

Improving the Refereeing Process: A Simple Proposal

We've all experienced the sinking feeling of opening a journal's reply to our paper submission only to see a rejection. Many, perhaps most, of us have had the misfortune to experience this in its most unpleasant and least defensible form: when the reply comes six months to a year or more after the submission and the referee report is so perfunctory it's clear that no effort was made to read past the introduction. I've certainly been through this a number of times myself, with twists ranging from the comic (as when a referee stated that he or she had mistaken the paper in question for a different one) to the tragic (as when some referees have clearly fundamentally misunderstood the main results). I've heard even more absurd and frustrating stories from friends. It's hard not to suspect that in some cases, referees take a negative report as the path of least resistance to avoid the hard work of a thorough refereeing job. In theory, an editor may recognize a shoddy report and seek out a better one, but in practice this rarely happens.

Some journals have taken an important step in the right direction by soliciting quick expert opinions prior to sending out a paper to referee, but more could be done and without very much difficulty. My simple proposal is to completely separate out the evaluation of the quality of the results and techniques—which can usually be carried out quite quickly and for the vast majority of papers is the sole factor determining ultimate acceptance or rejection—from the more painstaking process of attempting to verify correctness. Thus, a journal would first solicit one or more quick expert opinions on a paper and then make a binding decision on whether the paper is good enough to accept, assuming the mathematics is correct. Only in the case of a positive decision would the paper be sent out for a thorough refereeing for correctness.

Evidently, the great advantage of this system is that in principle rejections would always occur quickly, except in rare cases of unsalvageable mathematical errors. Also appealing is that it would be relatively easy to implement. Although it would create a modest amount of additional work for journals not already soliciting expert opinions, the benefits could be substantial not only to authors but to journals as well. As with the current expert opinion system, editors would be free to select more prominent and senior mathematicians for the evaluation of the importance of a paper while choosing younger mathematicians with more time for careful reading to evaluate correctness. One could easily imagine that the first journals to advertise the adoption of this system would see an increase in the quality of papers submitted, and once the system is adopted more widely, the average delay between the public posting of a paper and its publication would decrease,

counteracting to some extent the present march towards irrelevance of published versions of papers.

Switching to this system would also open the door to further experimentation. For instance, the second referee would not have to be anonymous, since he or she would not be evaluating the quality of the paper. This would allow real credit for improvements due to the referee (in contrast to the currently ubiquitous thanking of the “anonymous referee”) and, depending on how it was implemented, could put more pressure on referees not to simply rubber-stamp the paper.

Despite the present rapid pace of change in mathematical publishing, I believe that the peer-review process continues to play an important role in our profession. Unfortunately, in its present incarnation, it also causes a lot of collateral damage, much of it unnecessary. While the above proposal is no panacea, it offers a practical approach to addressing one of the most serious problems with the process. I hope that editorial boards will seriously consider adopting it.

—Brian Osserman
University of California, Davis
osserman@math.ucdavis.edu