Special Issue — Future of Publication

Modeling peer review: an agent-based approach

Stefano Allesina

Stefano Allesina (sallesina@uchicago.edu), Department of Ecology and Evolution and Computation Institute, University of Chicago, 1101 E. 57th Street, Chicago, IL, USA.

Abstract

The peer review system is under severe strain. Corrections have been proposed, but experiments to determine effective measures are difficult to perform. I propose a framework in which alternatives to the current peer review system can be studied quantitatively using agent-based modeling. I implement three possible systems. I show how, all other things being equal, these alternatives produce different results in terms of speed of publication, quality control, reviewers' effort and authors' impact. This modeling framework can be used to test other solutions for peer review, leading the way to an improvement of how science is disseminated.

Keywords: Peer-review, agent-based modeling, publishing, editorial rejection, simulations

Introduction

The number of scientific publications is constantly rising, increasing the strain on the review process. The number of submissions is actually higher, as each manuscript is often reviewed several times before publication (Aarssen et al. 2008). To face the deluge of submissions, top journals reject a considerable fraction of manuscripts without review (Nature Editorial 2006), potentially declining manuscripts with merit. The situation is frustrating for authors (Smit 2006), reviewers and editors (Schroter et al. 2004, Smith 2006, Hochberg et al. 2009, McPeck et al. 2009) alike. Recently, several editors wrote about the “tragedy of the reviewer commons” (Hochberg et al. 2009), advocating for urgent corrections to the system. Buchanan (2010) argued that the next revolution in science could come from

changing the way publication and funding are allocated, rather than by new discoveries in a specific scientific field.

The main question is how to modify the current system so that the negative and distortive effects are minimized. Almost every scientist has ideas on how to improve the system (e.g., Campanario 1997, Cintas 2004, Jaffe 2006, List, et al. 2006, Hauser and Fehr 2007, Martinez-Abraín 2009), but it is very difficult to perform large scale experiments to test which measures would be most effective. Many statistical approaches to study the effects of peer review and potential improvements have been proposed (Neff and Olden 2006, Bornmann and Daniel 2006, Bornmann and Daniel 2008a, Bornmann and Daniel 2008b).

Surprisingly, relatively few attempts have been made to model peer review directly (Bentley 2009) in a mechanistic way, so that relevant quantities can be manipulated. Here I implement a simulation framework in which ideas on peer review can be quantitatively tested. I incorporate authors, reviewers, manuscripts and journals into an agent-based model and a peer review system emerges from their interactions.

As a proof-of-concept, I contrast an implementation of the classical peer review system adopted by most journals, in which authors decide the journal for their submissions, with a variation in which editors can reject manuscripts without review and with a radically different system in which journals bid on manuscripts for publication. I show that even small modifications to the system can have large effects on publication quality, time to publication, and matching between manuscripts and journals.

Typically, a reviewer is asked to evaluate a manuscript according to three main characteristics: i) Topic (T): Is the topic of the manuscript in line with the

journal? ii) Technical Quality (Q): Does the work use the best techniques? Are references adequate? Is the statistical treatment of the data correct? And iii) Novelty (N): Is the work groundbreaking? Does it convey new ideas and methods? Accordingly, in the three models, each author and journal is defined by three beta probability distribution functions, one for each characteristic (T , Q and N ; Fig. 1i; see Appendix). A manuscript is simply a sample from the three distributions for the author (t , q , n ; Fig. 1ii). Once an author has produced a manuscript, she has to decide where to submit it. In the “classic setting”, an author is encouraged to choose optimally the journal so that the impact of the publication on other scientists (number of readers, citations) is maximized and so is the impact on her career (grants, tenure). Let us assume that the impact of a journal is measurable (impact factor, eigenfactor) and that the author is able to determine with precision the probability of acceptance of her manuscript in a given journal. Then the author can optimally choose the journal by maximizing the product of acceptance probability and impact (Fig. 1iii). Note that this procedure postulates that authors have a cost associated to the submission (e.g., time spent editing, reputation). If there is no cost, authors will tend to submit always to the highest impact journals.

