Peer-to-Peer Review: Authority in Digital Scholarly Networks

Kathleen Fitzpatrick, Pomona College
kfitzpatrick@pomona.edu

[What follows is a draft of a chapter from a book in progress, entitled "Planned Obsolescence: Publishing, Technology, and the Future of the Academy." Please contact the author before quoting.]

> *In a world where knowledge is being made available at a rate of millions of pages per day, it is comforting to know that some subset of that knowledge or science has been critically examined so that, were we to use it in our thinking for our work, we would be less likely to have wasted our time. — Ray Spier, "The History of the Peer-Review Process"*
>
> *[E]lectronic publishing distinguishes between the phase where documents are placed at the disposal of the public (publishing proper) and the phase where 'distinctions' are being attributed. It used to be that being printed was 'the' distinction; electronic publishing changes this and leads us to think of the distinction phase completely separately from the publishing phase.*
>
> *However, doing so changes the means by which distinction is imparted, and imparting distinction is a sure sign of power. In other words, those who now hold that privilege are afraid of losing it ('gate keepers') and they will [use] every possible argument to protect it without, if possible, ever mentioning it. — Jean-Claude Guédon and Raymond Siemens, "The Credibility of Electronic Publishing: Peer Review and Imprint"*

introduction

For the last two years, I have worked in collaboration with the Institute for the Future of the Book, my colleague Avi Santo, and a range of prominent scholars in media studies, on MediaCommons, an all-electronic scholarly publishing network. We've planned, we've blogged, we've held meetings, we've tested some small scale implementations of the technologies we hope the network will employ, we've published a few test-run articles — and in all of the feedback that we've received, in all of the conversations we've had with scholars both senior and junior, both beginning and established, one question has repeatedly resurfaced: what are you going to do about peer review?

I've suggested elsewhere that peer review threatens to become the axle around which the whole issue of electronic scholarly publishing gets wrapped, like Isadora Duncan's scarf, choking the life out of

many innovative systems before they are fully able to establish themselves.[1] This is a flippant response, to be sure; concerns about peer review are quite understandable, given that peer review is in some sense the *sine qua non* of the academy. We employ it in almost every aspect of the ways that we work, from hiring decisions through tenure and promotion reviews, in both internal and external grant and fellowship competitions, and, of course, in publishing. The work we do as scholars is repeatedly subjected to a series of vetting processes that enable us to indicate that the results of our work have been scrutinized by authorities in the field, and that those results are therefore themselves authoritative.[2]

But as authors including Michael Jensen of the National Academies Press have recently argued, the nature of authority is shifting, and shifting dramatically, in the era of the digital network.[3] Scholars in media studies have avidly explored such shifts as they affect media production, distribution, and consumption, focusing on the extent to which, for instance, bloggers have decentralized and even displaced the authority structures surrounding traditional journalism, or the ways that a range of phenomena including mashups and fan vids have shifted the previously assumed hierarchies that existed between media producers and media consumers, or the growing tensions in the relationship between consumers, industries, and industry regulators highlighted by file-sharing services and battles with the RIAA. These changes are at the heart of much of the most exciting and influential work in media studies today, including publications such as Siva Vaidhyanathan's *The Anarchist in the Library*, Henry Jenkins's *Convergence Culture*, and Yochai Benkler's *The Wealth of Networks*, projects that have grown

---

1 See Fitzpatrick, "MediaCommons: Scholarly Publishing in the Age of the Internet."

2 Mario Biagioli compellingly argues for an understanding of peer review as not simply productive of disciplinarity in an intellectual sense, but as a Foucauldian mode of disciplining knowledge itself, a mode that is "simultaneously repressive, productive, and constitutive" of academic ways of knowing (11). He pertinently distinguishes Foucault's discplinary reference points in medicine and the prison from the discipline of peer review, however, as only in the academy do we find "that the roles of the disciplined and the discipliner are often reversed during one's career" (12), indicating the ways that peer review functions as a self-perpetuating disciplinary system, inculcating the objects of discipline into becoming its subjects.

3 See Jensen, "Authority 3.0."

out of an understandable interest in the extent to which the means of media production and distribution are undergoing a process of radical democratization in the Web 2.0 era, and a desire to test the limits of that democratization.

To a surprising extent, however, scholars have resisted exploring a similar sense in which *intellectual* authority might likewise be shifting in the contemporary world.[4] One might see such a resistance manifested in the often unthinking and over-blown academic response to Wikipedia — for instance, the Middlebury College department of history's ban on the use of the online encyclopedia as a research source, and the debate that ensued — which seems to indicate a serious misunderstanding about the value of the project.[5] Treating Wikipedia like any other encyclopedia, by consulting only the entries, as John Seely Brown has been heard to say, runs the risk of missing the point entirely, as the real intellectual heart of the project lies on the history and discussion pages, where one can see the controversies inherent in the production of any encyclopedia entry enacted in public, rather than smoothed over into an untroubled conventional wisdom.[6] Centralized projects like Citizendium, which seek to impose traditional, hierarchical modes of review on a project like Wikipedia,[7] ignore the fact that, first, the wiki is in its very architecture a mode of ongoing peer review, and second, that not only the results of that review but the records of its process are available for critical scrutiny. Failing to engage fully with the intellectual merits of a project like Wikipedia, or with the ways in which Wikipedia represents one facet of a far-reaching change in contemporary epistemologies is a mistake that we

4 Perhaps unsurprisingly, this issue has been taken on by librarians, if not by scholars; see, for example, Regalado.

5 See Cohen. Many pro-Wikipedia commentators responded to the Middlebury ban by noting, quite sensibly, that college students shouldn't be citing encyclopedias in the first place. The locus of most of the concern about Wikipedia in this case, however, was the fact that "anyone" can edit its entries.

6 Panel presentation, Claremont Graduate University colloquium on Social Entrepreneurship and Design, 1 June 2007.

7 The creators of Citizendium claim that they hope to create a "an enormous, free, and *reliable* encyclopedia," which "aims to improve on the Wikipedia model by adding 'gentle expert oversight' and requiring contributors to use their real names." The suggestion, of course, is that uncredentialed contributors require such expert guidance, and expert status is conferred through traditional modes of authorization. See Citizendium, CZ:About.

academics make at our own peril. As one librarian frames the issue, "Banning a source like Wikipedia (rather than teaching how to use it wisely) simply tells students that the academic world is divorced from real-world practices" (Badke, qtd, in Regalado). The production of knowledge is of course the academy's very reason for being, and if we cling to an outdated system for the establishment and measurement of authority at the very same time that the nature of authority is shifting around us, we run the risk of becoming increasingly irrelevant to the dominant ways of knowing of contemporary culture.[8]

For this reason, what I am absolutely *not* arguing in what follows is that we need to ensure that peer-reviewed journals online are of equivalent value to peer-reviewed journals in print; in fact, I believe that such an equation is instead part the problem I'm addressing. Imposing traditional methods of peer review on digital publishing might help a transition to digital publishing in the short term, enabling more traditionally-minded scholars to see electronic and print scholarship as equivalent in value, but it will hobble us in the long term, as we employ outdated methods in a public space that operates under radically different systems of authorization. Instead, we must find ways to work with, to improve, and to adapt those new systems for scholarly use — but we must also find ways to convince ourselves, our colleagues, and our institutions of the value that is produced by the use of such systems.

traditional peer review and its defenses

David Shatz notes in the introduction to his 2004 volume on peer review that his text is not only "the first book-length study of peer review that utilizes methods and resources of contemporary

8 Janice Radway has argued that the rise of professionalization in the academy "had everything to do with specialization, with the growing emphasis on laboratory research, and with the creation of a communications infrastructure that enabled the publication, circulation, and discussion of research results not only among peers but within a larger society called upon to finance such research, to support it with students, and to understand its value" (217), thus reminding scholars that our very professional existences (and the support that we need in order to maintain such existences) may be dependent not just upon communication amongst ourselves, but on the inclusion of a broader public in that communication, such that they understand the value of academic ways of knowing.

philosophy," but also "the first wide-ranging treatment of the subject by a scholar in the humanities," a fact that becomes all the more surprising when he points out:

> Besides its ethical aspect, the topic also has dimensions of epistemological significance, since it implicates such concepts as truth, bias, relativism, conservatism, consensus, and standards of good argument. Philosophers and other humanities scholars have produced a voluminous literature on these subjects. Yet they have not applied their approaches to these topics to peer review itself, that is, to the very procedures and practices that produced much of the voluminous literature in ethics, epistemology, and so many other fields. (Shatz 4)

Shatz indicates a number of reasons why this may be so, including that the more nebulous (or, rather, problematized) understanding of "truth" in the humanities precludes such scholars from being able to "*show* that a peer review was wrong" (Shatz 6), and that a critical study of peer review might require empirical work of a sort for which humanists are neither trained nor rewarded. Beyond these factors, however, I'd argue that a critical study of the epistemological practices of peer review requires a form of self-analysis that, as Donald Hall has argued in *The Academic Self*, many of us resist. Such resistance might suggest an underlying anxiety about the outcome of the analysis, a concern that the time-honored procedures and standards of the humanities might be shown to be flawed — and thus that the work that has developed through those procedures and according to those standards might be even further marginalized within the academy's mission of knowledge-production. However, as Hall argues, genuinely "owning" our careers and the ways in which we conduct them requires taking the risk of applying our critical skills to an examination of "the textuality of our own profession, its scripts, values, biases, and behavioral norms" (Hall xiv). In the academy — as goes the joke about defenders of tradition in many realms — too many attitudes may be summed up in a mere eight words: "we have never done it that way before." The apparently intractable nature of the way things have always been done is precisely the kind of signal that, in other institutions, impels us to critical analysis; a refusal to turn such a critical eye on our own seemingly naturalized assumptions may create (or deepen) an

atmosphere of intellectual oppression and stultification, as we allow systems in which we do not genuinely have faith to dictate our engagements with the world, and with one another. Opening up the basis of those engagements through a thorough reconsideration of peer review may be precisely what we need in order to allow our work to help shape the ways of knowing of the contemporary world.

