

Academia doesn't scale

Eric Blair

14 December 2007

The current academic model is based on the academic societies of the 18th and 19th centuries, e.g., The Royal Society of London for Promoting Natural Knowledge¹, est 1660. These societies were everything you'd imagine: a bunch of wealthy white guys, all best of pals, generally brilliant, debating and experimenting. Many wore powdered wigs. Journals were often filled with cleaned-up letters or speeches given by members.

This is the root of the modern academic journal. Unfortunately, the model doesn't scale from a gentleman's club of up to a thousand members to the global mass that is modern academia. It reveals a specific worldview about where value comes from, which is not as advertised.

Ad hominem caveats This little column was hard to write because I don't want you to think I'm ranting about how all academics are evil. No, the simple thesis is that the current peer-reviewed journal system is ill-suited to the modern world. It works well for the small academic societies of the 1800s, and needs to be dropped now that the academic community is global and decentralized. After the initial statement of problem, I'll get to some suggestions of how modern academia could evolve past the 1800s.

As for my own experience, I've got my share of rejections from academic journals on the one hand, and a couple of papers published (or on their way to being published) on the other. I've also written two peer-reviewed books, partly because the peer-review process for books overcomes some of the small-circle problems that journal peer-review suffers. So I don't comfortably float around the peer-reviewed world, but I'm not a total outsider either.

That said, we can get to the key conflict, which is basically the eternal fight between the small-circle meritocracy and egalitarian democracy, except sometimes those with the most merit aren't in the small circle.

The conversation The academic literature is frequently described as an ongoing conversation. This is a fundamentally different concept from being a repository of the best current work. The ongoing conversation story neatly follows the tradition of the society journals, which were sometimes literally the record of ongoing discussions held by the individuals in person or via letters.

¹<http://www.wired.com/science/discoveries/news/2007/11/dayintech.1128>

As we've all experienced at parties, it's hard to walk in on an ongoing conversation, especially walking in on a conversation with people who don't know you. The typical first response is not *my, what an interesting new perspective* but *who the fuck are you?*. The same holds with a submission to an academic journal: the first question is who you are, where you're coming from, what perspective you have on the ongoing conversation, and then finally what you actually have to say. Law reviews, which are not really peer-reviewed in the traditional sense, take the direct approach and typically require a résumé with submission.

For a small society, the who-are-you stage of things goes pretty quickly, because if the society's members are willing to talk to you then you've already passed a crackpot test. But now that we're an egalitarian world and journals accept paper submissions from anybody, the editor needs to start with filtering crackpots, and then move on to evaluating the merits of the work.

Fear of crackpots The 'ongoing conversation' model is inherently conservative. After all, sometimes the topic of conversation really needs to change. To give a concrete example, there is a thread in the voter turnout literature over the claim that people turn out based on purely self-interested means: they find the likelihood that they'd be the pivotal voter, multiply by the expected personal gains from one candidate over another, and then choose to vote accordingly. The ongoing conversation is over reconciling this claim with the data, which contradicts it every step of the way. Gee, maybe the theory, which is countered by both intuition and empirical data, is just *wrong*. But it remains the baseline model, and if you want to write a paper about why people turn out to vote, you need to spend some portion of your time engaging in this ongoing conversation that really should have died a decade ago.

The conservativeness is rooted in a fear of crackpots, as I'd discussed in entry #224. Frankly, the number of times that somebody walks in from the proverbial left field and says something earth-shattering really is rare. So I agree with the general academic consensus that a paper that has no (or minimal) literature review of any sort is certain to be uninformed. However, that doesn't mean that a paper that is trying a new direction or using new techniques is necessarily wrong. Nonetheless, a change of topic that doesn't fit into the ongoing conversation has minimal odds of getting through peer review.

Speaking a New Word I submitted a model of network formation to *Economics Letters*, which gave me a one-paragraph rejection, the gist of which was: a model of network formation isn't Economics. I suppose this fact would be news to the editors of the Review of Network Economics. The intent of this example is to point out that the humans editing the journals evaluate the work based upon what they're comfortable with, not based upon some sort of objective criterion that they as humans aren't capable of achieving.