In the “classic setting”, once the manuscript is submitted to a journal, the editor assigns it to three referees. The referees are sampled at random from those whose expertise is best suited for the manuscript (see Appendix). The referees then provide the editor with an estimate of the three values characterizing the manuscript. The more familiar the reviewers are with the topic, the more accurate is their estimate (Appendix). The author revises the manuscript (increasing its quality and novelty; Appendix) and the review process is repeated (Fig. 1iv). Finally, the editor averages the t , q , and n values provided by the referees and uses this averaged estimate to decide the fate of the manuscript. This is done in the model by computing an acceptance probability. This probability is the product of the cumulative distribution functions of the journal for the estimate q and n , times a measure based on t that determines the interest of the journal in the topic of the manuscript (Fig. 1v; Appendix). The editor then accepts the manuscript with probability p_a and rejects it with probability $1 - p_a$. In the “classic setting”, if a manuscript is rejected then the author will submit it to the next best journal, unless it has been rejected 5 times (Appendix), in which case the author abandons the manuscript. The “editorial rejection” setting is exactly like the “classic”, but when the manuscript is submitted, the editors of the journal evaluate it rapidly: if the manuscript—according to the editors—does not have a sufficiently high probability of being accepted, it is rejected without consulting reviewers (Appendix). This variation

on the classic setting is actually adopted by most top-tier journals, with a percentage of rejection without review that can be as high as 83% (Nature Editorial 2006). The main goal is to avoid burdening referees with reviews that would not realistically lead to acceptance, while at the same time keeping the time from submission to decision extremely short.

i) Each Author is defined by three probability distributions.

ii) A new manuscript is a sample from the author's distributions.

iii) Authors submit their manuscript to the journal yielding the highest expected impact

Journal	Probability acceptance	Impact	Expected Impact	Rank
A	0.1	5	$0.1 \times 5 = 0.5$	IV
B	0.3	2	$0.3 \times 2 = 0.6$	III
C	0.2	10	$0.2 \times 10 = 5$	I
D	0.1	15	$0.1 \times 15 = 1.5$	II

iv) Reviewers evaluate the manuscript. The manuscript is then revised (increase q and n) and the process is repeated. The editor averages the estimates of the review (red box).

v) The editor computes the acceptance probability based on the averaged estimates. She then draws a random number and the manuscript is either accepted or rejected. In case the manuscript is rejected it goes back to step iii).

Figure 1. A flowchart of the classic setting of peer review.

While the editorial rejection system is a realistic variation on the classic setting, I am also examining a more radical alternative. In fact, the modeling framework can accommodate very different systems. In the “bidding setting” (Appendix), when an author produces a new manuscript, she will submit it to a first pool of manuscripts (e.g., a pre-print archive). However, to be able to submit one manuscript, the author must choose three manuscripts already in the pool for review. Therefore, more productive authors are also the more active reviewers. In the model, authors choose to review three of the manuscripts among those most suited for them. Once a manuscript in the first pool accrues three reviews, it is revised (increase in quality and novelty), and the reviewers are asked for a second evaluation. Then, the manuscript is moved to a second pool (ripe manuscripts). Every month, the editors of the journals evaluate the ripe articles. If an editor wants a manuscript for her journal, she will bid on it. At the end of the month, authors receive all the bids for their manuscripts in the second pool. In the case of more than one journal bidding on her manuscript, the author will choose that with the highest impact. If no journals bid on a manuscript, the author abandons it. Note that with “bid” here I mean simply the intention of a journal to publish the manuscript.

Results

I implemented the three settings using 50 journals and 500 authors. To produce a fair comparison, the same authors and journals are used in all the simulations. I ran the models for 10 years and tracked all manuscripts produced in this period (Appendix).

A first coarse comparison can be done among the number of papers published. In the classic setting, the 500 authors produced 15 233 unique manuscripts in 10 years. Of these, 10 287 have been accepted for publication (67.6%), while 4 936 have been abandoned after having been rejected by 5 journals. When editorial rejection is implemented, only 41.5% of the 14 984 manuscripts are accepted for publication, while 8 763 are abandoned after being rejected 5 times. Finally, in the bidding setting, 97.6% of the 14 984 manuscripts found their way into publication, while 356 were abandoned.

