

ARTICLE

Academic Journals, Incentives, and the Quality of Peer Review: A Model

Kevin J. S. Zollman¹ , Julian García²  and Toby Handfield³ 

¹Department of Philosophy, Carnegie Mellon University, Pittsburgh, PA, USA, ²Faculty of Information Technology, Monash University, Clayton VIC, Australia and ³School of Philosophical, Historical and International Studies, Monash University, Clayton VIC, Australia

Corresponding author: Toby Handfield; Email: toby.handfield@monash.edu

(Received 23 October 2022; revised 21 March 2023; accepted 25 April 2023; first published online 02 May 2023)

Abstract

We model the impact of different incentives on journal behavior in undertaking peer review. Under one scheme, the journal aims to publish the highest-quality papers; under the second, the journal aims to maintain a high rejection rate. Under both schemes, journals prefer to set very high standards for acceptance despite allowing significant error in peer review. Under the second scheme, however, in order to encourage more submissions of mediocre papers, the journal is incentivized to make its editorial process less accurate. This leads to both worse peer review and lower-quality articles being published.

1. Introduction

Philosophers of science are increasingly using credit economy models to understand the relationship between the individual incentives that scientists face and the epistemic aims of science. This modeling framework elucidates when the institutions of science facilitate or hamper our broader goal for science as a social enterprise (e.g., Heesen 2018; Strevens 2012; Zollman 2018). The vast majority of these models have focused on individual scientists as the actors of greatest interest. In this article, we turn our attention to a different actor: journals.

Some time ago, the principal purpose of journals was the physical transmission of papers. They were the medium by which new scientific results were circulated. With preprint servers, mailing lists, and the like, journals are no longer needed for the physical distribution of papers. However, journals continue to serve a social-epistemic purpose; they filter, sort, and certify academic papers. Both specialists and nonspecialists rely on journals to assist in judging the quality of a paper. First, the fact that a paper has been published in a reputable journal is taken as a mark of quality: it must have passed through peer review. Second, the ecology of journals might help to sort papers by quality. Papers published in, say, *Nature*, are thought to be of higher quality than papers published in a less prestigious journal.

This certification purpose serves an important social-epistemic function. There is far too much to read; specialists need a way to focus attention. Nonspecialists need this certification even more. A policymaker might not be able to distinguish good from bad science, but they must find the best work on a given topic.

Earlier literature has addressed the degree to which peer review serves this social-epistemic role (e.g., see Heesen and Bright 2021). Although important, we will sidestep this question and focus on a closely related one: Are journals, thought of as individual actors, incentivized to do the best they can? We suppose that peer review could, in principle, always be improved but that each improvement requires an increasing investment of time and effort by the journal. What determines how much effort the journal will ultimately expend?

Given their importance as certification bodies, it might be surprising that journals largely self-regulate in a *laissez-faire* way. No central authority regulates their standards. Some journals are owned by professional societies that exercise some control—although even then, most control is given over to the editor. Many journals are owned by for-profit companies that have no direct stake in the quality of scientific output; it only affects them insofar as it affects the price they can charge (Shideler and Araújo 2016; Björk and Solomon 2015). Even if the journal owner has little financial stake in the success of the journal, it stands to reason that the editors and owners would prefer to be involved with a journal judged to be better than one judged inferior.

What counts as a “good journal” is somewhat undefined and subject to social norms (Saha et al. 2003; Lowe and Locke 2005; Lee et al. 2002). In many fields, journal quality is determined largely by the quality of the papers published therein. Sometimes it will reflect “superficial” considerations, such as the presence of color figures and the quality of the graphic design. More commonly, metrics such as impact factor and other citation indices are calculated by averaging over the citations of published papers.

However, in some fields, the *selectivity* of a journal contributes to its reputation for quality—particularly given that it can be hard to observe the quality of the individual articles. If a journal is very selective, rejecting many submissions, it is likely to be judged better in quality than one that publishes a larger fraction of the papers it receives.¹

Given the important social-epistemic function of journals, and given that they are almost always self-regulating, we examine the impact of different incentives on journal behavior. We develop a series of models to address two interrelated questions.

¹ There is limited evidence, in some fields, of an actual correlation between rejection rate and certain objective measures of quality, such as impact factor (e.g., Aarssen et al. 2008), although there are also several cases where no correlation has been found or the evidence is inconclusive (e.g., Schultz 2010; Abrahams 1977). For philosophy, we analyzed rejection-rate data for journals with at least 10 observations in the American Philosophical Association (APA) journal survey (data obtained as of August 17, 2022, available here: <https://doi.org/10.26180/20499582.v1>) and found a correlation with the Scimago Journal Rank of 0.441 ($N = 80$, $p < 0.001$). Even if there is no correlation, it is widely *believed* that a journal's rejection rate is some indication of the journal's quality. See Regazzi and Aytac (2008) and Egbert (2007) for observations in this vein. Anecdotally, we also see evidence that editors sometimes value having a high rejection rate as a proxy for quality. For instance, in one editorial (Geddes 2010), the editor proclaims the “most happy” news that the journal he edits has recently seen its rejection rate increase by more than 30% (see also Loison et al. 2006).