Resistance to considering the merits of a more open mode of publishing often runs something like that expressed — in, I assume, an intentionally hyperbolic fashion — by Shatz:

> It is hard to say who would have the biggest nightmare were open review implemented: readers who have to trek through enormous amounts of junk before finding articles they find rewarding; serious scholars who have to live with the depressing knowledge that flat earth theories now can be said to enjoy 'scholarly support'; or a public that finds the medical literature flooded with voodoo and quackery. Let us not forget, either, that editors and sponsoring universities would lose power and prestige even while their workload as judges would be eliminated. (16)

The vehemence of such resistance reveals something about the nervousness of those who express it, and as in much psychotherapeutic discourse, it's only after some initial projection and displacement that the real source of that anxiety comes out: the loss of "power and prestige."[9] However, in responding to those earlier displacements of anxiety, one can provide certain kinds of reassurance. First, the computer technologies that make open review possible also make possible the implementation of analytical tools that can help filter "rewarding" articles from any "junk" they may be mired in, whether those tools employ the results of the open review system themselves or use other modes of sophisticated textual analysis and recommendation. Second, serious scholars depressed by the apparent anything-goes nature of open publishing can see to it, by participating in a system of open review, that "flat earth" theories

9 See also Harnad, who presents many of the same concerns: "Every editor of a learned journal, commentary journal or not, is in a position to sample what the literature would have looked like if everything had appeared without review. Not only would a vanity press of raw manuscripts be unnavigable, but the brave souls who had nothing better to do than to sift through all that chaff and post their commentaries to guide us would be the last ones to trust for calibrating one's finite reading time" (291). The implication, of course, is that without the power to determine whether a manuscript can be published or not, the prestige will drain out of the reviewing process, leaving scholars with only the opinions expressed by the *hoi polloi*.

obtain the reception that they deserve. Third, the public is already flooded with voodoo and quackery, a state easily revealed through the most cursory look at the relationship between the pharmaceutical industry and the bulk of publicly available medical information, or the dramatic shifts in direction of nutritional advice over the last ten years.[10] But finally, and most importantly, one must acknowledge that if the loss of power and prestige are our primary concerns in clinging to closed review, we would be best served by admitting this to ourselves up front. If, however, we enjoy the privileges that obtain from upholding a closed system of discourse sufficiently that we're unwilling to subject it to critical scrutiny, we may also need to accept the fact that the mainstream of public intellectual life will continue, in the main, to ignore our work. This can, of course, be rationalized as the inevitable, unenviable fate of genius in a world of mediocrity.

the history of peer review

It would be worthwhile, however, to explore several of the assumptions we make about the benefits of peer review, to avoid clinging to our present ways of working out of the mistaken sense that, as they have ever been thus, so they should remain. In fact, peer review as we currently know it has a different history than we might assume. Very little in the way of investigation of the historical development of peer review has been done, and the few explorations of peer review that attempt to present some sense of the system's history by and large cite the same few, brief texts.[11] Moreover, nearly all of the texts exploring the history of peer review focus on the natural and social sciences, and almost none of them mention peer review in scholarly book publishing.[12] It is, unfortunately perhaps, beyond

10 One might in particular note Roy and Ashburn, who indicate that it was not in spite of but rather *due to* the peer review process that published studies of the anti-inflammatory drugs Celebra and Vioxx excluded data about those drugs' potential for causing heart damage.

11 See, for instance, Fabiato, Meyers, Rennie, and Spier, among others, all of whom draw heavily from David Kronick's two-page "Peer Review in 18th-Century Scientific Journalism."

12 Given the overwhelming focus on the institution of peer review in the sciences, it might be worth asking whether the

the scope of this chapter to fill in all of those gaps, but it is worth noting a few wrinkles in the history of peer review as it is conventionally understood. Most often, authors date the advent of the thing we're talking about when we refer to editorial peer review today — the assessment of manuscripts by more than one qualified reader, usually not including the editor of a journal or press — to the 1752 Royal Society of London's creation of a "Committee on Papers" to oversee the review and selection of texts for publication in its nearly century-old journal, *Philosophical Transactions*.[13] A number of authors complicate this history by pointing to the existence of at least one earlier instance of formalized peer review in a scientific journal: the Royal Society of Edinburgh seems to have had such a system in place as early as 1731.[14]

However, Mario Biagioli argues in "From Book Censorship to Academic Peer Review" that a deeper excavation of the genealogy of peer review suggests that its origins may, significantly, lie in seventeenth-century academic book publishing, and that peer review of journal articles formed a much later stage in the process's development. Biagioli ties the establishment of editorial peer review to the royal license that was required for the legal sale of printed texts; this mode of state censorship, employed to prevent sedition or heresy, was delegated to the royal academies through the imprimatur granted them at the time of their founding. The Royal Society of London, for instance, took on that imprimatur by passing a resolution in December 1663, one year after its founding, which stated that "No book be printed by order of the council, which hath not been perused and considered by two of the council, who

---

adoption by the humanities of this same mode of assessment is further evidence of the desire to transform our fields into "human sciences," no doubt as a defense against claims — put forward with the greatest impact in university budgets — that our work is insufficiently rigorous and serious to be considered "research."

13 Prior to the establishment of this committee, the selection of manuscripts was in the sole hands of the society's secretary; this moment of transition was important both in the history of the society and of academic publishing, as this was the first time that the society made a public claim of its affiliation with and responsibility for the journal.

14 See Kronick, "Peer Review." However, note that Kronick indicates in *Devant le Deluge* that more than 20 percent of the attributed papers published in the journal while Alexander Monro was editor (1731 forward) were in fact written by Monro himself (181).

shall report, that such book contains nothing but what is suitable to the design and work of the society" (qtd. in Biagioli 21). The purpose of such review, as Biagioli emphasizes, is more related to censorship than to quality control: "As in traditional book licensing, the review was about making sure that a text did not make unacceptable claims rather than to certify that it made good claims" (Biagioli 23). Because the members of the royal academies were, if not literally part of the government, certainly dependent upon the state for their livelihoods, the concept of "peer review" in this instance indicates an early ambiguity between review by one's peers and review by a peer of the realm; as Biagioli suggests, "because of the 'pre-disciplining' of academicians, the simple requirement that manuscripts had to be reviewed by the whole academy or by a committee made it almost impossible that anything controversial would go to press" (15). Gradually, however, scholarly societies facilitated a transition in scientific peer review from state censorship to self-policing, introducing a degree of autonomy but simultaneously becoming, in the Foucauldian sense, a disciplinary technology, creating the conditions of possibility for the academic disciplines that it authorizes. Though peer review may have shed "its negative symbolic connections to early modern absolutism," as Biagioli concludes, and instead become "the new symbol of the relationship between science and liberal societies," its roots in early modern book censorship are revealed by its continued appeal to the imprimatur it grants, "though one that is now about technical accuracy, not legal approbation" (32, 34).

Peer review thus long pre-dates the invention of the scholarly journal, originating with the formation of the royal academies themselves.[15] Membership in these societies required scientists to demonstrate their *bona fides* in the form of publication, experimentation, or invention in order to be eligible for election — arguably subjecting their work to a form of peer review.[16] Further, early scientists

15 Spier argues, in fact, that the first known description of a peer review process may be found in a late 9th-century medical text.

16 See Kronick, *Devant le Deluge*, 96.

circulated letters amongst their peers, or read papers in society meetings, thus reporting the results of their investigations with the explicit intention of eliciting response.[17] The application of peer review processes to scientific journal publishing thus becomes a further extension of society business — reviewing and discussing the reports of work done by the society's membership. Moreover, Drummond Rennie argues that early journal peer review processes were less focused on quality control than we would now assume:

> systems of peer review, internal and external to journals, were put in place by editors during the eighteenth century in order to assist editors in the selection of manuscripts for publication. It was appreciated from the start that the peer review process could not authenticate or endorse because the editors and reviewers could not be at the scene of any crime... the journals from the beginning threw the ultimate responsibility for the integrity of the article squarely upon the author ("Editorial Peer Review" 2).

Early peer review in scientific journal publishing was meant to augment editorial expertise rather than to exercise more conventionally understood modes of quality control. Moreover, as Jean-Claude Guédon and Raymond Siemens indicate, while peer review developed in order to augment the expertise of the editor, the process "nevertheless rested on procedures that put the editor-in-chief in absolute control, albeit in an acceptable way" (18), via editorial control over the selection of reviewers. Thus, while we attribute the arbitration of value in scholarly publishing to the review process to which work has been subjected, that review process cannot guarantee the quality of the publications that appear through it, nor does it wholly diffuse the authority of the editor.