Also the number of reviews is extremely different. In the classic setting, on average 9.64 reviews per manuscript have been performed (the mean drops to 7.07 considering published manuscripts only), while in the editorial rejection, the number is 4.93 considering all manuscripts and 4.71 considering only published articles. Clearly, in the bidding setting, each manuscript receives exactly three reviews.

7.6% of the authors published more in the editorial rejection setting than in the classic (Fig. 2i), while

around 89.4% of the authors published more in the bidding setting than in the classic (Fig. 2i). Also, all the authors accrued more impact (computed summing the impact of the journal for each of the author’s articles) in the bidding setting, while only 13.6% had higher impact in the editorial rejection versus the classic setting (Fig. 2ii). Surprisingly, the average impact for each author was also highest in the bidding setting in all cases.

The time to publication is also very different among the three models (Fig. 2iii): the average is 22 months in the classic (1st Qu.:11, 3rd Qu.:30), 16.81 in the editorial rejection (1st Qu.:11, 3rd Qu.:22) and 8.65 in the bidding setting (1st Qu.:7, 3rd Qu.:10).

Which journals publish more articles in the various settings? In Fig. 2iv, I show how the publications in the three settings are partitioned among journals. I divided the journals in quartiles according to their “impact”. The bottom 25% of journals published more than half of the articles in the classic and editorial rejection settings, while only 20% of the articles in the bidding setting. Accordingly, bidding is the setting in which top journals publish a larger share of the articles. In all cases, the results for the editorial rejection model are intermediate between the classic (skewed toward lower-impact journals) and the bidding setting (skewed toward higher-impact).

Clearly, one factor any good system should control strongly is the quality of published work. If we measure the average merit of manuscripts as the product of their q and n values (Fig. 2v), we find that merit is higher for published papers in the editorial rejection model (mean: 0.407), intermediate for the classic (mean: 0.367) and lower for bidding (0.322). This is due to two effects: a) in the classic and editorial rejection settings, a larger fraction of manuscripts is abandoned, while in the bidding setting only extremely scarce articles are abandoned, and b) the number of revision rounds in the classic and editorial rejection settings is higher, and each revision improves the merit of the manuscript (Appendix).

This would mean that the classic and editorial rejection settings are better at controlling for what we may call “Type II error”: fewer manuscripts with scarce merit are accepted (Fig. 2v). As typically happens, there is a cost associated with it, namely rejecting manuscripts with merit (“Type I error”; Fig. 2v). The Type I error can be very high: in the classic setting, 23% of the manuscripts that are abandoned have more merit than the average published manuscript (Fig. 2v). Similarly, in the editorial rejection the fraction is 26%, while it is less than 0.1% in the bidding setting. To address this point more specifically, I took the top 25% (according to merit) of all the manuscripts for the three systems; of these, only 56% are published in the editorial rejection, 78% in the classic setting, while all but one are published in the bidding setting. Moreover, the impact of the journals in which the top manuscripts are published is

Figure 2. Results of the simulations. i) Number of publication for each author in the classic setting (x-axis) compared to editorial rejection (blue) and bidding setting (purple). ii) Impact for each author (sum of the impact of published articles) in the classic vs. the other settings. iii) Time to publication in the three settings. iv) Distribution of published articles according to the impact of the journals. The journals are divided in four quartiles according to their impact. v) Manuscripts’ “merit” (i.e., product of q and n) for the three systems. Manuscripts are divided in published (Pub) and abandoned (Ab). vi) Impact of the top 25% articles (according to merit) in the three settings.

higher in the bidding setting, while low in both the classic and the editorial rejection (Fig. 2vi).

In the classic and editorial rejection settings, abandoned manuscripts can be reduced by allowing the authors to submit to a greater number of journals before abandoning the manuscript. I repeated the analysis, allowing up to ten submissions for each manuscript. This led to higher publication rates, but decreased the average quality, and increased the number of reviews (data not shown).

