Selecting Scientific Papers for Publication via Citation Auctions

Josep Lluis de la Rosa, *University of Girona*
Boleslaw K. Szymanski, *Rensselaer Polytechnic Institute*

The scientific community has been suffering from peer review for decades. This process (also called *refereeing*) subjects an author's scientific work or ideas to the scrutiny of one or more experts in the field. Publishers use it to select and screen manuscript submissions, and funding agencies use it to award research funds. The goal is to get authors to meet their discipline's standards and thus achieve scientific objectivity. Publications and awards that haven't undergone peer review are often regarded with suspicion by scholars and professionals in many fields. However, peer review, although universally used, has many drawbacks.

We propose replacing peer review with an auction-based approach: the better the submitted paper, the more scientific currency the author will likely bid to have it published. If the bid correctly reflects the paper's quality, the author will be rewarded in this new scientific currency; otherwise, the author will lose this currency. We argue that citations are an appropriate currency for all scientists.[1] We believe that citation auctions will encourage scientists to better control their submissions' quality. It will likely also inspire them to prepare more exciting talks for accepted papers and to invite discussion of their results at congresses and conferences and among their colleagues. In the long run, citation auctions could have the power to greatly improve scientific research.

Peer review's drawbacks

While some believe that passing peer review is a certificate of validity, others are far more skeptical. One of the most common complaints is that peer review is slow: a submitted paper typically takes several months, or even years in some fields, to appear in print. Such a delay in a fast-growing field is devastating for the propagation of ideas and needs a solution. Another major drawback is its cost in terms of hours of volunteer work devoted to refereeing.

In addition, some sociologists of science argue that peer review makes publication susceptible to control by elites and to personal jealousy. Peer review might suppress dissent against "mainstream" theories. Reviewers tend to be especially critical of conclusions that contradict their own views but happily accept those that accord with them. At the same time, elite scientists are more likely than less established ones to be sought as referees, particularly by high-prestige journals or publishers. As a result, some argue, ideas that harmonize with the elite's views are more likely to be published and to appear in premier journals than iconoclastic or revolutionary ideas. Consequently, the whole process obstructs and delays the emergence of new ideas and scientific revolutions.

However, others have pointed out that scientists have many journals in which to publish, making control of information difficult. In addition, peer review's decision-making process, in which each referee gives an opinion separately without consulting others, is intended to mitigate some of these problems. Nevertheless, this process still doesn't address the cost of peer reviewing and publication delays.

Moreover, peer review tends to accept those weaker papers that have a mix of prestigious and unknown authors. This is because referees typically trust a paper with prestigious authors even if they don't fully understand its contributions, believing that "this must be true, must be good, and the final submission will address any potential drawbacks." This might not systematically be bad, but there are more elegant ways to give opportunities to new authors. For the sake of improving the quality of science, this trust in important coauthors should be reduced. Blind peer review, one possible way to deal with this problem, is still an imperfect solution because it doesn't decrease the cost of reviewing or publication delays.

Citations and their auctions

According to Philip G. Altbach, director of Boston College's Center for International Higher Education, the citation system was invented mainly to understand how scientific discoveries and innovations are communicated and how research functions.[2] On the basis of our research on the innovative use of auctions,[3,4] in May 2006 we considered using this approach to create an alternative to peer review.[5,6] Conferences often suffer from low participation

1541-1672/07/$25.00 © 2007 IEEE
Published by the IEEE Computer Society

Some Alternatives to Peer Review

In some fields, such as astronomy, much of the communication about new results no longer takes place through peer-reviewed papers but through preprints submitted to electronic servers such as arXiv.org.

The recently launched online journal *Philica* (www.philica.com) exemplifies one new way to redress many of peer review's problems. Unlike a traditional journal, it immediately publishes all submitted papers; review takes place afterwards. Reviews are still anonymous, but instead of an editor choosing the reviewers, any researcher who wishes to can review an article. Reviews appear at the end of each paper, thus giving the reader criticism or guidance about the paper instead of determining whether it's published. This means that reviewers can't suppress ideas with which they disagree.

Some authors (for example, Stefano Mizzaro[1]) suggest scoring papers and authors, letting readers act directly as referees, and receiving feedback on whether the readers provided good-quality judgments. So, good readers would earn good reputations.

