, can substantively improve decision-making.

Peer-to-Patent

On June 15, 2007, the USPTO launched an experiment, the “Peer-to-Patent: Community-Patent Review Pilot,” which could become a model for precisely this sort of collaborative governance. The program solicits public participation in the patent examination process via the Web. This system (the design and implementation of which I direct in cooperation with the USPTO) allows the public to research and upload publications—known in patent law as “prior art”—that will inform the patent examiner about the novelty and obviousness of the invention and enable her to decide whether it deserves a patent. This is truly revolutionary: In the 200 years since Thomas Jefferson founded the patent office, there has been no direct communication between the patent examiner and the public.

“For the first time in history,” David J. Kappos, vice president and assistant general counsel at IBM, says in the *Washington Post*, patent-office examiners will be able “to open up their cubicles and get access to a whole world of technical experts.” With the consent of participating inventors, this USPTO pilot allows the self-selecting public to review 250 software-patent applications from such companies as CA, Hewlett-Packard, General Electric, IBM, Intel, Microsoft, Oracle, Red Hat, Yahoo, and several smaller firms. The community not only submits information, but it also annotates and comments on the publications, explaining how the prior art is relevant to the claims of the patent application. The community rates the submitted prior art and decides whether or not it deserves to be shared with the USPTO. Only the 10 best submitted prior-art references, as judged on the basis of their relevance to the claims of

the patent applications by the online review community, will be forwarded to the patent examiner.

The USPTO pilot program is neither a blog nor a wiki. It is not a free-for-all for “software patents SUCK!” comments. Rather, the software is designed to provide a structured environment that solicits specific information targeted to the decision-making process. Eligible applications are posted online for review for up to four months. There is a shared discussion space, but it requires registration and joining the group committed to reviewing that application. The group can deliberate about the application’s quality, decide what research needs to be done, discuss where prior art may be found, and even divvy up the work of finding it. At the same time, the group has tools to “filter” comments in the discussion, identifying those that are most important, contain a request for follow-

up research, or are low-quality “noise,” thus mitigating the influence of the person who simply talks the loudest or the most. Private-sector websites—from the book reviews on Amazon to the movie reviews on IMDb to the news postings on Slashdot, the technology news site—have shown that these community ratings can be aggregated with

In the 200 years since Thomas Jefferson founded the patent office, there has been no direct communication between the patent examiner and the public.

surprisingly accurate results to sort and filter discussion.

While the discussion is designed to foster belonging and community and offer a space where ideas can be refined collaboratively, none of the discussion is forwarded to the agency, limiting the opportunity for undue influence. Instead, only the work that will be most helpful to the decision-making process, namely the community’s research and the prior art, are forwarded. Those who join an application community research the background to the application and share that research, helping to guide the governmental professional in her work. The research may help the examiner identify fruitful avenues for her search, decreasing the work done in that limited 20-hour window or, at the very least, shortening wild goose chases.

To be sure, the patent examiner still conducts a search. She has all the same information available to her as before. But now she also has the results of this “human database.” By structuring the request for feedback, the agency avoids inviting participation it cannot use. And the public has an opportunity to participate in a way that is directly relevant to and will shape decision-making.

What’s more, through the software administrators can measure the level of expertise of public reviewers and thus better understand how this online

participation process shapes that expertise. They can also measure the impact of public participation on examiner decision-making and on the resulting quality of the issued patent. Over time, with the benefit of greater experience and more data, it may become possible to introduce more refined algorithms for assessing the quality of information on the basis of the past performance of citizen-participants. We can also better understand and anticipate the ways in which those with an interest to do so will attempt to “game the system” (an unavoidable part of the process) and improve the practice accordingly.

Driving this pilot program is a combination of public attention to patent reform and the availability of the technology to do open peer review online. Patent reform has become the subject of intense, recent legislative debate in the United States and in Europe. The problem is not so much with the patent office itself as it is with the explosion in the number of applications, which has put enormous strain on an old system, resulting in the issuance of low-quality patents that subsequently become the subject of expensive and wasteful litigation. Given the doubling of the number of patent applications in the last decade, examiners currently have less than 20 hours to review an application about the most cutting-edge nanotechnology, the latest genetic bio-science, or the most controversial financial business method. In that time, they have to search the USPTO’s limited databases to determine if there is a prior publication that would reveal that the application lacks the requisite novelty or significance required. Yet publications—which go beyond traditional scientific journal articles and include websites, software code, and products—are not all to be found in this closed database. Examiners must also contend with poorly drafted applications. There is no legal duty incumbent upon inventors to do a thorough search of the prior art and submit it to the agency.

The problem is as obvious as it is intractable: It is hard to find information that is not there. In other words, the patent examiner has to scour the scientific literature to figure out if a particular chemical compound has already been invented by someone else, lest she award the patent to the wrong person who turns it into the next billion-dollar pharmaceutical blockbuster. And she can’t delay her research, as it would risk exacerbating a backlog that will soon approach one million applications.

Rather than going to the experts, the USPTO realized that it is easier to let the experts find it—and that is precisely what the new program does. The process is open to anyone. At the same time, just as Wikipedia’s two million entries are actually maintained by a few thousand knowledgeable and dedicated die-hards, this patent public participation program does not need to attract everyone, just everyone enthusiastic enough to do the hard work of sharing in the burdens of

decision-making. Of course, for those who work or study in a particular area of innovation, participation should be anything but hard.