Finally, the core of the peer review process for journals is a matching process: exciting, solid and novel contributions should end up being published in high-visibility journals, while less interesting or very specific articles should be found in low-impact journals. In the models, we can measure the correlation between the merit of published articles and impact of the corresponding journals. Surprisingly, in the classic, and editorial rejection settings, we find low correlations (0.34 and 0.38 respectively), while the bidding setting guarantees a much higher correlation (0.72).

Discussion

This exercise represents a highly idealized peer-review process for journals. All the reviews are useful, there is no bias, and each revision increases the overall quality of the manuscript, authors always revise their manuscript according to the referees' comments, reviewers are always punctual, and editors absolutely objective. Moreover, and even less realistically, authors can estimate exactly a) their manuscript value (t, q, n) and b) the probability of having the manuscript accepted in any journal. Note also that if the cost associated with a submission is very small, then authors are practically encouraged to always submit to top journals, regardless the probability of acceptance, rendering the system even more inefficient. The classic and editorial rejection settings can therefore be seen as best-case scenarios for the current system of peer review.

Results suggest that even small modifications, such as the adoption of editorial rejection, can lead to profound changes in the speed and quality of publication and on the goodness of the match between journals and manuscript. The editorial rejection setting shows an improvement in the time to publication, burden for referees and control for quality. The main cost is paid in terms of higher rejection rates, including the rejection of extremely good articles ("Type I error"). Also, authors in the editorial rejection setting are suffering from lower publication numbers and lower impact.

Both the classic and the editorial rejection settings embed a strong stochastic component, which leads to a poor match between quality of journals and quality of manuscripts.

In the bidding setting, publication is faster, reviews are efficiently distributed and proportional to productivity, authors publish more and in better journals, better journals publish more articles, and there is an overall good match between journals and manuscripts. The price to pay in this case is a higher "Type II error" and more work required to the editors.

Although the bidding setting may seem a radical change from current practices, some journals are adopting similar systems. For example, the Berkeley Electronic Journals in Theoretical Economics are three connected journals: whenever a manuscript is submitted, it is evaluated simultaneously by the three journals, which are ranked in tiers. However, acceptance is possible in only one of the three. To make things even more strikingly similar, when one author submits a manuscript, she agrees to act as a referee for two articles in the future or she can avoid refereeing altogether by paying a given amount of money. Other Berkeley Electronic publications are based on the same system. This multiple-journals submission could be easily implemented by other families of journals, such as the Nature or the PLoS family, or by scientific societies publishing several journals.

There could be "social" advantages to a system inspired to the bidding setting. Recently, several journals started providing "Open Access" (OA) content. Typically, authors have to pay a fee to have their articles published as OA. In the bidding system, OA would be free and guaranteed, as anybody could access the manuscripts in the two pools (as it happens now, for example, for the arXiv.org). Scientists in developing countries and institutions that cannot afford subscription costs would be on the cutting edge of science as are their colleagues in more fortunate situations. In fact, they could not only access articles that have already been sanctioned by journals, but also the manuscripts in the pools, i.e., those that will be published in the future. One could imagine that a scientist could read published articles in her discipline, "ripe manuscripts" in her sub-field, and even submitted manuscripts on specific topics. A parallel can be drawn with the release of open-source software: normal users will install "stable" versions, but those who want to experience new features can install "unstable" releases. Another advantage for the authors publishing in the bidding system is that their work becomes visible as soon as the manuscript is submitted, speeding up the process of accruing citations and settling eventual priority disputes. This "early view" effect could also benefit journals, as manuscripts could accrue more citations and faster (Moed 2007). Initiatives like the arXiv.org and other preprint servers go in this direction.

Finally, the bidding system already embeds some of the changes that have been proposed for the current

peer-review system (Campanario 1997), e.g., the creation of a database of reviews common to all journals (Martinez-Abraín 2009), and the idea of a “Community-based Assessment of Review Material” (Carmi and Koch 2007).