Reference

1. S. Mizzaro, "Quality Control in Scholarly Publishing: A New Proposal," *J. Am. Soc. Information Science and Technology*, vol. 54, no. 11, 2003, pp. 989–1005.

and little discussion of the presented papers, because authors focus on getting their papers to the conference rather than widely disseminating their results. Moreover, organizers extend extraordinary effort selecting the best conference paper, thus creating a heavy refereeing workload. To avoid these pitfalls, we propose an auction system combining citation's predictive value with authors' intimate knowledge of their papers. (For other alternatives to peer review, see the sidebar.)

We can characterize the situation as follows. Nowadays, a scientist wants to publish his or her results in conferences and journals to gain citations and reputation. As evidenced by the acceptance rates, highly ranked and read conferences and journals select only a small number of submitted papers—hopefully, those that will generate the most citations. In our approach, conferences and journals hold an auction to select those papers whose authors have bid the highest. Each bid represents an author's prediction of how many citations that paper will receive. To make this bid trustworthy, we consider the number of citations an author has received for previously published papers as that author's *citations wallet*—the amount of "cash" the author possesses. So, the number of citations in authors' wallets limits their bids, and every auction-winning paper withdraws a number of citations from its authors' wallets.

So, what should the wallet contain? Citations from one conference that has accepted a paper by the author? Citations from a group of related conferences? Citations of all the author's papers? We use the citations from either a group of conferences or all the author's papers. This lets an author collect citations in weaker conferences to get into stronger ones or into a journal. Consequently, rational authors will bid the highest number of citations they think acceptance will require, but to the limit of their wallet's cash (that is, their past performance). Authors might lose cash if their winning bid exceeds the number of citations the paper will generate. Conversely, they might win more cash if the published paper generates more citations than they bid. An author's ultimate goal will be to keep his or her wallet growing.

Our approach has two main benefits. First, it reduces refereeing costs because paper selection via a citation auction doesn't need prior understanding of the paper's content to evaluate its quality. Second, authors will be much more committed to their papers' quality. They will also focus much more on wide dissemination and detailed explanations of their papers to maximize the number of citations. In short, this novel approach emphasizes active promotion of ideas while reducing peer review's high expense.

We're fully aware that making this idea workable involves challenges. Here we discuss the three most important ones.

First, if each author receives credit for all citations to the paper, we'll have wallet inflation. A paper with five authors, five citations, and a bid of five citations will contribute 25 citations total to all the authors' wallets, but the payment will be just five citations from one author's wallet. To avoid the inflation, each citation will earn one unit of credit for the paper, regardless of the number of its authors. Additionally, we'll assign fractions of each citation to the individual authors according to their "citations contract." This contract might reflect the percentage of each author's participation. If no citations contract exists, all citations will be assigned to the first author, who will then decide how to share them. To encourage collaborations, we also let all the paper's authors contribute their citations to a joint bid. This will solve the inflation problem. Coauthors are often students, so they can build their wallets for a future independent career while working with their advisors on their theses.

The second challenge involves self-references. An author could safely bid a number of citations equal to the number of that author's self-citations in the paper. After the paper's publication, the author would automatically receive the credit for the citations, immediately rebuilding the wallet. To avoid this effect, none of the self-citations or citations of papers by a bid's contributors should be added to the author's wallet.

The third challenge regards the wallet's initial content. We assume that the initial content is zero citations. So, new authors (for example, graduate students) or authors who depleted their wallets via too-aggressive bidding wouldn't even be able to try to get their papers published. To avoid such situations, a certain fraction of papers should undergo peer review, but with much higher acceptance criteria. This will also avoid the problem of the boundary-quality papers, which are difficult and time consuming to evaluate.