On the one hand, peer review has its deep origins in state censorship, as developed through the establishment and membership practices of state-supported academies, and, on the other, peer review was intended to augment the authority of a journal's editor rather than assure the quality of a journal's

17 See Kronick, *Devant le Deluge*, in which he suggests that the letter "represented a form in which a scientific article could be disseminated for comment and may be considered equivalent to reading a paper at a scientific meeting before submitting it to a publisher or editor for peer review" (268).

products. Given those two disruptions in our contemporary notions about the purposes of peer review, it may be less surprising to find that the mode of formalized review that we now value in the academy seems not to have become a universal part of the scientific method, and thus of the scholarly publishing process, until as late as the middle of the twentieth century; *Science* and *The Journal of the American Medical Association*, for instance, did not vet manuscripts through outside reviewers until the 1940s.[18] The history of peer review thus appears to have been both longer and shorter than we may realize. And yet, because of the role that it has played in authorizing academic research — because we ourselves, as Biagioli suggests, are both the subject and the object of its disciplining gestures — it has become so intractably established that we have a hard time imagining not just a future without it, but any way that it could conceivably change.

the future of peer review

The issue of peer review's future has nonetheless been taken up in various forms by a number of recent publishing experiments. One such experiment is arXiv, an open-access "e-print" (or pre-print) repository, founded at Los Alamos and now housed at Cornell University, through which scientists have increasingly disseminated and obtained working papers in physics, mathematics, computer science, and quantitative biology. Such papers are very often submitted to arXiv before they are submitted to journals — sometimes because the authors want feedback, and sometimes simply to get an idea out into circulation as quickly as possible. However, a growing number of influential papers have *only* been published on the arXiv server, and some have suggested that arXiv has in effect replaced journal publication as the primary mode of scholarly communication within certain specialties in physics. As

18 See Burnham and Spier; see also Weller 3-6 for a suggestive list of scientific journals and the moments and modes in which they adopted editorial peer review.

Paul Ginsparg indicates, arXiv has had great success as a scholarly resource despite employing only a modicum of review:

> From the outset, a variety of heuristic screening mechanisms have been in place to ensure insofar as possible that submissions are at least *of refereeable quality*. That means they satisfy the minimal criterion that they would not be peremptorily rejected by any competent journal editor as nutty, offensive, or otherwise manifestly inappropriate, and would instead at least in principle be suitable for review (i.e., without the risk of alienating or wasting the time of a referee, that essential unaccounted resource). These mechanisms are an important — if not essential — component of why readers find the site so useful: though the most recently submitted articles have not yet necessarily undergone formal review, the vast majority of the articles can, would, or do eventually satisfy editorial requirements somewhere. (Ginsparg 12, emphasis in original)

In 2004, however, arXiv added a layer of author verification to its system by implementing an endorsement process that requires new authors to be vouched for by established authors before submitting their first paper to any subject area on the site. The site is at great pains to indicate that the endorsement process "is not peer review," but it is a process for the review of peers, and as such bears a direct relationship to the site administrators' desire to maintain the consistently high quality of submissions to the site, a means of verifying that "arXiv contributors belong [to] the scientific community" (http://arxiv.org/help/endorsement). The site administrators do note, however, that "Endorsement is a necessary but not sufficient condition to have papers accepted in arXiv; arXiv reserves the right to reject or reclassify any submission," suggesting that the open server is nonetheless subject to a degree of editorial control, if not in the form of traditional peer review.

The peer review experiment in scientific publishing that has recently received the most attention, however, is that undertaken in 2006 by *Nature*, which was accompanied by a debate, published on the journal's website, about the future of peer review. The experiment was fairly simple: the editors of *Nature* created an online open review system that ran parallel to its traditional anonymous review process. "From 5 June 2006," the editors wrote, "authors may opt to have their submitted manuscripts posted publicly for comment. Any scientist may then post comments, provided they

identify themselves. Once the usual confidential peer review process is complete, the public 'open peer review' process will be closed. Editors will then read all comments on the manuscript and invite authors to respond. At the end of the process, as part of the trial, editors will assess the value of the public comments" (Campbell). The experiment was closed in early December, after which time the editors analyzed the data resulting from it, and, later in the month, declared the experiment to have failed, announcing that "for now at least, we will not implement open peer review." The statistics cited by the editors are indeed indicative of serious issues in the open system they implemented: only 5% of authors who submitted work during the trial agreed to have their papers opened to public comment; of those papers, only 54% (or 38 out of a total of 71) received substantive comments. And as Linda Miller, the executive editor of *Nature*, told a reporter for *Science News*, the comments that the articles received weren't as thorough as the official reviews: "They're generally not the kind of comments that editors can make a decision on" (Brownlee 393).

Certain aspects of the experiment, however, raise the question of whether the test was rigged from the beginning, destined for a predictable failure because of the trial's constraints. First, no real impetus was created for authors to open their papers to public review; in fact, the open portion of the peer review process was wholly optional, and had no bearing whatsoever on the editors' decision to publish any given paper. Which points to the second problem, as no incentive was created for commenters to participate in the process: why go to all the effort of reading and commenting on a paper if your comments serve no identifiable purpose?

As several entries in the web debate held alongside *Nature*'s peer review trial made clear, though, the editors had not chosen a groundbreaking model; the editors of several other scientific journals that already use open review systems to varying extents posted brief comments about their processes. *Electronic Transactions on Artificial Intelligence*, for instance, has a two-stage process, a three-month

open review stage, followed by a speedy up-or-down refereeing stage (with some time for revisions, if desired, inbetween). This process, the editors acknowledge, has produced some complications in the notion of "publication," as the texts in the open review stage are already freely available online; in some sense, the journal itself has become a vehicle for re-publishing selected articles.

ETAI's dual-stage process highlights a bifurcation in the purpose of peer review: first, fostering discussion and feedback amongst scholars, with the aim of strengthening the work that they produce; second, providing a mechanism through which that work may be filtered for quality, such that only the best is selected for final "publication." Moreover, by foregrounding the open stage of peer review — by considering an article "published" during the three months of its open review, but then only "refereed" once anonymous scientists have held their up-or-down vote, a vote that comes only after the article has been read, discussed, and revised — such a dual-stage process promises to return the center of gravity in peer review to communication amongst peers.

ETAI's process thus highlights the relatively conservative move that *Nature* made with its open peer review trial. First, the journal was at great pains to reassure authors and readers that traditional, anonymous peer review would still take place alongside open discussion. There was, moreover, a relative lack of communication between the two forms of review: open review took place at the same time as anonymous review, rather than as a preliminary phase, preventing authors from putting the public comments they received to use in revision. And though the open review was on some level expected to serve as a parallel to the closed review process — thus Miller's disappointment that the comments weren't as thorough as traditional peer reviews — they weren't really allowed to serve a parallel function: while the editors "read" all such public comments, it was decided from the beginning that only the anonymous reviews would be considered in determining whether any given article was published.

anonymity

Perhaps this caution about open review was an attempt to avoid throwing out the baby of quality control with the bathwater of anonymity. The editors of *Atmospheric Chemistry and Physics*, however, presented evidence (based on their two-stage review process) that open review significantly increases the quality of articles a journal publishes:

> Our statistics confirm that collaborative peer review facilitates and enhances quality assurance. The journal has a relatively low overall rejection rate of less than 20%, but only three years after its launch the ISI journal impact factor ranked *Atmospheric Chemistry and Physics* twelfth out of 169 journals in 'Meteorology and Atmospheric Sciences' and 'Environmental Sciences'.
>
> These numbers support the idea that public peer review and interactive discussion deter authors from submitting low-quality manuscripts, and thus relieve editors and reviewers from spending too much time on deficient submissions. (Koop and Pöschl)[19]

Evidence such as this begins to suggest that traditional closed, anonymous peer review processes and quality control aren't quite as related as we often assume. In fact, it's arguable that the primary results of a closed peer review process are negative. As Fiona Godlee has argued, anonymous review "has the effect of giving reviewers power without responsibility" (65), since reviewers are freed by the veil of anonymity to behave, in some instances, in a variety of unprofessional ways, ranging from the relatively innocuous unleashing of snark on an undeserving target to several utterly unacceptable forms of academic dishonesty. Such behaviors are of course not the norm, but they occur frequently enough that they should give us pause.[20] On the other side of the review process, of course, are the authors, ostensibly equal participants in a conversation about their work. The anonymous peer review process, however, effectively closes the author out of the conversation, which instead becomes a backchannel discussion between the reviewer and the editor. As such, the author is hindered in her ability to learn from the review process *even if she is given a copy of the reviewer's comments*, as there is no forum in which

19 See also Pöschl, "Interactive Journal Concept."

20 See Godlee 66-67 for an account of one instance of reviewer fraud. See also Campanario.

she can respond to those comments in kind. By the time the comments arrive, generally speaking, the decision about the manuscript's fate has been decided, the conversation is over, and the author is too often left with no one listening.[21]

Reviewer anonymity, however, has long enough been a part of the process that many senior academics express alarm at the thought of that protection being removed, insisting that their anonymity as reviewers is necessary in order for them to have the freedom to say that a manuscript should not be published. Such a position would certainly be justifiable if the primary purpose of peer review is quality control, and if it can be demonstrated that the process is both scrupulous and effective. However, as Douglas Peters and Stephen Ceci famously uncovered in their 1982 article, "Peer Review Practices of Psychological Journals," reviewer reliability is not at all a given. In their experiment, Peters and Ceci selected one article from each of 12 journals in the field, published between 18 and 32 months previously, and resubmitted the article to the same journal, with some minor modifications: they changed the authors' names (but, significantly, not their sexes); they created new institutional affiliations for their authors (notably replacing "high-status" institutions with low- or no-status ones); and they slightly altered the phrasing, but not the meaning, of the articles' opening paragraphs. Only three of these 12 articles were discovered by either the editors or the reviewers to have been previously published, and of the 9 that went undiscovered, 8 were rejected, most on the grounds of "serious methodological flaws" (Peters and Ceci 202).[22] Their conclusion is that one of two things has occurred:

21 An exception to this state of affairs generally appears in scholarly book publishing, though only if the editor has decided based upon the reviews to take the manuscript to the press's editorial board for approval. In that case, the author's response to the reviews is requested; however, this response is generally directed not to the readers, but to the board, further complicating the flow of conversation.

