

11

Toward Participatory Expertise

Shay David

Most of us are familiar with the ways in which the Internet expands access to specialized knowledge such as medical information, legal advice, scientific research, and other domains of expertise. In many ways, the “databasing” of the world described by Geoffrey C. Bowker (Chapter 2, this volume) is also the process of disembedding information that was once more tightly bound to professional communities, with their tightly controlled forms of accreditation and membership. Several of the essays in this volume are concerned with the production of knowledge within more open, digitally mediated communities, and in particular with the ways in which technical architectures enable or constrain the production of new forms of expertise and authority within those communities. T. L. Taylor (Chapter 7, this volume), danah boyd (Chapter 8, this volume), and Robert F. Nideffer (Chapter 12, this volume) analyze different aspects of these broader *structures of participation* in our digitally mediated culture. This chapter focuses on online systems where the burden of legitimating knowledge production and incentivizing participation is explicitly embedded in the system architectures themselves—systems that manage reputations and rewards and structure editorial processes and community relationships. Through a series of short case studies of Wikipedia, Slashdot.org, Amazon.com, and Experts-exchange.com, this chapter explores the larger “solution space” of online knowledge communities and works toward a more general theory of the changing modalities of knowledge production and authority in digital culture.

The dynamic between accredited or expert knowledge and informal or practice-based knowledge has been a persistent focus of science studies, with numerous applications in science policy, workplace informatics, and system design. At a practical level, understanding the role of informal knowledge and patterns of communication within institutions and social systems has proven central to understanding and *improving the capacity for* change within those systems, ranging from changes in public policy to the successful introduction of new information systems within workplaces.¹ Such research has demonstrated that formal representations of knowledge flows—the chain of command or institutional flowchart are the classic examples—rarely provide adequate descriptions of how knowledge really circulates within institutions. The top-down application of expert decisions or institutional change may not acknowledge informal practices or have sufficient legitimacy to supplant them. This recognition has given rise not only to descriptive theories of knowledge systems but also to normative ones dedicated to

proving the value of extending the horizon of decision making beyond expert cadres.² Such work intersects a much broader tradition of democratic theory concerned with the administration of complex societies, with 20th-century American touchstones in the works of John Dewey and Walter Lippmann.

Our investigation of online knowledge communities raises two questions fundamental to this line of inquiry—one primarily political, the other sociological. First, what is the relationship between expertise and democratic governance in complex systems? Cadres of experts and institutions for training, certification, and accreditation are among the defining features of modernity and structure much of society's complex division of labor. Yet claims to superior knowledge sit uncomfortably with notions of democratic accountability. They operate in tension with the values of broad-based participation in decision making and public discourse, and of informed consent to authority.

For the sociologists Harry Collins and Robert Evans, this tension frames a dominant research paradigm within science studies. They call for the establishment of a research agenda dedicated to understanding the legitimacy of knowledge claims and, more specifically, expert knowledge claims in the encounter with larger publics (Collins & Evans, 2002, p. 235). For online knowledge communities that offer substitutes for the knowledge claims of professional groups, this problem is fundamental and two-sided. How, on the one hand, do online systems create legitimacy within their (large) user communities? And how do they establish it vis-à-vis (larger) external publics?

The second question puts us on terrain more familiar to sociologists of professions: Through what social processes do certain actors acquire "jurisdiction" over technical matters (Abbott, 1988)? In the cases of highly organized professions like medicine or law, this course involves processes of recognition by and institutionalization within society, with the grant of "exclusive rights" to members, professional bodies, and training and certification institutions. Professions manage knowledge—that is, they add to it but also, crucially, they regulate opportunities to acquire and exercise it, and, importantly, to exclude others from practicing their expertise within the regulated domain.

Across a wide range of fields, the freer flow of information made possible by the Internet has eroded these monopoly positions. Improved access to information and the empowerment of the end users is one side of this

story. New software-based community architectures for learning, gaining recognition, and developing and exercising expert authority is the other. Online knowledge communities encompass and blur these positions, making distinctions between production and consumption problematic. They struggle continuously and in highly concrete ways with what Collins and Evans call the problem of “extension,” that is, the problem of understanding and managing the effective limits of participation and reliance on expert authority. As large, unbounded communities become producers and mediators of specialized knowledge, how do they organize the collaborative processes and what is the role of experts within it? Increasingly, the answer lies in the development of “reputation systems” within these community sites, systems that create and allocate forms of recognition, hierarchy, privilege, and authority within the community.

Online communities are important sites for understanding these dynamics because, increasingly, their designs speak directly to problems of legitimation and extension. The challenge here is not, strictly speaking, the “open source” paradigm, in which information production is tackled through large-scale, loosely coordinated voluntary efforts, but rather the subset of projects that embed the process of aggregation, filtering, and quality control in the system itself. These systems operate as real-life experiments in managing reputation, experience, and incentives for participation, directed both toward learning and applying complex bodies of knowledge. The most successful of them have not only mobilized large-scale participation but they have also adjusted to the emergent characteristics of the systems themselves, as users explore the opportunities and constraints of their online environments.