This alternative system, however, could give rise to different concerns. For example, forcing the authors to review manuscripts could be counter-productive when there are no suitable manuscripts in the first pool. Also, manuscripts of less interest might not receive reviews for a very long time. Finally, nothing prevents the abuse of the system, where authors in cliques exchange positive reviews. Editors could request additional “traditional” reviews for manuscripts whose reviewers seem to be in a conflict of interest, but this would slow down the system. Moreover, a peer-review system must be financially viable for publishers. Nothing in the simulations proves that this would be the case for the bidding system.

The bidding setting seems better for scientists as authors: they publish more, faster, and in better journals. It is also the best system for scientists as referees: the burden of reviewing articles is minimized. However, because of the high acceptance rate, it is the worst system for scientists as readers: navigating the enormous number of publications becomes more difficult. The only positive result is that the high correlation between journal impact and article merit means that readers are more likely to find good articles in good journals when compared with the other settings.

These considerations relate to the main problem of peer review: what is the ultimate goal of the process? In fact, in this exercise there is no clearly stated goal function we can try to maximize (e.g., using model selection). Therefore, we can show the differences among systems, but cannot advocate the adoption of this or that setting. Funding agencies, committees for promotion and hiring, national governments and research institutions all do contribute to creating a set of expectations and pressure points that strongly influence the publishing behavior of scientists. Therefore, it is apparent that changing for example tenure procedures could have a large impact on the peer-review system. Recently, many articles dealt with the problem of measuring scientific productivity. It is clear that the adoption of certain practices (e.g., promoting the publication of few relevant manuscripts rather than dividing the work into “minimal publishable units”) could alter dramatically the landscape and the problems connected with peer review.

One important point that has not been discussed here is that it is critical to have incentives in place for referees to return timely, constructive, and useful comments. In the idealized models presented here, all reviews are useful and the authors strengthen their manuscripts by

answering the comments. Unfortunately, this is not what happens in the real world.

There is an infinite number of ways to implement a peer-review system, and this work is a first attempt to tackle the evaluation of alternatives in a quantitative way. I have examined only three options, but with further developments, the model can accommodate any variation on the theme. For example, the simulations could be made more realistic by embedding multi-authored articles, including a reputation trait to the authors, allowing referees to decline invitations to review. Also, journals are fixed in rank and have unlimited pages. Modifying these features would increase the realism of the models.

The central idea of this work is, however, that of modeling the process rather than proposing any specific model. I believe that this modeling effort can result in the identification of bottlenecks and trade-offs that are embedded in the various systems. Clearly, peer review is a very complex system that cannot be fully described using simple models. However, simple models have the virtue of highlighting the basic dynamics, features, and limitations of a given system, leading the way to its improvement. As scientists, we want to promote publishing practices that are transparent, efficient, and unbiased. The main goal of this work is to show that there is a quantitative, model-based way to select among candidate systems.

Acknowledgments

E. Baskerville, A.E. Budden, A.M. de Roos, E. Hackett, P.J. Lamberson, C. Lortie, C. Melian, J. Parker, J. Ranganathan and K. Woods provided valuable suggestions and comments. Part of this work was carried out when S.A. was a postdoctoral associate at the National Center for Ecological Analysis and Synthesis, a center funded by National Science Foundation grant DEB-0072909, and the University of California, Santa Barbara. This work was funded by NSF SciSP # 1042164.

Special Issue Editors

Christopher Lortie – lortie@yorku.ca
York University

Jarrett E. Byrnes – byrnes@msi.ucsb.edu
University of Massachusetts Boston

References

Aarssen, L., Tregenza, T., Budden, A., Lortie, C., Kori-
cheva, J. and R. Leimu. 2008. Bang for your buck:
Rejection rates and impact factors in ecological
journals. *Open Ecology Journal* 1:14–19. [CrossRef](#)