We can go back to my favorite question: where does value come from? It doesn't come from new knowledge about the world, but from the belief by other humans that such knowledge is important and relevant. That is, the peer review system establishes science as firmly subjectivist and relativist, no matter how much it pays lip-service to objectivity.

Or, we can go back to what I consider to be the fundamental rule of nonfiction writing: the work should make the reader feel smarter. The reader should know how to do something s/he didn't know how to do before, or learn facts and a means of structuring them that the reader hadn't known before, or otherwise make the reader feel more secure about his or her existing knowledge. In the political context, people are much more likely to read articles and books that agree with them than works that oppose them (and yes, this phenomenon has been extensively documented in peer-reviewed journals).

In the specific case of a peer-reviewed journal, it's not just anybody you're trying to make feel smarter, but the referee, who already has an established worldview and an established set of tools that he or she learned in graduate school. The process of presenting a new method, like presenting a network model to an old school economist, is exceptionally difficult, because new things are threatening and make the reader feel stupid until the reader has had time to absorb the full implications. Meanwhile, the rift between 'useful' and 'useful to the current members of the ongoing conversation' grows.

Let me again clarify that this isn't about being evil, it's about a well-known fact of human nature: we are more comfortable embracing the familiar than the foreign. Travis and Collins [1991] explain that this is not cronyism or an 'old boy' network, but the tendency to pick people from your intellectual school of thought over outsiders—which can often look a lot like cronyism and an old boy network.

Anonymity Anonymous review may make sense in a small society, where you may have to reject your best pal. In cases like these, you always know who wrote the rejection anyway. But when the author and reviewer may be continents away, anonymity just produces low-quality reviews.

Anonymous peer review doesn't scale, but it is necessary to perpetuate the system as it exists today. Pretty much every academic has a story about a rejection they got that was rude to the point of humorous. My own rejections have just been boring and typically indicated that the review didn't read the paper. E.g., I once received a rejection—after a year and a half wait—based on how I numbered my theorems. A rejection I received after I started writing this essay chided me for failing to discuss the multiple possible modes in a probit model—except the probit likelihood function is globally concave, and therefore always has exactly one mode. But hey, there are no peer review reviewers to check the facts or merits of the anonymous review.²

This one has at least a partial solution, because anonymity is an endogenous social norm: we can sign our reviews. I do, and make a point of never saying anything in a rejection letter that I wouldn't say to the author's face. The problem with this, and why it's not the norm, is that I can't be lazy. I can't reject based on ad hominem excuses or theorem numbering or a vague sense of dislike. In short, I can't pretend that I'm

²A not-anonymous peer points out that the editor should be checking the reviews for quality and acting accordingly. However, part of the referee's job is to save time and effort for the editor, and editors are generally inclined to trust the referee, so human laziness and trust generally prevent editors from overturning all but the truly worst reviews. I did once have an editor who told me that a referee report was so bad that it indicated more that the referee was a crackpot than that there were problems with the submission, so the editors are not entirely asleep.

being egalitarian and working from the merits of the paper while actually working to maintain a closed society.

Competition In both small society and the globe at large, there's always a conflict between striving for the greater good and for individual advancement. For a system where peers review your work, this is a central conflict, because the people most qualified to evaluate your work are your direct competitors.

In a small society, your direct competitors are probably also your pals. Any business organization is naturally something of a social organization as well, and apart from some famous disputes (e.g., Liebenitz v Newton), we find that all those letters between famous colleagues were generally collegial. So when somebody accepted an article from a competitor, it was at least from a competing pal, who would be able to give a leg up next time in return.

In a global society, the author and the person reviewing a paper are just competitors. The one who is doing the review was selected because s/he is considered to be part of the established system, and has something to lose, like real live funding, by expanding the membership of the established system.

I went to a delightful conference the other day regarding a new paradigm for demography. When the question of getting funding from the NIH came up, a couple of people suggested just not bothering. The people who read the grant applications are the people who were most successful in the last decade, and therefore are the ones who most stand to lose from a change in paradigm, and are in some ways the least qualified to evaluate a change of topic from their own life's work.