To reframe our starting question, then, this chapter explores how online knowledge communities address problems of legitimation and extension by privileging participation over prior accreditation at the level of the system architecture itself. The four cases examined here—Wikipedia, Slashdot.org, Amazon.com, and Experts-exchange.com—permit a rough sketch of the “solution space” that defines (and limits) these communities and their encounters with educational systems, professions, editorial norms, and other more established systems of knowledge production.

At a moment when our concept of democracy has not acquired much purchase over the technical architectures of our shared digital culture, these community systems are actively reconstructing concepts of accountability,

transparency, and public deliberation. How far they can go in deferring problems of scale (extension) and in surmounting their narrow project boundaries is—I would argue—a fundamental question for both the social and technical sciences that has only begun to be explored.

community expertise

Wikipedia is the largest encyclopedia in the world, with more than 1,000,000 entries in English (as of January 2007) and millions more entries in other languages. Wikipedia entries are open to and editable by the public. There is no centralized editorial control or allocation of entries to experts. Contrary to most norms of editorial process and quality control, Wikipedia demonstrates that—under some circumstances at least—large numbers of amateur contributors can create a dynamic in which “good” information drives out “bad,” leading to a form of conditional authority vested in the collaborative editorial process itself, rather than in the contributions of individual authors.

Slashdot.org is a technology website that advertises itself as “news for nerds.” It combines the functions of a news syndication service with those of a “letters to the editor” section. As a community architecture, Slashdot relies on a system of user feedback to dynamically allocate privileges to comment and contribute. This system establishes rank within the Slashdot community, measured in capacity to shape the content of the site and empower other users to do so. Slashdot has tremendous authority in the world of “geeks,” and claims more readers than *The New York Times*.

Amazon.com is the world’s largest book and CD store and is becoming an important retailer in other domains (mostly electronics and toys). It sells items directly but also facilitates a market for users to transact with one another. It does so, in part, via a complex system of user reviews, which have helped consolidate a form of community authority over markets in which information flows between peers (usually customers), instead of residing primarily in advertising copy. The Amazon review process allows users not only to review items but also to construct identities as expert reviewers and promote their social agendas.

Experts-exchange.com is a community for “trading” programming knowledge. It supports a marketplace for knowledge that uses a currency called “points,” which reflect ongoing collective judgments about the reputations and contributions of community members. Users spend points to obtain answers to questions and earn points by answering other users’ ques-

tions. The points economy anchors a system of recognition and reputation and provides a proxy for expertise and social capital.

At the heart of these systems is a new means of addressing a familiar problem of social scale. In small communities, expertise can be recognized and affirmed through personal contacts—news and reputation spread fast in small networks; reputation can be both enduring and “thickly supported” by multiple interactions. In large communities, in contrast, personal contacts are much less likely to be adequate for this role. Although the social distance between actors scales much more slowly than the population size—reflecting the “small-world” characteristics of human societies identified by network theorists—size and complexity have necessitated many forms of signaling of roles and expertise, from diplomas to uniforms. Systems of accreditation exist, in part, to solve the information problem inherent in large social networks: how to provide low-cost signaling of expertise in contexts where personal relationships map only a small portion of the population. The systems discussed here are attempts to bridge this difficulty, scaling up toward large populations while ignoring a priori accreditation as a basis for expert legitimation. Expertise in these communities is a continuous category, generated through a feedback loop between participation and community recognition. In practice, the distinction between expert and layperson is often diffuse or transitory. In some systems, the software architecture tries to associate expertise with community authority, the latter defined as the ability to shape the participation of others within the online space. In other contexts, expertise and authority are more granular, valid for a short time or within a subdomain. The power relationships that structure most expert–lay distinctions are filtered through this more fluid conception of roles and community, in which the interaction is not simply about problem solving by experts, but also an occasion for the transmission of knowledge and the reproduction of expertise.

Under these conditions, expertise is not strictly or solely vested in individuals. In many online communities, the notion of community expertise plays an important role. Community expertise resides in two competing dynamics: (1) the legitimation of “aggregate” opinion, as opinions tend toward an equilibrium, even on controversial issues; and (2) the openness of the community to dissenting opinions, with the potential to change the aggregate consensus. In successful online communities, these dynamics produce confidence in the knowledge-making process, rather than confidence

in any particular instance of expert opinion. The system is legitimate to the degree that it represents a properly constituted authority in the eyes of the community of users.

case study: encyclopedias without borders

Wikipedia's magnitude of many millions of articles dwarfs *Encyclopedia Britannica's* 80,000 articles and *Encarta's* 40,000 (Pink, 2005). Wikipedia is the highest-profile application of MediaWiki, a software tool that, like other "wiki"³ software, creates an open-community development space for viewing, editing, and linking web pages. Wikipedia is free to users; its low operating costs were initially covered by private donors and later supplemented by institutional contributions. For a site that shifts editorial and authorial functions onto a volunteer community, this openness is a condition of success. The system allows information consumers to become information producers at nominal cost. Of equal significance, it draws no distinctions between the latter group: There is no space for asserting credentials or other "outside" expert privileges; the system makes no judgments about the sources of knowledge held by participants.