22 Peters and Ceci specifically rule out the possibility that reviewers felt the work to be somehow redundant with the existing literature, even if they couldn't recall the exact source, as no indication of such appears in the reviewers' reports.

either the initial reviewers who approved the articles as originally published were incompetent — which seems unlikely — or "systematic bias was operating to produce the discrepant reviews" (202).[23]

One of the correctives that has been suggested in response to evidence of such "systematic bias," as well as some of the other more egregious abuses of peer review, is a further layer of anonymity: blind review, in which the identity of the author is cloaked, as well as that of the reviewer. Blind review is imagined by many (including Shatz) to be a mode of avoiding certain forms of reviewer bias, for instance, preventing the continuation of an "old boys network" that excludes the work of women, or ensuring that personal grudges play no role in the review process. However, the effectiveness of blind review in genuinely masking authorial identity has been subject to some critical scrutiny, by authors who suggest, for instance, that blinding is futile: "Alas, anyone capable of evaluating research in a given specialty generally knows that specialty sufficiently to identify the probable author of the manuscript under review" (Guédon and Siemens 18).[24] In many cases, in fact, the author has previously presented and discussed the material in public, whether via informal networks or in more formal conference settings. Moreover, blind review can only correct for *ad hominem* bias on the part of reviewers, and cannot compensate for the reviewer who operates within a cloud of intellectual bias, dismissing any arguments or conclusions that disagree with his or her own.[25]

It's also necessary to note that neither reviewer nor author identity are hidden from the editor, who may have his or her own biases; as Godlee notes, "[e]vidence suggests that editors may be susceptible to the pull of prestige," citing the results presented by Zuckerman and Merton, which suggest

---

23 Godlee suggests thirty years later that science has, since the time of Peters and Ceci's experiment, become "less clubby and more competitive" (72), while nonetheless indicating that reviewer bias with respect to the institutional prestige of an author remains operative.

24 See also Godlee 74-75, and Zuckerman and Merton 86.

25 See, for instance, Blair, et al, for an exploration of one remarkable instance of such intellectual or ideological bias amongst blind reviewers.

that "if a paper had higher-ranking authors, editors were more likely to come to a decision without sending it out for peer review" (Godlee 73). Moreover, the editor's selection of reviewers for a manuscript may be influenced by the author's identity, and the editor's evaluation of the reviewers' reports may similarly be affected by the differing levels of prestige of reviewer and author.

Finally, one cannot help but wonder about the logic of correcting for the abuses of anonymity on one side of a conversation by establishing anonymity on the other, creating further barriers between peers rather than encouraging open, effective, productive discussion of intellectual issues. As Drummond Rennie argued in 1994, "We have an ample history to tell us that justice is ill served by secrecy. And so it is with peer review. Two or three hundred years ago, scientific papers and letters were often anonymous. We now regard that as quaint and primitive. I hope that in 20 years, that's exactly how we will look on our present system of peer review" (Commentary 1143).

credentialing

If closed peer review processes aren't serving scholars in their need for feedback and discussion, and if they can't be wholly relied upon for their quality-control functions — if they appear, at least to some, "quaint and primitive" — why do we cling so ferociously to them? Arguably, the primary purpose that anonymous peer review actually serves today, at least in the humanities,[26] is that of institutional warranting, of conveying to college and university administrations that the work their employees are doing is appropriate and well-thought-of in its field, and thus that these employees are deserving of ongoing appointments, tenure, promotions, raises, and so forth. As Rennie has noted, "editorial peer review is seen by investigators and research institutions as a convenient quality control mechanism, for

26 A qualifier that points to a need for much further exploration of the different requirements with respect to peer review in the different disciplines. A study conducted by Zuckerman and Merton in 1971 investigated the differing outcomes of peer review across disciplines, noting that the rejection rate in the humanities was far higher than that in the social or natural sciences; this outcome might suggest the importance for the humanities of a continuing review of its communication and credentialing processes.

which they usually do not have to pay" ("Editorial Peer Review" 10). This mechanism, on the level of the academic book, has been described by Lindsay Waters as a means for departments to "outsource" the evaluation of junior scholars to university presses; the existence of a book by a reputable press comes to serve as a convenient binary signifier of the quality of that scholar's work.[27]

To some degree, one must wonder whether using the results of peer review as a shortcut in faculty performance evaluations isn't misguided in and of itself; much of the most important work published by scholars today is already published in forms that aren't subject to conventionally-understood modes of peer review, such as edited volumes. Moreover, understanding a successful navigation of peer review as a sufficient sign of quality work is a category error of sorts. As Paul Ginsparg has argued, the mere existence of an author's peer reviewed publication is insufficient evidence, for hiring and promotion purposes, of the scholar's level of accomplishment; "otherwise there would be no need to supplement the publication record with detailed letters of recommendation and other measures of importance and influence. On the other hand, the detailed letters and citation analyses would be sufficient for the above purposes, even if applied to a literature that had not undergone that systematic first editorial pass through a peer review system" (9). In other words, our institutional misunderstanding of peer review as a necessary prior indicator of "quality," rather than as one means among many of assessing quality, dooms us to misunderstand the ways that scholars establish and maintain their reputations within the field.

Another obvious question to ask is whether peer review as it is currently practiced is really able to support credentialing in the ways we assume. It's at least imaginable, if as yet untested, that the

27 This concern about the shift in responsibility for reviewing the work of younger scholars is echoed in the final report of the MLA Task Force on Evaluating Scholarship for Tenure and Promotion, which, while at pains to dissociate the reliance on press judgments from peer review itself, nonetheless acknowledges that "this apparatus of external peer review also created the conditions whereby individual departments can practically abdicate their responsibility to review the scholarly work of the very colleagues they have appointed to tenure-track positions" (56).

intellectual purposes that we expect of peer review — most importantly, quality control — could be undermined by this functionalist use of the process's results, as some extremely well-meaning reviewers, all too aware of the stakes of their evaluations, could unconsciously tend toward a sort of scholarly grade-inflation. And many scholars work with a sense, however vague, that certain publications use peer review as a means of supporting pre-determined ideas about a field's in-group, suggesting that the credentialing cart may have been put before the peer-review horse.

The internet, as Guédon and Siemens indicate in one of this article's epigraphs, has in any event disrupted our ability to draw an association between the fact that a scholarly text has been published and the quality of work it may therefore contain. The result, conventionally, has been the dismissal by many faculty and administrators of all electronically published texts as inferior to those that appear in print — and thus the reinforcement of "the way things have always been done" at the expense of experimental modes that might produce new possibilities. Such conservatism shouldn't come as much surprise, of course; those faculty and administrators who are in the position of performing assessments of the careers of other, usually younger, faculty are of necessity those who have sufficiently benefitted from the current credentialing system as to rise to that position. As Guédon and Siemens suggest, those who hold such privilege will find ways to keep it, preferably without drawing attention to their having done so, precisely by making a virtue — and a besieged one, at that — out of the status quo.

But I want to emphasize at this point that, while I have spent a great deal of time in this chapter on the various abuses and shortcomings of the peer review process as currently constituted, the core notion behind it — that one's work as a scholar should be reviewed and assessed by one's peers — is of course a good one; the problem is in the implementation of that notion as an exercise in gatekeeping, and its subsequent transformation into a means of creating authority in and of itself. Those two shifts have the potential not only to interfere with peers' ability to communicate directly and fully with one

another, but they also create enormous amounts of extra, unproductive work for everyone involved. Scholars pour hours upon hours into peer review each year, time which is not only usually uncompensated but which also results in a product for which reviewers can receive no "credit," as peer reviews, unlike post-publication reviews, cannot ever themselves be counted among the reviewer's published work. For all of these reasons, I want to suggest that the time has come for us to consider whether, really, we might all be better served by separating the question of credentialing from the publishing process, by allowing everything through the gate, and by designing a post-publication peer review process that focuses on how a scholarly text should be received rather than whether it should be out there in the first place. What if peer review learned from social software systems such as Slashdot and Digg, and became peer-to-peer review?

the reputation economy

The notion of "peer-to-peer review" that I have been circulating for the last couple of years draws upon the convergence of the kinds of discussion many scholars would like peer review to produce and the decentralized peer-to-peer networks that have sprung up across the Internet. In fact, just as Biagioli suggested a shift, across the early modern development of the scientific academy, in the definition of the term "peer" — from a member of the royal court to a scholarly colleague — so Chris Anderson has argued that the term is once again being redefined in online communities: "In the Internet age, 'peer' is coming to mean everyman more than professional of equal rank. Consider the rise of peer-to-peer networks and the user-created content of 'peer production,' such as the millions of blogs that now complement mainstream media" (Anderson). Anderson uses this transformation in the notion of a peer to suggest that the academy might fruitfully find ways to open its review processes to "the wisdom of the crowds," allowing new models of authority in online information distribution to

augment more traditional review systems. For instance, Anderson's reading of Wikipedia contradicts many of the conventional academic assumptions about the project, calling it "not so much anti-elitist as… 'anti-credentialist,'" a distinction that indicates that site editors' "contributions are considered on their merit, regardless of who they are or how they became knowledgeable. If what they write stands up to inspection, it remains; otherwise it goes" (Anderson).[28] Such systems of communal knowledge-production are thus far from the free-for-all that many have assumed — and, in fact, are at least in theory bringing into being a new mode of authority production; those editors whose work consistently "stands up" to community inspection may be accorded a kind of clout within the community that then affects assumptions about their future work.