- Bentley, P. 2009. The game of funding: modelling peer review for research grants. Proceedings of the 11th annual conference companion on Genetic and evolutionary computation conference', ACM. Pages 2597–2602.
- Bornmann, L. and H. Daniel. 2006. Potential sources of bias in research fellowship assessments: effects of university prestige and field of study. *Research Evaluation* 15:209–220. [CrossRef](#)
- Bornmann, L. and H. Daniel. 2008a. Selecting manuscripts for a high-impact journal through peer review: A citation analysis of communications that were accepted by *Angewandte Chemie International Edition*, or rejected but published elsewhere. *Journal of the American Society for Information Science and Technology* 59:1841–1852.
- Bornmann, L. and H. Daniel. 2008b. The effectiveness of the peer review process: Inter-referee agreement and predictive validity of manuscript refereeing at *Angewandte Chemie. Angewandte Chemie International Edition* 47:7173–7178. [CrossRef](#)
- Buchanan, M. 2010. Come the revolution. *Nature Physics* 6:2. [CrossRef](#)
- Campanario, J. 1997. The journal scout. *The Scientist* 11:9.
- Carmi, R. and C. Koch. 2007. Improving peer review with CARMA. *Learned Publishing* 20:173–176. [CrossRef](#)
- Cintas, P. 2004. Confidential reports may improve peer review. *Nature* 428:255. [CrossRef](#)
- Nature Editorial. 2006. Quantified: Peer review. *Nature* 441:xiii. [CrossRef](#)
- Hauser, M. and E. Fehr. 2007. An incentive solution to the peer review problem. *PLoS Biology* 5:e107. [CrossRef](#)
- Hochberg, M., Chase, J., Gotelli, N., Hastings, A. and S. Naeem. 2009. The tragedy of the reviewer commons. *Ecology Letters* 12:2–4. [CrossRef](#)
- Jaffe, K. 2006. Peer review could be improved by market forces. *Nature* 439:782. [CrossRef](#)
- List, A., Zhang, S., Tobias, S., Hughes, D., Weiss, N. and W. Laurance. 2006. Reviewers' reports should in turn be peer reviewed. *Nature* 442:26. [CrossRef](#)
- Martinez-Abraín, A. 2009. Improving the efficiency of manuscript selection. *Animal Biodiversity and Conservation* 32: 49–50.
- McPeck, M., DeAngelis, D., Shaw, R., Moore, A., Rausher, M., Strong, D., et al. 2009. The golden rule of reviewing. *The American Naturalist* 173: E55–58. [CrossRef](#)
- Moed, H. 2007. The effect of "open access" on citation impact: An analysis of arxiv's condensed matter section. *Journal of the American Society for Information Science and Technology* 58:2047–2054. [CrossRef](#)
- Neff, B. and J. Olden. 2006. Is peer review a game of chance? *BioScience* 56:333–340. [CrossRef](#)
- Schroter, S., Black, N., Evans, S., Carpenter, J., Godlee, F. and R. Smith. 2004. Effects of training on quality of peer review: randomised controlled trial. *British Medical Journal* 328:673. [CrossRef](#)
- Smit, C. 2006. Peer review: time for a change? *Bioscience* 56:712–713. [CrossRef](#)
- Smith, R. 2006. Peer review: a flawed process at the heart of science and journals. *Journal of the Royal Society of Medicine* 99:178–182. [CrossRef](#)

Appendix

Author. Each author is defined by three beta distributions, T , Q and N . Each distribution X requires two parameters: α_x and β_x . If both parameters are equal to 1, the distribution is uniform $U[0,1]$. The beta distribution is quite flexible, so that various shapes can emerge. For all the simulations presented here, I created 500 authors according to three types. Fifty were “broad interests” authors whose parameters were extracted at random from uniform distributions: α_T , β_T from $U[1,5]$, α_Q and α_N from $U[50,100]$, β_Q and β_N from $U[5,10]$. These authors work on a broad range of topics with very high technical quality and novelty. I then added 150 “specialist” authors: α_T , β_T from $U[10,100]$, α_Q and α_N from $U[5,10]$, β_Q and β_N from $U[1,5]$. The remaining 300 authors were assigned values using $U[1,10]$ for α_T , β_T , α_Q and α_N , while using $U[5,10]$ for β_Q and β_N .

All the authors have 1/4 probability of producing a manuscript at each month, so that the number of manuscripts for each author is going to be binomially distributed with an average of 3 manuscripts/year. All manuscripts are single-authored. In the classic and editorial rejection settings authors re-submit a manuscript until it has been rejected five times, in which case they abandon the manuscript as not publishable.