In practice, almost all major articles are collaboratively authored, sometimes by dozens or hundreds of contributors. Contributors can register and acquire "verified" identities, or remain unregistered and anonymous. This distinction has underwritten a dynamic in which most new entries come from registered users, whereas anonymous users are more often associated with the editing of established entries.⁴ Many entries follow a predictable pattern of expansion and consolidation as participants add to articles, and as this (often disorganized) accretion raises incentives for major consolidating edits. The open structure leads to volatility when proponents of hard-to-reconcile positions dispute entries. This dynamic tends to favor the most committed posters, who can devote more time and energy to monitoring the editorial process. (Like other wikis, MediaWiki allows users to roll back pages to earlier versions.) Persistent conflict in some Wikipedia subcommunities led to the establishment of a code of etiquette in 2004, defined primarily by bans on personal invective and a "three revert rule" that limited the number of reversions allowed within a 24-hour period. A court of arbitration was established to exclude offenders from the site.

Although some disputes have occasioned considerable ill will, Wikipedia continues to grow because the community process operates sufficiently

well to resolve most conflicts within posting communities. The community for any entry forms in a back channel to the entry page—the Talk page—where content changes are explained and negotiated. By most accounts, major disputes are infrequent, and contentious topics tend to stabilize around consensual language. Cases of vandalism are usually self-healing, often within a few minutes (IBM, 2005).

Wikipedia has been used as an example of paradigm shift in the production of knowledge goods, from production by hierarchical organizations (e.g., firms) to what Yochai Benkler (2002) calls the rise of “commons-based peer production,” characterized by a broad-based community model coordinated by new communications technologies. Although open source software is the usual reference point here, the history of encyclopedias can also be reframed through this lens:

In the beginning, encyclopedias relied on the One Smart Guy model. In ancient Greece, Aristotle put pen to papyrus and single-handedly tried to record all the knowledge of his time. . . . [In] the 1700s, Diderot and a few pals (including Voltaire and Rousseau) took 29 years to create the Encyclopédie, ou Dictionnaire Raisonné des Sciences, des Arts et des Métiers. With the industrial revolution, the One Smart Guy approach gradually gave way to the One Best Way approach, which borrowed its principles from management science and the lessons of assembly lines. Encyclopedia Britannica pioneered this approach in Scotland and honed it to perfection. Large groups of experts, each performing a task on a detailed work chart under the direction of a manager, produced encyclopedias of enormous breadth. [Wikipedia represents] a third model—call it One for All. Instead of one really smart guy, Wikipedia draws on thousands of fairly smart guys and gals. . . . Instead of clearly delineated lines of authority, Wikipedia depends on radical decentralization and self-organization—open source in its purest form. Most encyclopedias start to fossilize the moment they’re printed on a page. But add Wiki software and some helping hands and you get something self-repairing and almost alive. A different production model creates a product that’s fluid, fast, fixable, and free. (Pink, 2005)

Pink’s conclusion echoes Wikipedia’s own description of its virtuous circle between process and content. For Pink and the Wikipedia staff, Wikipedia is not just a new way to write an encyclopedia but also a better epistemological model, rooted in greater responsiveness to change and a range of less tangible assertions, including the claim that the accuracy of the encyclopedia improves over time:

As anyone can edit any article, it is of course possible for biased, out of date or incorrect information to be posted. However, because there are so many other people reading the articles and monitoring contributions using the Recent Changes page, incorrect information is usually corrected quickly. Thus, the overall accuracy of the encyclopedia is improving all the time as it attracts more and more contributors. You are encouraged to help by correcting articles and passing on your own knowledge. (Wikipedia FAQ, 2005)

This set of claims should be read as Wikipedia's central social and epistemological hypothesis: Unstructured collaborative authorship yields not just community, but quality. No one doubts that significant parts of Wikipedia realize this promise, although there has been no systematic survey or effort to establish criteria of comparison. It is clear that not all entries enjoy this virtuous dynamic; Wikipedia's coverage is strongest where subjects attract high degrees of attention. Although Wikipedia covers an astonishing range of obscure and minor topics, many of these are the products of single or small-group authorship. In these circumstances, outliers in quality and "bias" are frequent.