The alternatives Now and then, people tell me that the peer review system is a revered part of science. This is even political—the PRISM coalition³ is a group of academic publishers who oppose open access journals under the presumption that it is the private, traditional peer review system that ensures high quality in publishing. The claim as I stated it here implies that peer review has evolved into what it is based on centuries of refinement and improvement. Rather, it's a throwback: it's a system that primarily emerged among small academic societies and has entirely failed to adapt as modern academia stopped being about a few inner circles and became an open, global meritocracy.

So, what can we do? There are already many online repositories to be had. My favorites are the Arxiv⁴ for math/stat/physics and the self-descriptive Social Science Research Network⁵. The SSRN is especially important because of the absolutely pathetic speed of peer review in the social sciences. That paper above that got rejected over nonexistent nonconvexities in the probit (it's December): I submitted that paper in January. Amusingly enough, the reviewer criticized my literature review for not citing papers published in May. So if we relied solely on the formal journals for social science research, the entire system would quickly grind down, as everybody would be a year behind all the time. Instead, we spend more time at places like the SSRN.

³<http://www.prismcoalition.org/>

⁴<http://arxiv.org>

⁵<http://ssrn.org>

So the SSRN is already eating the journals' lunch with regards to the work of archiving and dissemination. But the SSRN lacks peer review, and an endorsement system is still valuable and important. It's still the case that 90% of everything is crap, and some papers are more important, better written, or otherwise of higher quality than others. We humans with limited time on this Earth need some sort of guidance toward what is worth our time.

Could Arxiv and the SSRN implement peer review? Sure, in a heartbeat. Especially if we give up on the gentleman's society rule of anonymity, a paper's web page could include endorsements or comments from others. Readers will then have more than enough to evaluate whether the paper is useful, accurate, and so on. Depending on how reviewers are assigned, authors who write about relatively new methods may be more likely to find another party who doesn't feel dumber when confronted with that specific method. [The Arxiv already has a very weak endorsement system, but it doesn't yet provide as much information as users need.]

We'd like authors to revise based upon comments and improve things accordingly; which means that there'd need to be some sort of revision control system in place. The author may have the right to publicly respond to public peer review, in which case the ongoing conversation would happen right there on the page. Since a peer review is now an invitation from the editor to publish a short article, colleagues now have a half-decent incentive to actually do peer review beyond the vague sense of responsibility that is the sole, insufficient motivator now.

[And yes, people suck, and there are bad academic apples who would say mean things and try to ruin the system for everyone. But this is still a system that restricts commenting access to named and identified peer reviewers, not YouTube. All of these details are a low-grade kind of problem which SSRN and Arxiv could easily surmount with a few days of coding and some vigilance on the part of the editors.]

A system like this would also take the semi-sacred significance off of peer review, which is a good thing. The popular media often refer to peer-reviewed papers as if they are unquestionably valid, and not-peer-reviewed papers as necessarily pseudoscience, but with all the problems underlying the system above, the signal is not so clear. A public endorsement system would guide the reader toward good papers but not imply that the paper is the gospel truth—just that two or three knowledgeable but fallible humans found it to be of high quality.

I expect that the concept of a paper—a single unit of scholarship that others can read and refer to—will continue to exist, but its delivery and evaluation will have to change. Delivery has already changed: nobody goes to the library to pick up dusty bound volumes from the 1800s, since they're in PDF format on Jstor. Nobody in the current ongoing conversation of social science even bothers with new journals as they are mailed out parcel post, because they're just the archiving of research from a year or two ago. The archiving process will be online no matter what.

The endorsement system as it stands will live a lot longer, because it clearly benefits the incumbents, provides a means for the small inner circle to keep itself small, and provides a shield of anonymity that many reviewers continue to use as a crutch. Especially in social sciences, this is a sad state of affairs: we have a dozen journals devoted to mechanism design and making sure that people's incentives are aligned with our overall social goals, and yet we still base decisions on anonymous comments from those who are most likely to lose funding and relevance to something new.

References

G D L Travis and H M Collins. New light on old boys: Cognitive and institutional particularism in the peer review system. *Science, Technology, & Human Values*, 16 (3):322–341, 1991.