Journal. Also journals are defined by three beta distributions. I created a list of 50 journals using the following parameters: 5 “broad interest” journals (α_T and β_T from $U[1,5]$, α_Q and α_N from $U[50,100]$, β_Q and β_N from $U[5,10]$); 15 “specialist” journals (α_T and β_T from $U[10,100]$, α_Q and α_N from $U[5,10]$, β_Q and β_N from $U[1,5]$); 30 “normal” journals (α_T , β_T , α_Q , α_N from $U[1,10]$, β_Q , β_N from $U[5,10]$).

Journal impact depends on its distribution. It is computed as I (see Box 1, Equation 1), where $B(\alpha, \beta, x)$ is a beta distribution and z is a measure of the density of a beta distribution around x . Note that in this case x is simply the mean of the beta distribution. If $x+0.1 \leq 1$ and $x-0.1 \geq 0$, then z is simply obtained by integrating the probability density function, z (see Box 1, Equation 2). However, in order to ensure that the computation of

the impact is fair for journals whose average T is close to 0 or 1, z must be redefined to depend on the value of x (see Box 1, Equations 3 and 4).

This definition of z ensures that journals with mean topic close to 0 or 1 are not penalized. Therefore, the impact of a journal is simply the product of the expected value for Q and N (first two terms) divided by the degree of specialization of the journal. If most of the density is concentrated around the mean (highly specialized journals), then the impact will be lower; if instead the density around the mean is low (the minimum is 0.2 for a uniform distribution, i.e., a completely generalist journal) then the journal will have higher impact.

Classic Setting: submission. In the classic setting, authors maximize the expected impact of their submission. If a journal X has impact I_X and the probability of acceptance in the journal (see below) for an article whose values are (t, q, n) is p_X then the “score” of the journal is $p_X I_X$. Authors rank the journals according to the score for their manuscript and submit to the highest ranked first. The submission happens 1 month after the manuscript is created.

Classic Setting: reviewers’ choice. Journals choose the reviewers for a manuscript according to their expertise. The score of a reviewer X , whose T distribution is defined by α_T^X and β_T^X for a manuscript whose topic is t is $1/z(\alpha_T^X, \beta_T^X, t, 0.1)$ where z is the function introduced above. Reviewers who are more familiar with the topic of the manuscript will therefore receive a higher score. Reviewers are sorted according to their score (descending order) and three reviewers are sampled from the top 20.

Classic Setting: reviews. Every month each author completes each of the assigned reviews with probability $1/2$. The review process is simply an estimate of the “true” values of a manuscript (t, q, n) . The estimate has an error that is inversely proportional to the familiarity of the reviewer with the topic of the manuscript. We define δ as $\delta = (1 - z(\alpha_T, \beta_T, t, 0.1))/2$ where the α_T and β_T are the parameters for the reviewer, t is the real topic value for the manuscript to be reviewed and z is the function introduced above to compute the density of a beta distribution around a given value (t). The estimate t_1 of the topic for the first reviewer is extracted from the uniform distribution $U[\max(t - \delta, 0), \min(t + \delta, 1)]$. In the same way q_1 is extracted from $U[\max(q - \delta, 0), \min(q + \delta, 1)]$, and n_1 from $U[\max(n - \delta, 0), \min(n + \delta, 1)]$. The three reviewers follow the same procedure, and finally the editor estimates are obtained by averaging the three estimates for the reviewer: $t_e = (t_1 + t_2 + t_3)/3$ and so forth.

Changing the way the editors estimate the triplet—for example using the geometric average, or using the

minimum—to mimic the fact that negative reviews impact a decision more than positive ones, does not change the results presented here qualitatively.

Classic Setting: revision. When a manuscript is revised, its quality and novelty are improved. The improvement tends to be smaller when the manuscript has been revised several times. The quality of the revised manuscript q' is computed as $h(q, 0.1, k)$ where k is the revision number and $h(a, b, c)$ is defined by Equation 5 (Box 1). The manuscript can therefore improve up to 10% after the first revision, but the subsequent improvements will be smaller and smaller. The same happens with the novelty. After the manuscript has been revised, the reviewers give a second estimate. Both the revision and the second round of reviews are instantaneous.