Because of this variability, Wikipedia has been met with recurrent questions about trust and authority, especially as its explosive growth brings it into contact with communities with other habits and expectations of expert reference. Representatives of traditional editorial cultures have weighed in, the most visible and visceral being Robert McHenry (2004), a former editor in chief of *Encyclopedia Britannica*, who compared Wikipedia to a public restroom, which one can use in times of need, not knowing who used it before. But questions have also been articulated from within the community of users, partly in an effort to articulate the goals of Wikipedia and examine the consequences of its openness. How authoritative is Wikipedia? For whom? Compared to what? What does trust mean in this context?

Wikipedia's lack of a system for distributing authority among its participants—whether based on software-defined "karma" or community recognition or some other mechanism—makes it uniquely open but also uniquely weak in its ability to articulate enforceable standards over practice or content. The back channel negotiation over entries fulfills this function in a limited, local sense, but at any time dialogic authority can be trumped by those wielding more free time. The lack of endogenously developed leaders who can wield greater power within the community is a design choice, one that reflects Wikipedia's claims about the relationship between breadth of partici-

pation and quality of outcome. For Larry Sanger (2004), a Wikipedia founder and subsequent critic of the open editorial process, this belief is inseparable from a broader anti-elitism in the Wikipedia community, which not only insists on the invisibility of credentials but also actively disrespects expert opinion. He notes the possibility of a negative dynamic in which experts are driven away through their parity with ignorant contributors. Sanger envisions an eventual “fork” in the project, in which Wikipedia is reconstituted around an editorial committee; more recent conversations about Wikipedia have discussed the possibility of a print version built on editorial review, or of a rolling “accreditation” of vetted versions of important pages. Such compromises become important as Wikipedia’s owners begin to conceive it as more like *Britannica*, with corresponding responsibilities toward verifiability, stability, and “citability,” and less like the vast but unpredictable knowledge base of the Internet. How such steps would affect the community dynamic that sustains Wikipedia is unclear. Distancing a premium product from the community process will alter the sense of community responsibility over entries. Freezing or publishing “high-quality” entries reintroduces problems of the scalability of production that the wiki enterprise was intended to overcome. Wikipedia, like the larger open system of the Internet, is charting new territory in which widespread use is not contingent on widespread trust, at least on the terms established by earlier expert systems.

case study: distributing the news— all the news that fits the screen (and then some)

Slashdot.org’s tagline, “News for Nerds—Stuff that Matters,” states a challenge familiar to many media organizations: how to decide what is “fit to print,” in the sense of meeting the expectations of its community of readers. Slashdot started in 1997 as a simple bulletin board for referencing technology-related news articles; it encouraged readers to link to articles and to add unmoderated commentary. As the site grew, its nondiscriminating editorial policy began to break down. The site was overwhelmed by submissions: It was unclear how to keep content fresh as old material remained active, or how to prioritize good over bad content. Twenty-one-year-old Rob “CmdrTaco” Malda, Slashdot’s founder, understood two things about this process: (1) The main value of the site resided in the user comments, not the stories (which were published first elsewhere); and (2) the only way to scale Slashdot was to leverage the community’s cognitive power. Rather than

create an editorial hierarchy, Malda developed a community-based selection process built into the architecture of the site that relies on moderation and “karma points” in order to control the way users submit and rate comments. The system rewards both participation on the site and community judgment of that participation. As the Slashdot FAQ (2005) explains:

Karma is the sum of your activity on Slashdot. This means posting, moderation, story submissions. . . . Karma is used to remove risky users from the moderator pool, and to assign a bonus point to users who have contributed positively to Slashdot in the past.

Your karma is a reference that primarily represents how your comments have been moderated in the past. Karma is structured on the following scale “Terrible, Bad, Neutral, Positive, Good, and Excellent.” If a comment you post is moderated up, your karma will rise. Consequently, if you post a comment that has been moderated down, your karma will fall. . . . In addition to moderation, other things factor into karma as well. You can get some karma by submitting a story that we decide to post. Also, meta-moderation can cause your karma to change. This encourages good moderators, and ideally removes moderator access from bad ones. . . . Note that being moderated Funny doesn’t help your karma. You have to be smart, not just a smart-ass.

Every 30 minutes, the system distributes “tokens” to users based on the number of comments posted. These tokens turn users into moderators and are valid for three days. Comments float up or down on the story page based on the aggregate judgments of the moderators. Comments whose ranking drops below a certain threshold are no longer displayed. Skill in ranking stories and comments creates a feedback loop in the form of karma points, which allow users to build both reputation and influence within the system. The point system is layered to permit the acquisition or loss of karma across several forms of participation on the site. It is important to note that the system is not completely open: initial submissions are filtered by a core group of “authors” (i.e., editors). The community does the rest, deciding what matters and what does not.