I say "in theory" because one of the most important criticisms that has been leveled at Wikipedia is its acceptance of anonymous contributions, which hinders the ability of readers to assess particular Wikipedians' work based upon their reputations. Reputation in this sense should be understood as separate from credentials; the point is not whether a particular Wikipedia editor has a degree in the appropriate subject area, but rather whether his or her work on the site has repeatedly stood up to community scrutiny. There is, of course, no small irony in the fact that the academic outcry against the anonymous nature of much of Wikipedia's production occurs at the very same time that we cherish our own anonymity as peer reviewers, and we might take the implications of this contradiction to heart.

In a recent experiment with community-based peer review, Noah Wardrip-Fruin published the manuscript of his book-in-progress, *Expressive Processing*, in a CommentPress-based format on his co-authored blog, *Grand Text Auto*, seeking review from the community of GTxA's readers, at the same

28 Recent controversies between the so-called "deletionist" and "inclusionist" Wikipedians complicate Anderson's model significantly, of course. See "The Battle for Wikipedia's Soul."

time that MIT Press sent the manuscript to traditional anonymous peer reviewers. Though a number of articles, including most notably one in the *Chronicle of Higher Education*, represented this experiment as a "head-to-head" competition between open and closed peer-review systems,[29] Wardrip-Fruin was clear that such a contest was not his goal. The important aspect of the experiment was in getting feedback from a community he trusted:

> In most cases, when I get back the traditional, blind peer review comments on my papers and book proposals and conference submissions, I don't know who to believe. Most issues are only raised by one reviewer. I find myself wondering, "Is this a general issue that I need to fix, or just something that rubbed one particular person the wrong way?"...
>
> But with this blog-based review it's been a quite different experience. This is most clear to me around the discussion of "process intensity" in section 1.2. If I recall correctly, this began with Nick's comment on paragraph 14. Nick would be a perfect candidate for traditional peer review of my manuscript — well-versed in the subject, articulate, and active in many of the same communities I hope will enjoy the book. But faced with just his comment, in anonymous form, I might have made only a small change. The same is true of Barry's comment on the same paragraph, left later the same day. However, once they started the conversation rolling, others agreed with their points and expanded beyond a focus on The Sims — and people also engaged me as I started thinking aloud about how to fix things — and the results made it clear that the larger discussion of process intensity was problematic, not just my treatment of one example. In other words, the blog-based review form not only brings in more voices (which may identify more potential issues), and not only provides some "review of the reviews" (with reviewers weighing in on the issues raised by others), but is also, crucially, a conversation (my proposals for a quick fix to the discussion of one example helped unearth the breadth and seriousness of the larger issues with the section).

In the end, he notes, "the blog commentaries will have been through a social process that, in some ways, will probably make me trust them more" (Wardrip-Fruin). Knowing the reviewers' reputations, and seeing those reputations as part a dynamic process of intellectual interaction, produces the authority of the comments, and will thus affect the authority of the book that Wardrip-Fruin finally publishes.

Given this, we might begin to posit an intimate relationship between reputation and authority in the intellectual sphere. This relationship has of course long existed within the academy, manifested in

29 See Young.

our various mechanisms of assessment and review, but digital networks give us new modes of determining reputation, as well as new requirements for such reputation-determining metrics. Not all networked publishing structures are concerned with reputation, of course: Wikipedia, for instance, only makes tangential use of a reputation-based system in assessing the authority of its entries. Other systems, most notably online retailers such as eBay, rely heavily on customer feedback in evaluating the reliability of service provided by individual merchants within the network. And the news and discussion forum Slashdot, most famously, uses a system of rating contributions to assess the reputations of individual contributors.

The Slashdot system evolved out of a more traditional system of comment moderation, in which 25 people weeded out the nonsense and highlighted the valuable; when the work became too much for those moderators, they selected 400 more moderators based on the reputations they'd developed as users of the site. However, this hierarchical moderation system, in which some users had power that others didn't, quickly led to abuses, and the site's owners began developing what they refer to as a system of "mass moderation." In this system, nearly every active contributor to the site has the potential to receive, for a period of time, a degree of power to rate the site's contributions, through being given a number of "points of influence"; each time the contributor rates a comment on the site, he expends one influence point.[30] These influence points expire rapidly if unused, and contributors cannot rank comments in threads in which they actively participate, thus preventing influence from becoming a currency within the system, and preventing moderators from controlling the discourse. The power to moderate, moreover, is only granted by the system based upon the contributor's "karma" within the site, based

30 This footnote should be read as an acknowledgment of the issues that my use of "he" in describing this system highlights; Slashdot is felt by many to be a highly male-dominated, if not downright misogynist, environment. One might, for instance, see what happens in the comment thread when a poster asks for advice on handling being the lone woman working for an IT firm: <http://ask.slashdot.org/article.pl?sid=06/08/31/1755259&threshold=-1>.

upon the ways that the contributor's own comments have been moderated, which is understood to be a community-based assessment of whether or not the contributor's comments have been a helpful, positive addition to the community.[31]

There are, of course, weaknesses in a reputation system such as Slashdot's, in which the value of a user's contributions to the community can become subject to manipulation and attack, potentially replacing substantive discourse and engagement with a networked popularity contest. As noted by one user of Advogato, an online forum for free and open-source software developers,

> If you believe that "in any sufficiently large crowd, the majority are idiots," then this can be applied to Slashdot moderators too. All moderators have equal powers and the system is supposed to work as a kind of democracy. But if the majority does not think very much about what they are doing (because of lack of time, lack of interest, lack of intelligence or many other reasons), then it becomes easy to abuse the system.... I hope that something similar to the trust metrics implemented on Advogato could help. (Quinet)

Advogato's "trust metrics" are intensively computational, evaluating each "node," or user, within the network via its interconnections with the network's many other nodes, certifying each node through three levels of trust (apprentice, journeyer, and master). One of the benefits of this system, as its developer writes, is its "resistance to catastrophic failure in the face of a sufficiently massive attack" (Levien). Reputation, in this implementation, cannot be hacked; on the other hand, it is entirely objectively calculated, leaving little to no room for subjective evaluation.

However, while such "trust metrics" might seem inappropriate as a model for reconsidering peer review, they may nonetheless help point us in the direction of a more sophisticated, partially computational, partially review-based system for determining authority in networked scholarly publishing, the kind of model Michael Jensen imagines under the rubric of "Authority 3.0." Such a system, whatever its particulars, must operate in accordance with three key principles. The first is that it

31 See Slashdot's FAQ: <http://slashdot.org/faq/com-mod.shtml#cm520>.

must be as non-manipulable as possible, preventing the importation of in-group favoritism, logrolling, and other interpersonal abuses from traditional peer review into the new system. Second, the system must achieve a critical mass of participation, and thus will need to operate within an ethos of "quid pro quo"; in contrast with Slashdot's system, in which users earn the right to become reviewers by publishing within the system, scholars must earn the right to publish within these new electronic publishing networks by actively serving as reviewers. And finally, and most significantly: the key activity of such a peer-to-peer review system must be not the review of texts, but the review of the reviewers. It is the reviewers, after all, that a reader within such a network needs to trust, and as Jonathan Schwartz, the COO of Sun Microsystems, has argued in numerous interviews, "trust is the currency of the participation age."[32]

It's no accident, of course, that trust is here defined through an economic metaphor; while the "currency" that reputation affords within the academy is far less spendable than is that within the corporate world, there's nonetheless an economic reality at its root, and thus at the root of the peer-review mechanisms through which reputation is currently granted. Print-based publishing operates within an economics of scarcity, with its systems determined in large part by the fact that there are a limited number of pages, a limited number of journals, a limited number of books that can be produced; the competition among scholars for those limited resources requires pre-publication review, to make sure that the material being published is of sufficient quality as to be worthy of the resources it consumes. Electronic publishing faces no such scarcity; there is no upper limit on the number of pages a manuscript can contain or the number of manuscripts that can be published, or at least none determined by available resources, as the internet operates within an economics of abundance. We

32 Schwartz should certainly understand the value of trust in the digital world, given the need to rebuild its reputation that Sun faced after the dot-com bust; see Falkow. Schwartz's phrase has become the tagline for the Open Media Commons service operated by Sun; see <http:// www.openmediacommons.org>.

might think, for a moment, of Cory Doctorow's "Whuffie," in *Down and Out in the Magic Kingdom*, a currency of sorts that measures the esteem one is held in, a system designed specifically for an economics of abundance.[33] As Doctorow explained in an interview, Whuffie becomes important in the digital sphere precisely because such a sphere "isn't a tragedy of the commons; this is a commons where the sheep s*** grass — where the more you graze, the more commons you get" (SFGate.com). Such is the abundance of the internet, and given this abundance, imposing artificial scarcity through a gatekeeping model of peer review makes little sense.