It is assumed that competitive self-interest will be one of the drivers causing people to get involved. Far from distorting the process, when someone from IBM or Microsoft (both of which have publicly announced that employees may participate in the program while on the job) has relevant information to contribute, ultimately the public benefits. Graduate students and those with deep knowledge but little status may also be inclined to participate, whether out of moral indignation that a patent is being applied for in an area of research that they know or out of the desire to gain status in, and professional recognition from, their community of interest. Incentives could eventually be created through prizes and rewards, such as a monetary bounty from an inventor for those who find the best prior art, or public recognition (an “Oscar” for public participation) to the reviewer and the team who contribute the most useful information. It is also an opportunity to generate interest in and a market for the invention.

The UK patent office has announced plans to do its own, much-wider pilot, adapted to its legal regime and cultural practices and the USPTO is contemplating expanding the scope of its own pilot. This will make it possible to begin to build a global community with expertise and interest and, at the same time, to test different practices, gather real data about how expertise can best be obtained in response to this complex subject matter, and refine the process for patents. It will also provide valuable learning about how such practices could be adapted for other areas of governance.

Evolving Institutions

Our institutions of governance are characterized by a longstanding culture of professionalism in which bureaucrats—not citizens—are the experts. Until recently, we have viewed this arrangement as legitimate because we have not practically been able to argue otherwise. Now we have a chance to do government differently. We have the know-how to create “civic software” that will help us form groups and communities who, working together, can be more effective at informing decision-making than individuals working alone. We know from James Surowiecki’s book, *The Wisdom of Crowds*, as well as Simon and Schuster’s new MediaPredict project (which encourages readers to guess which manuscripts will become best-sellers), that technology can be used to aggregate predictions. But Peer-to-Patent is teaching us that we can go beyond the tallying of votes. While the general public has good instincts about value-based decisions and could be engaged better to identify “big mistakes” (such as

egregiously unfair media ownership rules), there are specific people out there who possess specific information about patents or trucks or chemicals whom we can now incorporate into our decision-making.

Nor does it have to stop at patents. Scientists who currently give their time to review grant applications might be just as willing to contribute their knowledge to decision-making about the environment, transportation, nuclear power, and agriculture. Frequent travelers have useful information to share with homeland security officials about how to best organize security at airports. Economists, businesspeople, and lawyers know a great deal about financial markets, securities, and consumer protection. State Department officials do not possess better information than select graduate students in computer science about RFID chips for passports. Immigrants and welfare recipients have information based on lived as well as learned experience to contribute.

In order to make it possible for ordinary, busy people to participate, and for government to make use of their knowledge, government must design the practices of public participation to enable groups, not just individuals, to participate. Agencies must crystallize

the questions they ask of the public and embed those targeted practices into software. Of course, not even the best-designed civic software is going to stop government from making bad decisions or ignoring—willfully or otherwise—good information dropped in its lap. And there will be manipulation and gaming, to be sure. But these are not reasons to shy away from opening up the practices of governance, especially when those practices can evolve over time to respond to problems that may arise.

Opening up closed decision-making also introduces a greater degree of transparency and accountability than we have had before. For example, even if Peer-to-Patent does not yield good research or prior art every time, many eyeballs on the application still encourage the inventor to do a better job of writing it and produce public debate and discussion about the application. It drives more information into the open and encourages the “liberation” of government data, not for its own sake, but as an enabler for engagement. And public involvement reminds the agency official that he is working for (and being watched by) the public.

To bring about the new revolution in governance, the next president ought to issue an executive order requiring that every government agency begin to

We have the know-how to create ‘civic software’ that will help us form communities who can be more effective than individuals working alone.

pilot new strategies for improved decision-making. For example, he or she could require that each agency, as part of their Semi-Annual Regulatory Agenda delivered to Congress and as part of a new collaborative governance report to the Office of Management and Budget (OMB), set forth at least one “Peer-to-Policy” experiment to see how it could make its decision-making practices more collaborative.

Experimentation with community feedback would provide the impetus for independent rulemaking agencies, such as the Environmental Protection Agency, to post questions to the scientific community before enacting regulations. It would encourage the Department of Labor to create an online network to solicit the expertise of those with disabilities. These new opportunities for participation are not limited to rulemaking activity. The FCC can seek targeted

but open advice on the economics of spectrum policymaking as part of its efforts to decide which rules to draft. The National Oceanic and Atmospheric Administration might create an online marketplace for the collaborative creation of weather maps. The National Institute for Standards and Technology could award venture funding for scientific

By enabling agencies to try different practices, we can learn a great deal about how to construct new and better practices of public participation.

innovation and research, which would benefit from public collaboration and input. With public collaboration, decisions about everything from health care initiatives to housing programs to efforts to attract foreign direct investment could be improved by collaboration with public experts.

And imagine if this executive order were backed by an industry “X-Prize,” a large monetary reward for the private or public sector entity that designs the best new civic software tools to support more open decision-making. This X-Prize—or Peer-to-X-Prize—would provide a matching grant to support agency pilot programs and allow OMB to give this software to agencies. It might also go to support the development of tools to be disbursed on a state and local level. The aim would be to create turn-key technologies that can be deployed by any government decision-maker to enlist the help of her “public staff” and give her the tools she needs to send out a short questionnaire, enable collaborative drafting on a document, invite proposals and suggestions, or perform any number of useful activities to engage the interested and expert community. By enabling agencies to try different practices—while bringing down the costs of procuring technology—we can learn a great deal about how to construct new and better practices of public participation.

WIKI-GOVERNMENT

By being explicitly experimental with new forms of digital institution-building, we have an opportunity to increase the legitimacy of governmental decisions. The tools—increasingly cheap, sometimes free—will not replace the professionals. Technology will not, by itself, make complex regulatory problems any more tractable, or eliminate partisan disputes about values. What this next generation of civic software can do, however, is introduce better information by enabling the expert public to contribute targeted information. In doing so, it can make possible practices of governance that are, at once, more expert and more democratic. ■