Through this process, the expert function of editors in the selection and qualification of content is shifted to the Slashdot community. Like Wikipedia, participation is not conditioned by any a priori qualifications (like Wikipedia, there is a free process of registration.) Unlike Wikipedia, however, Slashdot’s system is designed to produce a hierarchy among the

contributors. Through the acquisition of good karma, moderators can influence the visibility of stories; their actions carry disproportionate weight in a context in which comments on stories often number in the hundreds. The moderator is not defined by substantive expertise—although he or she may possess it—but by a stronger investment in the community and a superior ability to cater to its information needs.

Since its founding almost a decade ago, revisions to the Slashdot architecture have been motivated chiefly by challenges of scale—how to maintain a relatively open community and a relatively efficient filtering mechanism for information as membership grew (from 10,000 articles in the year 2000 to 10,000,000 postings in 2004). There have been several such revisions in this period, including the major shift to the karma-based system of metamoderation in 1999. An unintended consequence of this growth and efficiency is the emergence of the “Slashdot effect”—the “roving random distributed denial of service attack,” according to Wikipedia, that can occur when participants visit websites identified in popular posts en masse (usually the victims are small websites with limited server capacity or bandwidth). Slashdot, unlike Wikipedia, Amazon, and Experts-Exchange (discussed below), does not internalize its transactions—its product is the collective attention of its members focused on stories outside the site. In economic terms, the Slashdot effect is a negative externality of the production process. And although it was first identified on Slashdot, it comes into play in many large online communities dedicated to efficiencies in the attention economy—to the filtering of news or other information on the Web.

case study:

book reviews and the politics of experience

New York Times Book Review editor Charles McGrath recently asked, only half in jest, whether “there had ever been a book that wasn’t acclaimed” (Safire, 2005). McGrath’s comment was directed at the “praise inflation” in book promotion and the perceived resulting decline in the usefulness of reviews, blurbs, and other commentary as indicators of quality. Traditionally, the critics employed by institutions like *The New York Times* or other leading newspapers and trade magazines served as cultural gatekeepers, shaping wider public perception of the quality of a work. In recent years, however, the overabundance of books, CDs, and other media has precluded comprehensive coverage by small groups of paid experts. Filling the gaps are a wide

variety of trade publications and websites that cover more specialized subfields of the culture industries and increasingly employ systems that harness the power of user communities in offering critical assessments of works. In many areas of cultural production, user reviews are mushrooming as an alternative to traditional expert reviews.

Amazon's online review system is a pioneering example of an increasingly common class of commercially oriented community review systems. The Amazon system manages an "accreditation community" around its products, with multiple layers of contributors, avenues for participation, and mechanisms for dynamically assigning visibility and credentials to popular reviews and reviewers. The system is structured in terms of what I have elsewhere called the "six degrees of reputation" of online expert systems (David & Pinch, 2005). Amazon's representation of a work is shaped by this dynamic process, which integrates the reputation effects of (1) the authors or creators themselves, whose name recognition and reputation accompanies the work; (2) promotional book reviewers, drawn from media sources or produced for Amazon by its employees; (3) community reviewers and participant-authored "best-of" lists, which also provide an opportunity to rate the work on a numeric scale; (4) reader judgments of the book reviews (useful/not useful), with an option for reporting inappropriate reviews; (5) a dynamic feedback mechanism for reviews that adjusts their visibility based on community opinion (usefulness) and reviewer rank; and (6) a dynamically adjusted system for reviewer rank, reflecting the amount of reviews posted (level of contribution) and the usefulness quotient (quality of contribution.) Reviewers that reach the upper ranks of the list receive visual accreditation in the form of an icon next to their name: "Top 100 reviewer" (or 500, 1,000, etc.). Top 20 reviewers receive a personal profile that accompanies their reviews.⁵

Amazon has a clear stake in managing the economy of trust within its review system. Recent qualitative research suggests that the review dynamic can positively affect sales (Chevalier & Mayzlin, 2004), providing a definitive commercial metric for the system. Amazon's reviewer ranking—especially the personalization of reviews that accompanies high rank—is an effort to replicate the perceptions of trust that underpin a managed editorial process. This takes precedence over, and in most instances displaces, the community building that accompanies the back channel exchanges on Wikipedia, as well as the social networks and system of interpersonal recognition that underpin editorial privileges on Slashdot. It is much closer to the signaling function of

formal credential systems in large offline networks (e.g., university degrees or uniforms), where the primary function is not to ground a social relationship but to enable the low-cost negotiation of reliable information or services. In other words, Amazon does not have a community in the senses described above; it has customers, and consequently does not confuse process with product. Like Slashdot and similar systems, it supports an open process of accreditation that operates without reference to an established norm or corpus of knowledge. Authority is generated through the process, and other members of the reviewer community are the only points of reference.