First, are journals incentivized to make accurate decisions about the quality of papers submitted to them? Second, does it make a difference whether we judge a journal by its quality of published articles versus its selectivity? That is, would we expect journals that strive to improve on one of these dimensions to be systematically better in some sense compared to a journal that is incentivized on the other?

Our article argues that although journals are sometimes incentivized to maintain high-quality peer review, this often fails for two reasons. First, a journal might use self-selection by authors as a substitute for peer review. If the author of a bad paper chooses not to submit, then the journal need not worry about peer reviewers recognizing its failures. This has a complicated relationship with journal incentives, which we discuss later in the article. We argue that this use of self-selection can be epistemically productive in some sense but might have negative consequences as well.

Second, journals that are incentivized by their selectivity have less desirable properties as collective epistemic resources. They have an incentive to *discourage* self-selection because a paper that is not submitted cannot be rejected. This results in a strange process whereby journals make peer review worse in an attempt to induce bad papers to submit, but they maintain sufficiently good peer review to ensure that a large proportion of those bad papers will probably be rejected. We present these results through a series of game-theoretic models in the sections that follow.²

2. Nonstrategic author model

We will begin with a simple model of the journal-selection process. There is a universe of papers that will submit to a single journal. We'll assume that each paper has a quality, q , that is represented by a real number in $[0, 1]$. We remain agnostic about what this quality represents. It could represent something intrinsic to the paper, such as the epistemic quality of the work. It might also be a judgment about something extrinsic to the paper, such as the number of citations the paper will receive. In an empirical field, it might be the probability that the paper replicates. In a theoretical field, it might represent the importance of the theoretical advance. The only significant assumption is that quality can be represented on a single dimension.

Although it won't matter until section 3, we also assume that the author is aware of the quality of her paper: she knows how good her paper is. This is an idealizing assumption for the purposes of this article, but we expect our results could be generalized to any setting where the author knows substantially more about their paper's quality than the journal.³

For simplicity, we will assume that there is a paper for every real number in $[0, 1]$. This represents a setting where the papers are uniformly distributed over that range.

² There may be other gatekeeping activities that could be interpreted in light of this model, for example, external firms vetting candidates for hiring or college admissions screening potential students. We do not explore those possibilities explicitly in this article.

³ For the nonstrategic author model, this assumption is unnecessary because the authors make no decisions. For the later models, however, this assumption is critical. This stands in contrast to other models of peer review, where the reverse assumption is made: authors don't know the quality of their papers, but journals can—with effort—determine it (Heintzelman and Nocetti 2009; Ellison 2002). In one respect, our article can be seen as an exploration of what happens when this asymmetry is reversed.

In our first model, we will assume that every paper is submitted to the journal regardless of any decision made by the journal. Hence, the authors are behaving *nonstrategically*.

Journals make two decisions. First, they decide on a *quality threshold*, denoted by Q_T , which is the minimally acceptable paper. A journal can say, “We’re only going to accept papers that are in the best 10% of the field” ($Q_T = 0.9$), or “in the best 1%” ($Q_T = 0.99$), or “only the very best paper” ($Q_T = 1$). At the other extreme, they might say, “We’ll publish anything” ($Q_T = 0$).⁴ The second decision is the quality of peer review, which determines the probability that any particular quality paper is accepted or rejected, given the journal’s quality threshold.

The peer-review process is modeled as a noisy quality-estimation process. That is, the peer-review process returns a number that represents the estimated quality of the paper. The estimated quality is $q + e$, where e is normally distributed with mean 0 and variance ε . We model the decision to adopt a given quality of peer review as the decision to set ε to a particular value: high ε entails large errors and thus poor-quality review, and vice versa.⁵

If the paper’s apparent quality $q + e$ is higher than or equal to the journal’s threshold, Q_T , the journal publishes the paper. If the paper’s apparent quality is below the threshold, then the journal rejects. In this idealized scientific world, there is no process to revise and resubmit.⁶

To sum up, the journal makes two decisions: First, what’s the threshold? And second, how much effort will we put into peer review?