Classic Setting: decision. The probability of acceptance for the journal X with parameters $\alpha_T^X, \dots, \beta_N^X$ of the manuscript defined by the triplet (t_e, q_e, n_e) is defined by Equation 6 (Box 1). Once the probability of acceptance for a manuscript has been computed, the editor draws a random number from $U[0,1]$ and if the number is lower or equal than the probability, she accepts the article for publication, otherwise she rejects the manuscript. The decision for one manuscript happens the month after the reviews are completed. This acceptance function is used by each author to decide where to submit the manuscript: the author applies the function to her manuscript for all the journals, which then are ranked as explained above.

Editorial rejection: summary of algorithm. The editorial rejection setting is identical to the classic setting, with the addition of an extra step before the manuscript is sent to the reviewers. The editor of the journal screens all the manuscripts and does a quick review. The review done by the editor works exactly like that of the referees’ explained above with the differences that the editor uses the journal’s T distribution and that the error associated with the evaluation is doubled (to mimic the fact that the review is quick and the editor not a specialist in the specific topic of the manuscript). After this screening, the editor computes a probability of acceptance: if the probability is lower than 30%, the manuscript is rejected without review. All these operations are done instantaneously when the manuscript is submitted, to mimic the very fast turnaround typical of top journals.

Bidding Setting: summary of algorithm. In the bidding setting, a manuscript is created as in the classic setting. The following month, however, instead of being submitted to a journal, the manuscript is sent to the first pool. When the author submits a manuscript, she has to

commit to choose three manuscripts for review the next month (this is required in order to keep the minimum time before acceptance equal for all settings). The author chooses the manuscript to review among the 20 that are best suited for her profile (i.e., the 20 that have a topic that is most familiar to the author).

Once a manuscript has been reviewed, it is revised, the reviewers give a second opinion (instantaneous), and the manuscript is moved to the second pool. The

following month, all the journals run the acceptance procedure on all the manuscripts in the second pool. If they are willing to accept any one of them, they bid on it (i.e., they accept the manuscript for publication provided the author wants to publish it in their journal). At the end of the month, authors choose the journal with the highest impact among those that bid on it for their manuscripts. If no journal bids on a given manuscript, its author abandons it.

Box 1. Definitions for variables and distributions described in the Appendix.

Equation 1.
$$I = \frac{\alpha_Q}{\alpha_Q + \beta_Q} \frac{\alpha_N}{\alpha_N + \beta_N} \frac{1}{z(\alpha_T, \beta_T, \alpha_T/(\alpha_T + \beta_T), 0.1)}$$

Equation 2.
$$z(\alpha_T, \beta_T, x, 0.1) = \int_{x-0.1}^{x+0.1} B(\alpha_T, \beta_T, t) dt \quad 0.1 \leq x \leq 0.9$$

Equation 3.
$$z(\alpha_T, \beta_T, x, 0.1) = \int_{x-0.1}^1 B(\alpha_T, \beta_T, t) dt + \int_0^{0.1+x-1} B(\alpha_T, \beta_T, t) dt \quad x \geq 0.9$$

Equation 4.
$$z(\alpha_T, \beta_T, x, 0.1) = \int_0^{x+0.1} B(\alpha_T, \beta_T, t) dt + \int_{0.9+x}^1 B(\alpha_T, \beta_T, t) dt \quad x \leq 0.1$$

Equation 5.
$$h(a, b, c) = a + \frac{b}{c}(1 - a)U[0, 1]$$

Equation 6.
$$p(\text{accept}|X, t_e, q_e, n_e) = z(\alpha_T^X, \beta_T^X, t_e, 0.1) \cdot \int_0^{q_e} B(\alpha_Q^X, \beta_Q^X, y') dy' \cdot \int_0^{n_e} B(\alpha_N^X, \beta_N^X, y'') dy''$$