Alternatively, authors could receive "sponsorship." For example, when one such author looks for credit, another researcher might loan some citations from his or her wallet (likely after the paper's review and after the lender has suggested improvements). For a supervisor or advisor, the

Table 1. A simulation of four citation auction strategies.*

	No. of citations											
	1999		2000		2001		2002		2003		2004	
Author and paper	Bid	Earned	Bid	Earned	Bid	Earned	Bid	Earned	Bid	Earned	Bid	Earned
Author 1 (aggressive)												
Author 1's wallet	13		13		13		8		7		7	
Paper 1												
Paper 2	0											
Paper 3			0		6	1						
Paper 4							4	3				
Paper 5									1	1		
Author 2 (cautious)												
Author 2's wallet	0		6		7		7		5		3	
Paper 1	2 (loan)	8										
Paper 2	1 (loan)	1										
Paper 3			2	3								
Paper 4			1	1								
Paper 5					2	2						
Paper 6							1	1				
Paper 7							1	0				
Paper 8							1	0				
Paper 9									1	0		
Paper 10									1	0		
Author 3 (very cautious)												
Author 3's wallet	0		1		1		2		2		1	
Paper 1	1 (loan)	3										
Paper 2	1 (loan)	0										
Paper 3			0		1	0						
Paper 4					2 (loan)	4						
Paper 5							1	1				
Paper 6									1	0		
Author 4 (no risk)												
Author 4's wallet	0		0		0		0		0		2	
Paper 1	0		0		1 (loan)	1	0					
Paper 2									1 (loan)	3		

* Blue squares indicate unsuccessful bids.

motivation for a loan is clear: lending part of the prestige represented by the citations in his or her wallet is a kind of investment. Hopefully, the student will eventually return the citations, perhaps with extra citations as interest, which an agreement between the two could formalize. We don't believe that this approach would benefit "bad students with famous supervisors" over "good students with unknown supervisors." That's because a donor system based on one scientist's rational criteria

Table 2. Author rankings generated by classic measures, the H-index, and citation auctions.

Author	Accepted papers	Total citations	Earned citations	Total productivity	H-index	Wallet	Earnings	Bid productivity	Losses
A1	4	18	5	1.25	2 or more	7	–6	–1.50	6
A2	10	16	16	1.60	2	3	5	0.50	0
A3	6	8	8	1.33	2	1	1	0.17	0
A4	2	4	4	2.00	1	2	2	1.00	0
The best	A2	A1	A2	A4	A1	A1	A2	A4	A2, A3, A4
Second best	A3	A2	A3	A2	A2 or A3	A2	A4	A2	A2, A3, A4
Third best	A1	A3	A1	A3	A2 or A3	A4	A3	A3	A2, A3, A4
The worst	A4	A4	A4	A1	A4	A3	A1	A1	A1

should be more efficient than bureaucratic systems based on collectives of scientists who have reached a consensus of what's good or bad. Finally, even unrelated researchers might want to invest their citations in a paper by a promising young talent for a profit of future citations.

It's beyond this article's scope to discuss the citation auction diffusion and promotion mechanism, which requires deep insight into the auction model.[3] Clearly, we'll need a new economic model derived from the citation auction to foresee how auctions increase the quality of research, decrease peer review costs and publication time, and so on.

A hypothetical case

Table 1 represents a possible scenario with a simulation of four researchers using different auction strategies (aggressive, cautious, very cautious, and avoiding all risk). The Bid column indicates an author's bid; blue squares indicate unsuccessful bids. The Earned column indicates the number of citations received by each paper published via the auction. For each year in which a paper was published, the Earned column shows the number of citations that as of May 2006 were made to that paper (to increase realism, we used actual data from a group of conferences). For simplicity, we use a *first-price sealed-bid* auction, where every author submits in a closed envelope a bid stating how many citations he or she is offering. This bid supposes to be lower than the number of citations the author will collect after the paper's publication. Some authors who couldn't bid higher than 0 because their citation wallets were empty borrowed citations from colleagues and paid back the loans from the citation earnings.