We compare two different ways that one can incentivize a journal. We might judge a journal by its quality. In this case, we will consider a journal that wants to maximize the *average quality of papers published* in the journal (call this the *quality-incentivized journal*). We will represent the average quality of papers published in the journal as \bar{q} . The other way journals are incentivized is by selectivity. We represent a journal’s rejection rate as r , and in our second type of model, we assume journals strive to maximize this (call this the *selectivity-incentivized journal*).⁷

Figure 1 shows two panels representing equivalent choices of Q_T and ε from the journal’s perspective. The horizontal axis is the announced quality threshold, Q_T . The vertical axis here relates to the reviewing quality, ε , which is the variance of the journal’s error distribution. At $\varepsilon = 0$, peer review is perfect. The lines in this space

⁴ Our model is a variation of that by Azar (2015) that allows us to analyze strategic questions regarding peer-review quality (which he treats as exogenous) and types of incentives (which he fixes as what we call “quality”). Azar’s model assumes that the journal must publish a certain number of papers (which we do not). In his model, he demonstrates that the threshold strategy is an optimal one among all potential strategies. Instead of proving this in our context, we merely assume it.

⁵ This way of modeling peer review treats it like a black box. The peer reviewers are not agents in the model. The peer reviewers are just like a measurement instrument, with a specified error rate. It is not important now, but critical for later models, that the peer reviewers are not being modeled as Bayesian agents who attempt to infer which papers might be being submitted to the journal.

⁶ For a model that explicitly looks at strategic choices centered around revisions, see (Ellison 2002).

⁷ For our purposes, we calculate both types of incentivization by considering the expected average quality of published papers and the expected average rejection rate. In both cases, we must define what happens when the journal receives no submissions. Because we do not want journals to actively attempt to discourage all submissions, we stipulate that when no papers are published, the average quality is zero, and when no papers are submitted, the rejection rate is zero.

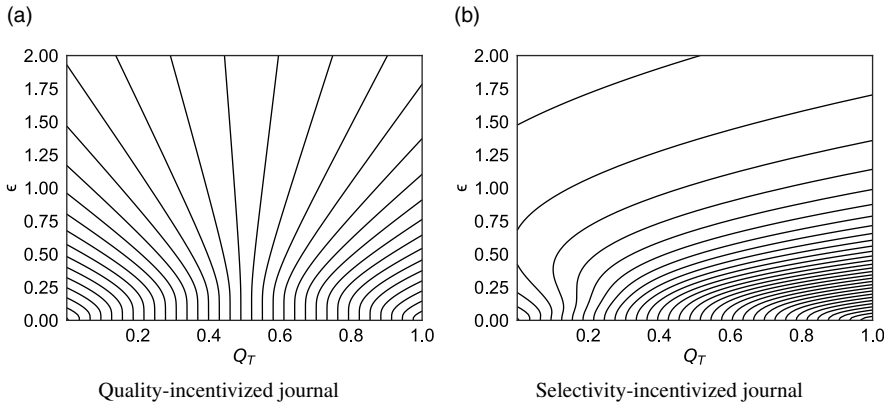


Figure 1. Indifference curves for the journal when choice of Q_T and ε is unconstrained. Within each panel, each line represents choices that yield an equivalent payoff for the journal.

represent lines of equivalent goodness from the perspective of the journal. The right-hand plot is for a journal that cares about its selectivity. From the journal’s perspective, for any target rejection rate, it is indifferent between any points along the relevant line.

The plot on the left-hand side represents the journal incentivized by the average quality of papers published in the journal. This looks a little bit different, of course, because now it’s not concerned about how many papers it’s rejecting. It’s concerned about the quality of the papers it’s accepting. Both panels illustrate one important point: that journals can achieve equivalent results, from either perspective, with several different pairs of choices. They might achieve the same payoff with very good peer review and a slightly lower Q_T or by having bad peer review and a different Q_T . This basic point will underwrite much of what is to come.

So far, there is no cost to any choice the journal makes, so in both cases, the overall optimal behavior for the journal is to choose $Q_T = 1$ (only accept the very best paper) and $\varepsilon = 0$ (have perfect peer review). This will result in a journal that has an average quality of 1 and a rejection rate of 1, both the highest values possible.

In reality, however, there is a cost for high-quality peer review. Identifying appropriate reviewers, soliciting reviews for a paper, and reading and evaluating the quality of the reviews take time and effort. Some journals even pay their reviewers. All of these things are costly. In order to account for this, we must include a cost for the quality of peer review. If we introduce this cost, we get an interestingly different result.