Is the information generated in this system trustworthy? Recently, a technical fault on the Canadian division of Amazon revealed the identities of several thousand of its “anonymous” reviewers, and alarming discoveries were made. It was established that a large number of authors had received glowing testimonials from friends, husbands, wives, colleagues, or paid professionals. A few had reviewed their own books or launched attacks on rivals. The fact that the system survived this scandal doubtless reflected a variety of factors, such as the complex shaping of taste and opinion that make the review system only a partial (and highly substitutable) factor in book purchases. But it also arguably reflected a number of overlapping dynamics particular to the architectural choices of Amazon and other commercial sites: (1) the thinness of community investment in the review process, which made corruption by individuals less of a collectively felt affront; (2) the force of numbers and consensus in the review process, which underpinned a stable—if not dialogic—structure of community expertise; and (3) the metaranking mechanism itself, which offered a modest check on reviewer power.⁶

case study: experts-exchange allowing experts to exchange expertise and reputation

Experts-Exchange.com is a peer-to-peer community specializing in the trading of technical support and other IT-related knowledge on topics ranging from operating systems to networking, programming, web design, security applications, and much more. The community is built on question-and-answer transactions between members filtered through a system of “points,” which operates as a form of internal pricing system. Points ground a reputation system that rewards those who can provide answers. Over time, the accumulation of questions and answers has produced a large knowledge base available to all members of the community (as of January 2007, the site listed

over 15 million questions). The site is oriented toward technology professionals and has more than 220,000 registered users (as of January 2007). Questions range from the trivial (the second most popular question is “How do I recover a lost Windows administrator password?”) to complex programming issues. The large number of participants raises the likelihood that even highly specialized questions will find an answer very quickly (something this author can attest to through participant observation).

Experts-Exchange thus has a strong problem-based orientation, a community dynamic, and a system of rank that distributes visibility (like Amazon), but not power (like Slashdot) within the community. Unlike Wikipedia, it requires neither user collaboration on the end products (answers), nor does it appeal to an abstract epistemology (neutrality, objectivity) to ground the legitimacy of answers. Questions give rise to dialogue between participants, and answers are valid when the poser of the question is satisfied. The technical nature of the material ensures that the knowledge transacted is verifiable: An answer either solves the problem or it doesn't, and this can be determined very rapidly.

The site's knowledge economy revolves around two types of points: “question points” and “expert points.” Registered members receive five question points every day, cumulative to 500, which they can assign to new questions. The points system enables some flexibility in signaling needs. The number of points assigned to a question may not indicate the difficulty of the answer, but rather a need for rapid response. The greater the number of points assigned to a question, the higher the incentive for community members to answer it. When another member takes up a question, an online conversation ensues until the question is answered to the satisfaction of the poser. Sometimes other experts join the thread. If the poser accepts an answer as valid, the points attached to the question are transferred to the appropriate responder. In addition, the poser submits a numerical grade of the quality of the answer in function of the difficulty of the question and grants the responder expert points equal to the question points multiplied by the number grade. Although the difference seems subtle, the second type of point permits recognition based not just on an a priori valuation of the knowledge sought, but also on the quality of the interaction and the virtuosity of the response. Expert points are the glue of the community.

Members are ranked based on total points, and these ranks are visible to the community. Like on Amazon, rank equates with personal visibil-

ity to the larger community. These signaling mechanisms include entry into a Hall of Fame, with personal profiles, printable certifications that reflect points accumulated (ranging from Master at 50,000 to Genius at 1,000,000), and awards like Expert of the Year and Most Valuable Experts. Such recognition structures mimic traditional forms of accreditation in some respects but, importantly, are intended to incentivize participation, not serve as a prerequisite to it.

Because questions and answers accumulate in a searchable knowledge base, Expert-Exchange fosters the growth of a knowledge commons that benefits the whole community. The knowledge commons includes answers to a vast array of common questions, which no longer need to be processed through the community dynamic. Because Experts-Exchange's role is limited to that of a transaction broker propelled by a modest reputation system, it manages to separate expertise from authority. Authority over the community is negligible in this structure: Expert status conveys no ability to shape the participation of others, nor can question posers steer their requests toward reputable experts. Members self-select on both sides, providing an example of the virtues of open systems in the allocation of limited resources: time and expertise.

toward a taxonomy of online knowledge communities

The above comparisons suggest that the design of online knowledge communities—at least in this early phase of development—operates within a relatively small “solution space,” marked by a few recurrent tensions and tradeoffs. Our very preliminary account suggests important distinctions between:

- Reputation systems that reward performance with *visibility within* the community versus those that distribute *authority over* the community (Experts-Exchange vs. Slashdot). The two are not mutually exclusive, but mark a distinction between “soft” and “hard” power in these communities.
- Reputational rewards based on substantive knowledge of a topic versus reputation based on the ability to reflect community preferences (Experts-Exchange or Wikipedia vs. Slashdot). The categories blur where process expertise involves catering to a community based in substantive expertise, as in the case of specialized editorial and filtering work.
- Systems that favor the experience of community versus those that privilege the production of knowledge goods. Arguably, both the problems and strengths of Wikipedia reflect a structural lack of clarity on this issue.