However, in a self-multiplying scholarly commons, some kind of assessment of the material being published (or having been published) remains important, but not because of scarce resources; instead, what remains scarce are time and attention. For this reason, peer review needs to be put not in the service of gatekeeping (determining what should be published for any scholar to see) but of filtering (determining what of the vast amount of material that has been published is of interest or value to a particular scholar). In using a computational filtering system, of course, the most important thing to understand is its algorithm — what criteria, in what balance, it's using in making decisions to include or exclude various pieces of data.[34] Similarly, in using a human filtering system, the most important thing to have information about is not the data that is being filtered, but the human filter itself. Thus, in a peer-to-peer review system, the critical activity is not the review of the texts being published, but the review of the reviewers.

---

33 Notably, John Holbo appropriated the term in a blog post speculating on a cooperative electronic publishing model, entitled "Will Work for Whuffie."

34 See, for instance, discussions of Google's PageRank algorithm, which arguably measures popularity of pages through an analysis of inbound links (Regalado), but which others interpret as "inherently conservative," granting further authority to the already popular (Vaidhyanathan, "Where Is This Book Going?").

community-based filtering

One might see a relatively simple example of such a system in Philica, which bills itself as "the journal of everything." Philica is an open publishing network, co-founded by British psychologists Ian Walker and Nigel Holt, which invites scholars from any field to post papers, which are then made freely available for reading and review by any interested user. Philica describes itself as operating "like eBay for academics. When somebody reviews your article, the impact of that review depends on the reviewer's own reviews. This means that the opinion of somebody whose work is highly regarded carries more weight than the opinion of somebody whose work is rated poorly" (Philica Tutorial).[35] Account registration is open, though members are asked to declare their institutional affiliations if they have them, and encouraged to obtain "confirmation" of their status within the academy by sending the site administrators a letter on institutional letterhead, or a letter detailing appropriate credentials as an independent researcher. The site's FAQ indicates that membership is in theory restricted to "fully-qualified academics," though without confirmation, one could simply claim such a status, and thus the system makes an unconfirmed membership "much less useful than a confirmed membership, since (a) unconfirmed members' reviews carry less weight than confirmed members' reviews and (b) readers are less likely to trust research from unconfirmed authors. In other words, there's not really much point joining if you do not go on to prove your status" (Philica FAQ). Reviewing articles published on Philica is open to registered, logged-in members, whether "confirmed" or not, though confirmed members' reviews are noted with a check mark. Articles are evaluated by reviewers both quantitatively (rating "originality," "importance," and "overall quality" on a 1-to-7 scale) and qualitatively, via comments.

35 The choice of eBay is perhaps a bit unfortunate, resulting in faintly crass images of intellectual commerce, but there's something apt in the comparison as well, suggesting that electronic scholarly publishing might function as a marketplace of ideas in which producers and consumers can find one another without the need for an intermediary. Lindsay Waters, however, argues that the marketplace "is not a concept that should be considered the ultimate framework for the free play of ideas" (9). See also Shatz for a more elaborated argument against the marketplace metaphor.

Article authors each have a page that details their work on the site, including the number of articles and notes that they have published, the mean peer-review ratings their work has received, and the number of reviews and comments that the author has contributed to other work. The site notes that the author's ratings "will change whenever a new review of this author's work appears, as well as whenever somebody reviews the work of anybody who has reviewed" the work of the author in question.[36]

While Philica's system presents some compelling possibilities for the future of scholarly publishing, it nonetheless has a number of apparent shortcomings: though the articles uploaded to the site are reviewed, and reviews are valued based on the assessed quality of the work of the reviewers, the quality of the reviews themselves isn't assessed, and thus these reviews don't count among the "work" used in determining the value of a reviewer's comments. In part this is due to the fact that while the comments made by a particular reviewer are associated with one another, they are not associated with their authors by name, but are rather submitted anonymously. Each review entry page contains the following notice: "Unless you sign your review, which you are welcome to do if you wish, it will be anonymous to the author and to other Philica readers. Nevertheless, the administrators can see who you are if necessary so please be sure your review is not abusive" (Philica). Thus, Philica only opens the comments produced by peer review to public scrutiny; though reviewers are accountable to the site's administrators, they are not directly accountable to the article's authors, or to the network's community as a whole. And while the reviewers' own peer-review ratings affect the way the system weights the ratings they assign to others, the working of this algorithm remains partially hidden behind the veil of anonymity.

Further, as a "journal of everything," Philica runs the risk of precisely the kind of overflow that makes Internet skeptics worry; if "everything" is published there, how will researchers find what they

36 See, for instance, <http://philica.com/user_details.php?user_code=1>.

need — and will they, as Shatz suggested, be required to "trek through enormous amounts of junk before finding articles" that are at all "rewarding" (16)? Such concerns are well-founded, in this case, as work published on Philica is organized by discipline, but only 27 such disciplinary categories currently exist on the site, with no further subdivisions, tags, or other metadata allowing the reader to find relevant material. The site thus suffers from a too-general mode of organization; the "humanities" as a whole, for instance, represents a single field on Philica. The result, however, has not been overflow but, if anything, underflow; only 124 articles or notes were published on Philica between March 2006 and March 2008, a mere 4 of which were in the humanities. Such a miniscule rate of participation, like that experienced in the *Nature* open review trial, could be taken to indicate a general resistance among academics to new publishing models — and yet, it's hard to imagine that a traditional, closed review, print-based "journal of everything" would fare much better. The purpose of scholarly publishing, after all, is not merely making the results of research public, but making those results public *to the appropriate community*. Because Philica has no particular disciplinary focus, it hasn't been able to build a community.

The development and maintenance of such a community is key to the scholarly publishing network of the future, and in particular to its implementation of peer-to-peer review, because while the post-publication filtering mechanisms that such a system will require may in part be computational, they cannot be wholly automated; the individual intelligences and interests of the members of this social network are the bedrock of community-based filtering. One might, for instance, look at Chris Anderson's explanation for the success of MySpace as a promoter of exclusively "Long Tail" music, where other such networks like MP3.com had failed: "The answer at this point appears to be that it is a very effective combination of community and content. The strong social ties between the tens of millions of fans there help guide them to obscure music that they wouldn't otherwise find, while the content gives

them a reason to keep visiting" (*Long Tail* 149). The absence of the kind of community that MySpace fosters — a user base committed to the site as a means of self-expression, whose relationships with one another are built precisely around that self-expression — prevented MP3.com from becoming a flourishing site for the exploration of new and obscure music, precisely because the absence of social ties among users left them no way of assessing the recommendations others were making. And the more niche-based the mode of cultural communication becomes — the further down the "tail" that communication moves — the more important such community-based knowledge becomes.

Given the case of Philica, in fact, one might begin to speculate that, in electronic scholarly publishing, the community is necessary not just to the post-publication review and filtering process but to the production of content itself. Scholarly communication, generally speaking, is *all* tail, aimed at a comparatively small niche group of similarly focused readers; for that reason, the technologies of the internet seem particularly well-positioned to enable those readers both to find and communicate with one another, as well as to set community-based standards for the evaluation of their work. Only once it is clear to scholars that the standards of this community are their standards — that this is a community to which they belong — will many of them venture to contribute their work to it. In order for such community to be established, however, its individual members must know one another, at least by reputation, and thus the process of review — the setting of standards by the community — must itself be open to continual review.

It seems self-evident: the more open such systems are, the more debate they foster, and the more communal value is placed on participating in them, the better the material they produce can be. However, all of these aspects of the community must be carefully fostered in order for it to avoid turning into what Cass Sunstein describes, in *Infotopia*, as a deliberative cocoon, in which small groups of the like-minded reinforce one another's biases and produce unspoken social pressures toward

conformity with what appears to be majority opinion, resulting in a mode of "group-think" that propagates errors rather than correcting them. Sunstein points out that new internet-based knowledge aggregation systems such as wikis, open source software, and blogs "offer distinct models for how groups, large or small, might gather information and interact on the Internet. They provide important supplements to, or substitutes for, ordinary deliberation" (Sunstein 148), enabling correctives for the errors that small groups of decision-makers can produce. Using such new technologies for purposes of deliberation, however, requires that all members of the network be equally empowered — and in fact, equally compelled — to contribute their ideas and voice their dissent, lest the network fall prey to a new mode of self-reinforcing group-think.