To measure papers' impact, we can examine several rankings of authors. The first is the citation wallet ranking, based on how many citations remain in an author's wallet after an auction. In this scenario, author 1 clearly leads with seven citations, followed by author 2 with three citations, author 4 with two, and author 3 with one. The wallet values are lower than the total number of citations for each author, which leads to another ranking (author 1, author 2, author 3, and author 4). However, the citation wallet ranking has the advantage of reflecting the actual number of citations (not just the number of expected citations) that the entire set of publications has generated for an author in a given time period. This ranking can track the behavior along a scientist's career, unlike the H-index,[7] which just provides a cumulative analysis of the scientist's best publications. So, these two measure different things. If an author's papers are ranked in descending order of their numbers of citations, the H-rank is the largest rank that is smaller than the number of citations that the corresponding paper has. Hence, the H-index tries to capture the impact of the key papers throughout a scientist's history. Citation wallet ranking tries to catch the impact of all the scientist's publications, which could be equal to the H-index in the case of outstanding scientists but will differ for other scientists (most of us). Table 2 compares several rankings, including the H-index.

Table 2 also shows the outcome compared to expectations from the auctions shown in table 1. The Earned Citations and Total Productivity columns accurately measure the accepted papers' quality. So, authors have incentive to gain the maximum number of citations, at least as many as they invested in the bid. Moreover, this system is self-regulating. If an author repeatedly underperforms in citations (the author's papers receive fewer citations than what he or she bid), the author's wallet will eventually approach zero. This might happen even to productive authors, such as author 1 in table 2, whose bid productivity is negative although the citation productivity is quite high. Such authors will have difficulty assuring the publication of their papers, because they will lose many auctions. Conversely, authors can quickly increase their wallet's size by submitting papers that will be highly cited, thus making it easy to publish papers in the future. Although the wallets of the most conservative strategies tend to grow, the other strategies might oscillate. This effect clearly differs from existing citation measures, which always grow through time.

Making citation auctions usable will involve many steps. The most important are to

- create a proof-of-concept of citation auctions, using a few selected publications;
- explore the most appropriate auction mechanisms;
- develop the technology for a *citation bank* using standard citation engines; and
- study the use of "taxes" for stabilizing the resulting economic model.

The adoption of citation auctions might take decades. However, the increased efficiency of scientific activities arising from their use should let them prevail over the less efficient peer review. ■

Acknowledgments

This work was supported partly by the EU project grant 34744 (ONE: Open Negotiation

Environment, FP6-2005-IST-5, ICT-for Networked Businesses), and by US National Science Foundation grant OISE-0334667. This article's content doesn't necessarily reflect the sponsors' position or policy; no official endorsement should be inferred or implied.

References

1. C. Carrillo et al., *Social Currencies and Knowledge Currencies*, IOS Press, 2006, pp. 266–274.
2. P.G. Altbach, "The Tyranny of Citations," *Inside Higher Ed*, 8 May 2006, http://insidehighered.com/views/2006/05/08/altbach.
3. J.-S. Lee and B.K. Szymanski, "A Novel Auction Mechanism for Selling Time-Sensitive E-services," *Proc. 7th Int'l IEEE Conf. E-Commerce Technology* (CEC 05), IEEE Press, 2005, pp. 75–82.
4. J.-S. Lee and B.K. Szymanski, "Auctions as a Dynamic Pricing Mechanism for E-services," *Service Enterprise Integration*, C. Hsu, ed., Kluwer, 2006, pp. 131–156.
5. J.L. de la Rosa, *Outline of Citation Auctions*, IOS Press, 2006, pp. 299–307.
6. J.L. de la Rosa and B.K. Szymanski, "Citation Auctions as a Method to Improve Selection of Scientific Papers," *Proc. 2nd Int'l Conf. Digital Information Management* (ICDIM 07), IEEE CS Press, 2007, pp. 479–486.
7. J.E. Hirsch, "An Index to Quantify an Individual's Scientific Research Output," *Proc. National Academy of Sciences*, vol. 102, no. 46, pp. 16569–16572; www.pnas.org/cgi/reprint/0507655102v1.pdf.

Josep Lluis de la Rosa is the founding director of the University of Girona's EASY Innova R+D Transference Centre for Innovation and Business Development. Contact him at peplluis@eia.udg.edu.

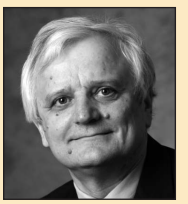

Boleslaw K. Szymanski is the Clair and Roland Schmitt Distinguished Professor of Computer Science and the founding director of the Center for Pervasive Computing and Networking at the Rensselaer Polytechnic Institute. Contact him at szymansk@cs.rpi.edu.