- **“Bounded” knowledge systems that internalize their products versus unbounded ones that repackage and affect resources distributed across the web. The “internal transaction” orientations of Amazon and Experts-Exchange versus the news-filtering function of Slashdot illustrate this dynamic. Google is arguably the extreme case: a mechanism for distributing visibility so ubiquitous that it virtually defines its environment.**

The new structures of participation in online communities are built around these parameters, and create new—if still only modestly exploitable and experientially thin—opportunities for learning, sociability, and advancement. It seems likely that the scope and character of these opportunities will grow as the architectures for online communities become richer and more sophisticated. Stronger approximations of “offline” systems of identity, dialogue, experience, and trust will continue to offer tempting directions for development. At present, these systems offer very thin representations of human experience; the exclusion of a priori credentials is also, in large measure, the exclusion of the authority of the past. The premium placed on participation is a premium on *current participation*, whose benefits erode rapidly with disuse. Although these systems have articulated new strategies for addressing problems of scale within community-based reputation systems (thereby pushing the problem of extension further back than previously thought practical), it may be that the limits of the attention economy prove less flexible. Advancement within these systems, to date, is often an exercise in singular and constantly renewed devotion, a fact that may limit their growth or raise incentives for new, portable, persistent forms of reputation.

The disempowering of the past associated with these systems is accompanied, ironically, by a near-complete capacity to record and search that past. These systems have long “community memories” that preserve not only records of achievement, but also the complex social negotiations that produced them, including dissenting opinions, past revisions of articles, rejected answers to technical questions, low-rated commentary, and so on. Collectively, these records trace the community’s shared experience in addressing both substantive questions and process concerns. Theoretically, these histories are available as a resource to the community, but the practical impact of this preservation is unclear. Online knowledge communities lack historians.⁷ More precisely, they lack history functions: systems that would expose the historical record of the community and integrate it into present practice.

toward a reputation society

Like other aspects of digital culture, online knowledge communities are social laboratories where community values, goals, and modes of interaction must be designed rather than received. The variations on the theme are also explorations of how basic social parameters translate into the much thinner context of digital sociability (danah boyd [Chapter 8, this volume] has explored this question in relation to social networks). The reputation systems that structure these communities illuminate a much broader dynamic between reputation or authority derived from a community, and credential systems that bring authority to communities. Online knowledge communities are not microcosms of these larger social systems but new formations within and in continuity with them. Looking forward, one of the most interesting questions for online knowledge communities is how this contact—and even competition—with established knowledge systems will be negotiated. The trials of Wikipedia are probably just an early taste of this process of integration. It involves both the (currently minimal) portability of personal reputations built within these communities and the reputation of the community knowledge base itself vis-à-vis other, more established knowledge producers. Most definitely, it involves the market power of these large new participatory publics and the resulting incentive for traditional media and expert communities to accommodate themselves to emergent forms of accreditation and reputation. A cursory look at book ads in the *New York Times Book Review* suggests that this process has begun: For example, a recent novel *Twilight* is advertised as “A *New York Times* Bestseller,” “A *Publisher’s Weekly* best book of 2005,” and “An Amazon ‘Best book of the decade . . . so far.’”

Many participants in and observers of these communities have high expectations for this encounter, and indeed envision community-based reputation systems as a potentially powerful force in opening forms of technocratic decision making to democratic participation and scrutiny. Hassan Masum and Yi-Cheng Zhang (2004) make this connection in their “Manifesto for the Reputation Society”:

Our judgments on any complex topic are inevitably transmitted to each other in an incomplete and distorted fashion. The task of reputation system designers is then to set up incentives that minimize inaccuracies and maximize productive collaboration, so that wherever possible the judgment of a group—or indeed, of an entire society—becomes better than the judgment of its individual members.

What would be involved in using community-based “karma” to manage reputation and authority in traditional professional domains? Such extensions are not impractical: the site RateMyProfessor.com, for instance, aggregates and quantifies the array of informal college student–circulated information about classroom teaching. It tries to objectify an aspect of professional performance that has no organized metric or feedback mechanism within the academy. Other systems, like Wikipedia, more directly confront existing professional prerogatives. Looking forward, the test of community-based knowledge communities in the broader public domain is not whether they generate relevant information, but how they negotiate the border with established systems of knowledge production, expertise, and credentialization. In thinking about a normative theory of expertise—of what kinds of expertise we should empower and respect—online knowledge communities show us, at the very least, how to place a premium on those modes that report to a large community. After the fundamental principle of open source code—of technical architectures exposed to public scrutiny—online knowledge communities may be the second democratic innovation of digital culture.