The key to avoiding such group-think, however, is not heightened intellectual individualism — separating oneself from the network — but paradoxically placing the advancement of the community's knowledge ahead of one's own personal advancement. Sunstein presents evidence that the propagation of errors is "far less likely when each individual knows that she has nothing to gain from a correct *individual* decision and everything to gain from a correct *group* decision" (205). Such a turn toward a communally distributed mode of knowledge production, however, will not come easily in a culture in which credentialing processes focus precisely on individual achievement. I'll turn my attention more fully to the issue of collaboration and community in chapter 2, but for now will simply suggest that the success of a community-based review system will hinge on the evaluation of one's contributions to reviewing being considered as important as, if not even more important than, one's own individual projects. Genuine peer-to-peer review will require prioritizing members' work on behalf of the community within the community's reward structures.

mediacommons and peer-to-peer review

This need to focus on the communal aspects of peer-to-peer review — in particular, the review of the reviewers — led me and my colleagues, in the early stages of our planning, to start thinking about MediaCommons as less a digital scholarly *press* than a digital scholarly *network*. Though the social aspects of MediaCommons are not its primary product, we've increasingly come to believe that they're a precondition for the success of the publishing aspects of the network. Too many digital publishing experiments, like Philica, have lagged due to an assumption that might be summed up as "if you build it, they will come." In fact, such publishing experiments would often benefit from examining the relative success of MySpace in comparison with MP3.com, thus placing a greater focus on getting users to come in the first place, on drawing them in by demonstrating the ways that the network's connections will benefit their work. For this reason, the first part of MediaCommons that we are building is the community, in order to create a network of trust between authors and reviewers. For our purposes, a more appropriate analogy between MediaCommons and other "web 2.0" systems, rather than "eBay for academics," might well be "Facebook for scholars," as we are focused on building a network structure that allows people, and not just texts, to interconnect. And the most salient point of that comparison is this: as some scholars have argued, the success of Facebook, as opposed to previous social networking systems such as Friendster and Orkut, derived in no small part from the decision its developers made in keeping the network relatively closed by limiting its use, in its early days, to students at a small number of colleges and universities and by focusing on the connections within those institutions.[37] The emphasis, in other words, was not on allowing users to create new social networks, but rather on helping

37 Numerous pundits insisted that opening Facebook to any user might, in the end, prove to be the service's undoing, though many were primarily lamenting their loss of exclusivity. danah boyd, however, has argued that the success of social networking systems has largely hinged on the ability to control the social context in which one's profile appears. See boyd, "Viewing American Class Distinctions" and "Loss of Context."

them extend their existing offline social networks into digital environments. MediaCommons will thus begin by facilitating the relationships among scholars who are already connected — who attend the same conferences, publish in the same journals, and read one another's work.

This text is being written while the MediaCommons systems are still very much in development, and therefore what follows retains a speculative, hypothetical tone; the actual functions of the network may well wind up being a bit different than what I here project.[38] We are fairly certain, however, that the peer-to-peer system that will be the backbone of MediaCommons will involve the development and implementation of a number of interconnected modules, including:

- a networked user profiling system that will enable scholars to define their interests in taggable, complexly searchable ways;
- a portfolio system that will maintain a comprehensive record of any user's writing within the site, both formal and informal, allowing scholars both to maintain publicly accessible versions of their work and to receive some sort of academic "credit" for the kinds of work — including peer reviews and participation in online forums — that too often remain invisible;
- a sophisticated recommendations system that uses the information in a member's profile, along with robust textual analysis of documents in the network, to present the user with frequently updated suggestions for texts to read, discussions to participate in, and collaborators to work with;
- a "reputation" system that will allow users of the network to review the reviewers, to assess the "value" of a particular scholar's work within the network.

Many systems like these have been developed in isolation from one another, both in open-source and proprietary variants, but they have not as yet been brought together to create such a dynamic community structure, nor have they been put to the uses that scholars might make of them. A social networking system such as Facebook, for instance, allows its users to create profiles and join groups, but

38 As I'll discuss in chapter 2, however, one of the most exciting aspects of a digital publishing environment such as the one the electronic version of this text will be published in is that the text will be updatable in order to reflect MediaCommons's actual state, and yet versionable, to preserve for the historical record what I'd thought it would look like at this point in time.

its publishing tools are limited in effectiveness.[39] Drupal, an open-source content management system, allows users to create limited profiles, and tracks user participation in a site, but that information remains relatively static rather than being used dynamically to help generate connections amongst users. Recommendation systems of varying stripes are in use at a number of commercial sites (usually of the "customers who bought x also bought y" type), but they usually rely upon keywords rather than full textual analysis, and little use of such systems has been made in the organization and dissemination of scholarly research. Then there are "reputation" systems, such as that in use in a large-scale discussion forum like Slashdot, which have proven effective at filtering out unhelpful or nuisance commentary, but their potential use in a system of scholarly review has as yet gone unexplored.

MediaCommons intends to bring such systems together, providing scholars with a range of tools through which to connect with one another, to produce and publish networked, multi-modal texts, to review those texts, and then, most crucially, to review the work of the reviewers, enabling the community to determine its own standards and adjudicate their implementation. In a peer-to-peer reviewing system, "reputation" will be determined not simply through an assessment of the scholar's own production but through an assessment of her reviewing practices. Reviews might, for instance, be rated on numerical scales that measure both their incisiveness and their helpfulness, resulting in a reviewer reputation score. Reviews written by scholars with better reputations would then be accorded more weight in determining the status of texts published through the network.

The emphasis in MediaCommons's peer-to-peer reviewing system is thus not simply on being smart, but on being helpful. The primary point of network members' commitment will need to be the

39 Beyond this, of course, lies user frustration with the sudden overflow of Facebook applications that resulted when the site developers opened up the system's API. In very short order, Facebook went from being a focused and contained, if limited, platform to a wild mishmash of annoying and seemingly pointless content. Perhaps a peer-to-peer reviewing system for Facebook apps — a community-based filtering system — might have helped stem the overflow; see Iskold.

advancement of the community as a whole, rather than the advancement of their own work; only in that way can both the individual scholar and the field as a whole succeed. In order to promote such a commitment, MediaCommons will need to find a way to implement a pay-to-play system, requiring community members to become active participants in the network's review processes in order to take advantage of it publishing capabilities. This might be done by constructing a point system, in which a scholar must earn credits by reviewing, which can then be spent on publishing, but it might also be done by linking the scholar's reputation as a reviewer to her own published texts, encouraging authors to improve their "karma," in Slashdot-speak, and thus the rankings of their work as a whole, by publishing more, and better, reviews of texts by others.

credentialing, revisited

But the idea of texts and authors being "ranked" and "rated" within the system raises several important concerns, most notably about the quantification of assessment across the academy. Faculty in the humanities in particular are justifiably anxious about the degree to which accrediting bodies and the U.S. Department of Education are demanding empirical, often numerical, accounting of things like "student learning outcomes," even in fields in which the learning itself isn't empirically-driven, but rather focused on interpretation and argument. Certainly we don't want our own work to be subject to the same kinds of "bean-counting" principles, in which statistics overtake more nuanced understandings of significance; as Lindsay Waters suggests, the danger in assuming that all knowledge can be quantified is that "[e]mpiricism makes people slaves to what they can see and count" (9), and the values of the humanities are largely non-countable. Moreover, our colleagues in the sciences might provide a bit of a cautionary tale: even in fields whose methods and evidence are largely empirically produced, concerns

about the reliance on citation indexes and impact factors as metrics of faculty achievement are growing.[40] We certainly don't want to suggest to tenure and promotion review committees that the data produced through a process of online peer-to-peer review is a more accurate evaluation of faculty performance simply because it is numerical.

On the other hand, we're already relying upon a system that's even more reductive than the kinds of metrics that the web can provide; the results of the current system of peer review are a simple binary: either the article was published in a peer-reviewed venue or it was not. There is precious little nuance in such a mode of evaluation, little room for considering whether a text published in a non-traditional format has been important in its field, little means of assessing the value of a scholar's contributions to a field outside of standardized modes of publishing. Network-based peer-to-peer review can provide us with certain kinds of information that can help complicate this practice, including of course the quantitative, such as numbers of inbound links, of comments, of citations, and of course statistical analysis of community-based review practices, but also including a wide range of qualitative, evaluative, interpretative commentary from the other authors and readers interacting with the texts we produce. No single measure can demonstrative proof of scholarly effectiveness, but a range of such information, including both the numerical and the narrative, the empirical and the ephemeral, can help illuminate the wide variety of ways that texts interact with the community of scholars.

The question remains, of course, whether the various credentialing bodies that currently rely on peer review's gatekeeping function will be satisfied with the kinds of information that such a system can provide. This is the point at which I must fall back on polemic, and simply insist that they must — that

40 See, as only two among many possible citations, Seglen and Richard Smith. Don Brenneis has likewise drawn my attention to the grave concern in the UK about chancellor Gordon Brown's decision to replace the Research Assessment Exercise, which previously determined funding for British universities, with a very narrow set of metrics including citation indexes; see Alexandra Smith.

we must say to hiring committees, tenure and promotion review bodies, and, most importantly, ourselves, that the fact that ostensibly anonymous reviewers didn't determine whether an article or monograph was worthy of publication shouldn't matter. A system of peer-to-peer review won't give us an easy binary criterion for determining "value" — but then, if we're honest, it never has. It will, however, give us invaluable information about how a scholar is situated within her field, how her work has been received and used by her peers, and what kind of effect she is having on her field's future. Requiring an up-or-down measurement of such a complex set of factors, or even relying on computationally produced metrics, can never provide an adequate substitute for the real work that such credentialing bodies must do: reading and assessing the scholarship, and engaging with expert analysis on the relationship between the scholarship and the field.[41] It is in part our desire for shortcuts, for a clear and quantifiable set of benchmarks by which we can judge "quality" without having to do the labor ourselves, that has gotten the academy into its current predicament, in which the very systems of production on which it relies are crumbling. Until institutional assumptions about how scholarly work should be assessed are changed — but moreover, until we come to understand peer-review as part of an ongoing conversation among scholars rather than a convenient means of determining "value" without all that inconvenient reading and discussion — the processes of evaluation for tenure and promotion are doomed to become a monster that eats its young, trapped in an early twentieth century model of scholarly production that simply no longer works.

works cited

Anderson, Chris. *The Long Tail: Why the Future of Business is Selling Less of More*. New York: Hyperion, 2006.