references

- Abbott, A. (1988). *The system of professions: An essay on the division of expert labor*. Chicago: The University of Chicago Press.
- Anthony D. (2005, September). *Explaining quality in open source: Zealots and good Samaritans in the case of Wikipedia*. Paper presented at the Social Economics Conference, Ithaca, New York.
- Benkler Y. (2002). Coase's penguin, or Linux and the nature of the firm. *Yale Law Journal*, 112(3), 369-446.
- Chevalier, J., & Mayzlin, D. (2004). *The effect of word of mouth on sales: Online book reviews*. Available online at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=432481
- Collins, H. M., & Evans, R. (2002). The third wave of science studies: Studies of expertise and experience. *Social Studies of Science*, 32(2), 235-296.
- David, S., & Pinch, T. (2005, July). *Six degrees of reputation: The use and abuse of online recommendation systems*. Paper presented at the Annual Conference of the Society for Social Economics, Budapest. Available online at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=857505
- Epstein, S. (1995). The construction of lay expertise: AIDS activism and the forging of credibility in the reform of clinical trials. *Science, Technology, and Human Values*, 20(4), 408-437.
- Hilgartner, S. (2000). *Science on stage: Expert advice as public drama*. Palo Alto, CA: Stanford University Press.
- IBM. (2005). *History flow: Visualizing dynamic, evolving documents and the interactions of multiple collaborating authors*. Available online at <http://researchweb.watson.ibm.com/history/index.htm>
- Masum, H., & Zhang, Y. (2004). Manifesto for the reputation society. *First Monday*, 9(7). Available online at http://www.firstmonday.org/issues/issue9_7/masum/index.html
- McHenry, R. (2004, November 15). The faith-based encyclopedia. *Tech Central Station*. Retrieved May 5, 2005, from <http://www.techcentralstation.com/111504A.html>
- Orr, J. E. (1996). *Talking about machines: An ethnography of a modern job*. Ithaca, NY: ILR Press.
- Pink, D.H. (2005). The book stops here. *Wired Magazine*, 13(3). Available online at http://www.wired.com/archive/13.03/wiki_pr.html
- Safire, W. (2005, May 1). Blurbosphere (on lanuage). *The New York Times, Sunday Magazine*, p. 12.
- Sanger, L. (2004, December). Why Wikipedia must jettison its anti-elitism. *Kuro5hin*. Available online at <http://www.kuro5hin.org/story/2004/12/30/142458/25>
- Terdiman, D. (2005, December 5). Growing pains for Wikipedia. *News.com*. Available online at http://news.com.com/Growing+pains+for+Wikipedia/2100-1025_3-598119.html?tag=st_lh
- Wynne B. (1989). Sheep farming after Chernobyl: A case study in communicating scientific information. *Environment*, 31(2), 10-15.

- 1 See, for example, Orr (1996) on the role of informal knowledge sharing among Xerox repair men, and Wynne (1989) on the response by local sheep farmers to post-Chernobyl nuclear fallout in northern Britain.
- 2 See Epstein (1995) on the importance of the use of lay experts in the early days of HIV medication clinical trials. See also Hilgartner (2000) on the construction of expertise as a form of public performance.
- 3 *Wikiwiki* is the Hawaiian word for quick.
- 4 The distinction does not appear to translate into measurable differences in quality: The quality of anonymous contributions is comparable to that of registered users (Anthony, 2005). Anonymity has shielded periodic, and occasionally much remarked, abuses of the system, including the notable false biography of former presidential aide John Seigenthaler (Terdiman, 2005), which occasioned a minor scandal in November 2005. In order to discourage malicious posts, Wikipedia founder James Wales announced in December 2005 that all contributions would require user registration. It is less clear how this will meaningfully verify identity, since false registration is quite simple under the current system.
- 5 The six-degrees model is useful for analytical purposes, showing how reputation is built into different layers of the system. In reality, many users “game” the system, trying either to promote an agenda or build their own identities. A full account is given in David and Pinch (2005).
- 6 Compared with Slashdot, Amazon’s feedback mechanisms (in 2005) are fairly rudimentary: Usefulness is a binary ranking condensing a wide range of possible opinion (e.g., something can be useful but not accurate, or not useful but funny). This metaranking affects the order of the reviews on the site.
- 7 Distinct from Crane’s vision (Chapter 3, this volume), of a vastly enriched experience of history built on location-based technologies, there is growing concern in the academic history community about the lack of tools for understanding online and other electronic forms of social experience, which have no spatial referents or necessary connection to broader matrices of human experience. Although online activities are, in many respects, more transparent and recordable than many “real-world” events, as a practical matter there is little systematic capacity or perceived need to preserve this history, or render it tractable to analysis. The Center for History and New Media at George Mason University (<http://chnm.gmu.edu/about.php>) hosts a project on the history of Wikipedia and of open source software projects that represents one effort to address this growing gap.