41 Lindsay Waters: "Reading the papers themselves! How quaint! How medieval!" (20).

——. "Technical Solutions: Wisdom of the Crowds." *Nature* Peer-Review Web Debate. 2006. <http://www.nature.com/nature/peerreview/debate/nature04992.html>.

"The Battle for Wikipedia's Soul." *The Economist*. Technology Quarterly (6 March 2008). Online: <http://www.economist.com/search/displaystory.cfm?story_id=10789354>.

Biagioli, Mario. "From Book Censorship to Academic Peer Review." *Emergences* 12.1 (2002): 11-45.

Blair, Carole, Julie R. Brown, and Leslie A. Baxter. "Disciplining the Feminine." *Quarterly Journal of Speech* 80 (1994): 383-409.

boyd, danah. "Loss of Context for Me on Facebook." *Apophenia* (10 August 2007). <http://www.zephoria.org/thoughts/archives/2007/08/10/loss_of_context.html>.

——. "Viewing American Class Divisions Through Facebook and Myspace." *Apophenia* (24 June 2007). <http://www.danah.org/papers/essays/ClassDivisions.html>.

Brown, John Seely. Panel presentation, Claremont Graduate University colloquium on Social Entrepreneurship and Design, 1 June 2007.

Brownlee, Christen. "Peer Review under the Microscope." *Science News* 170.25 (16 Dec 2006): 392-93.

Burnham, John C. "The evolution of editorial peer review." *JAMA* 263.10 (March 9, 1990): 1323-9.

Campanario, Juan Miguel. "Peer Review for Journals as It Stands Today, Part 2." *Science Communication* 19.4 (June 1998): 277-306.

Campbell, Phillip. "Nature's Peer Review Trial." *Nature* Online (June 2006). <http://www.nature.com/nature/peerreview/>. Accessed 2 February 2008.

Cohen, Noam. "A History Department Bans Citing Wikipedia as a Research Source." *New York Times* (21 February 2007). Online: <http://www.nytimes.com/2007/02/21/education/21wikipedia.html>. Accessed 16 February 2008.

"CZ:About." *Citizendium*. <http://en.citizendium.org/wiki/CZ:About>. Accessed 16 February 2008.

Falkow, Sally. "The Flip Side of Corporate Blogging." *The Navigator*. Cision, Online: <http://navigator.bacons.com/>.

Fitzpatrick, Kathleen. "MediaCommons: Scholarly Publishing in the Age of the Internet." *MediaCommons* (29 March 2007). <http://mediacommons.futureofthebook.org/scholarlypublishing/>.

Ginsparg, Paul. "Can Peer Review Be Better Focused?" *Science and Technology Libraries* 22.3-4 (2002): 5-17.

Godlee, Fiona. "The Ethics of Peer Review." In Jones, Anne Hudson and McLellan, Faith, eds. *Ethical Issues in Biomedical Publication*. Baltimore, Md.: Johns Hopkins University Press, 2000. 59-84.

Guédon, Jean-Claude, with Raymond Siemens. "The Credibility of Electronic Publishing: Peer Review and Imprint." *TEXT Technology* 11.1 (2002): 17-35.

Hall, Donald. *The Academic Self: An Owner's Manual.* Columbus: Ohio State University Press, 2002.

Harnad, Stevan. "Learned Inquiry and the Net: The Role of Peer Review, Peer Commentary, and Copyright." *Learned Publishing* 11.4 (1998): 283-292.

Holbo, John. "Will Work for Whuffie - Or - Anything You Can Do, I Can Do Meta." *The Valve* (19 Novenber 2005). <http://www.thevalve.org/go/valve/article/will_work_for_whuffie_or_anything_you_can_do_i_can_do_meta/>.

Iskold, Alex. "Facebook: What if More is Less?" *ReadWriteWeb* (27 September 2007). <http://www.readwriteweb.com/archives/facebook_what_if_more_is_less.php>.

Jenkins, Henry. *Convergence Culture: Where Old and New Media Collide*. New York: New York University Press, 2006.

Jensen, Michael. "Authority 3.0: Friend or Foe to Scholars?" *Journal of Scholarly Publishing* 39.1 (October 2007): 297-307. Rpt. of "The New Metrics of Scholarly Authority," *Chronicle of Higher Education* 53.41 (15 June 2007): B6.

Koop, Thomas, and Ulrich Pöschl. "Systems: An open, two-stage peer-review journal." *Nature* Peer-Review Web Debate. 2006. <http://www.nature.com/nature/peerreview/debate/nature04988.html>.

Kronick, David A. *"Devant le Déluge" and Other Essays on Early Modern Scientific Communication*. Lanham, Md.: Scarecrow Press, 2004.

——. "Peer review in 18th-century scientific journalism." *JAMA* 263.10 (March 9, 1990): 1321-2.

Kroski, Ellyssa. "Authority in the Age of the Amateur." *Info Tangle* (20 February 2006). <http://infotangle.blogsome.com/2006/02/20/authority-in-the-age-of-the-amateur/>.

Levien, Raph. "Advogato's Trust Metric." *Advogato* <http://www.advogato.org/trust-metric.html>.

Peters, Douglas P. and Stephen J. Ceci. "Peer Review Practices of Psychological Journals: The Fate of Published Articles, Submitted Again." Originally published in *Behavioral and Brain Sciences* 5.2 (1982). Reprinted in Shatz, David, *Peer Review: A Critical Inquiry*. Lanham, Md.: Rowman & Littlefield, 2004. 191-214.

Philica FAQ. *Philica*. Online: <http://philica.com/faq.php>.

Philica Tutorial. *Philica*. Online: <http://philica.com/tutorial.php>.

Pöschl, Ulrich. "Interactive journal concept for improved scientific publishing and quality assurance." *Learned Publishing* 17.2 (April 2004): 105-13.

Quinet, Raphaël. "Social Experiment on Slashdot Moderation." *Advogato* (24 January 2000). <http://www.advogato.org/article/27.html>.

Radway, Janice. "Research Universities, Periodical Publication, and the Circulation of Professional Expertise: On the Significance of Middlebrow Authority." *Critical Inquiry* 31.1 (Autumn 2004): 203-28.

Regalado, Mariana. "Research Authority in the Age of Google: Equilibrium Sought." *Library Philosophy and Practice* (2007). <http://www.webpages.uidaho.edu/~mbolin/regalado.htm>.

Rennie, Drummond. Commentary on Fabiato, A., "Anonymity of reviewers." *Cardiovascular Research* 28 (1994): 1142–1143.

——. "Editorial Peer Review: Its Development and Rationale." In Godlee, Fiona and Jefferson, Tom, eds. *Peer Review in Health Sciences*. London: BMJ Books, 2003. 1-13.

Modern Languages Association. *Report of the MLA Task Force on Evaluating Scholarship for Tenure and Promotion*. December 2006. Online: <http://www.mla.org/pdf/task_force_tenure_promo.pdf>.

Roy, Rustum, and James R. Ashburn. "The Perils of Peer Review" (letter to the editor). *Nature* 414 (22 Nov. 2001): 393-94.

Seglen, Per O. "Why the Impact Factor of Journals Should Not Be Used for Evaluating Research." *British Journal of Medicine* 314 (1997): 497.

Shatz, David. *Peer Review: A Critical Inquiry*. Lanham, Md.: Rowman & Littlefield, 2004.

Smith, Alexandra. "Scientists Concerned over Research Assessment Changes." *The Guardian* (7 December 2006). Online: <http://education.guardian.co.uk/RAE/story/0,,1966570,00.html>

Smith, Richard. "Commentary: The Power of the Unrelenting Impact Factor—Is it a Force for Good or Harm?" *International Journal of Epidemiology* 35.5 (2006): 1129-1130.

Spier, Ray. "The History of the Peer-Review Process." *TRENDS in Biotechnology* 20.8 (August 2002): 357-58.

Sunstein, Cass R. *Infotopia: How Many Minds Produce Knowledge*. New York: Oxford University Press, 2006.

Vaidhyanathan, Siva. *The Anarchist in the Library: How the Clash Between Freedom and Control is Hacking the Real World and Crashing the System*. New York: Basic Books, 2004.

——. "Where is This Book Going?" *The Googlization of Everything* (25 September 2007). <http://www.googlizationofeverything.com/2007/09/where_is_this_book_going.php>.

Wardrip-Fruin, Noah. "EP Meta: Chapter Four." *Grand Text Auto* (16 February 2008). <http://grandtextauto.org/2008/02/16/ep-meta-chapter-four/>.

Waters, Lindsay. *Enemies of Promise: Publishing, Perishing, and the Eclipse of Scholarship*. Chicago: Prickly Paradigm Press, 2004.

Young, Jeffrey R. "Blog Comments and Peer Review Go Head to Head to See Which Makes a Book Better." *Chronicle of Higher Education*. 22 January 2008. Online: <http://chronicle.com/free/2008/01/1322n.htm>.

Zuckerman, Harriet, and Robert K. Merton. "Patterns of Evaluation in Science: Institutionalization, Structure, and Functions of the Referee System." *Minerva* 9 (1971): 66-100.