Our two utility functions are as follows:

$$\begin{aligned} u_R(Q_T, \varepsilon) &= r - c_j(\varepsilon) \\ u_Q(Q_T, \varepsilon) &= \bar{q} - c_j(\varepsilon), \end{aligned} \tag{1}$$

where u_R represents the utility for the selectivity-incentivized journal, and u_Q represents the utility for the quality-incentivized journal. The function $c_j(\cdot)$ is arbitrary and represents the cost of improving peer review. We will assume that c_j is decreasing in ε , meaning that the cost increases as peer review gets better.

With that relatively weak assumption, we can prove that a journal incentivized by its rejection rate will want to set $Q_T = 1$ (this is proven in the [appendix](#)). The journal will set the threshold for publication at the highest value it can. It does this because changing Q_T is free, whereas improving peer review is not.

Once the journal sets Q_T at 1, it will then optimize ε to balance the costs and benefits from improving peer review. This will depend on the functional form of $c_j(\cdot)$, but in many plausible cases, it will result in an intermediate value of ε that represents neither perfect nor completely unreliable peer review.⁸

Interpreting these results, we get the following conclusions: the journal wants to claim to be the perfect journal. It claims that it will only accept the very best papers, but it won't go all the way to perfect peer review. It will choose an intermediate quality of peer review that represents the trade-off between the benefits and the cost of its peer-review process.

To say more beyond this observation, we assume a particular functional form for $c_j(\cdot)$. For the remainder of this article, we assume

$$c_j(\varepsilon) = \frac{1}{(1 + \varepsilon)^k},$$

where k is a parameter that allows us to vary the steepness of the cost function. For “low” values of k , the cost increases steeply as peer review becomes better. For “high” values of k , cost increases less steeply, so good peer review is relatively cheaper.

Figure 2 is the optimal error in peer review, from the journal's perspective, for various values of k . Very low k values are not plausible because they lead to uninteresting results.⁹ For what we think of as plausible values of k , which are around 6 and higher, a journal incentivized by its quality will do better at peer review than a journal incentivized by its rejection rate. In all subsequent plots, we set $k = 8$.

When a journal is judged by its rejection rate, it is indifferent about which papers it rejects. If it rejects good papers, the journal is rewarded just the same as it is if it rejects bad papers. However, if we judge a journal by its quality, this is not the case: a journal is punished for rejecting good papers and rewarded for rejecting bad ones. As a result, when there is a cost to peer review, journals judged by their quality will

⁸ In our computational study of this and related models that follows, we assume that the error for judgment quality is normally distributed. In our findings, all journals (regardless of how they were incentivized) set $Q_T = 1$. This might not remain true for all assumptions about how error is generated, but it is nonetheless robust enough for our purposes.

⁹ When k is low (approximately 0.5–4.0), the cost of review is exceptionally high, and the journal adopts the lowest possible quality of review. This corner solution does not represent reality (at least for nonpredatory journals, which we do not aim to capture in our model). If k is too large, then it will be cheap to have very accurate reviewing. It certainly seems unlikely that peer review is “easy” to do well, given an abundance of evidence about its low reliability (e.g., Bornmann 2011; Cicchetti 1991; Snodgrass 2006; Heesen and Bright 2021; Brembs 2018; Campanario 1996; Campos-Arceiz et al. 2015; Cullen and Macaulay 1992; Deveugele and Silverman 2017; Ernst et al. 1993; Howard and Wilkinson 1998; Jackson et al. 2011; Kravitz et al. 2010; Mahoney 1977; Peters and Ceci 1982; Reinhart 2009; Rothwell and Martyn 2000; Rubin et al. 1993; Siegelman 1991; Siler et al. 2015). Many of these articles contain suggestions to improve peer review, but they are all costly to implement. Some have even suggested paying peer reviewers (Engers and Gans 1998). When k is extremely low (approximately less than 0.5), our chosen cost function becomes unrealistically flat, such that the journal can obtain maximal quality peer review for practically the same cost as maximally poor-quality peer review.




Figure 2. Optimal ε value, under two incentive regimes, as a function of k . Note: In these calculations, we cap ε at 1 because that is already a very low reliability of peer review, whereby even the worst paper has approximately a one-in-six chance of being accepted in a journal with the maximum quality threshold.

dedicate more effort to improving peer review. A quality-incentivized journal better serves the epistemically valuable sorting function than a selectivity-incentivized journal.

However, there is still a significant concern about both the quality-incentivized and selectivity-incentivized journals. Both journals are setting $Q_T = 1$. In effect, the journals are saying, “We are the very best journal. Our quality standards are the highest they could possibly be.” But they choose some intermediate value for the quality of peer review, which creates a gap between the quality of papers that are published in the journal and the announced quality of the journal. The journal, in some sense, claims to be enforcing the highest standards—but the claim is somewhat hollow, given the low investment in peer review. In terms of serving a certification function for the public, it might be more beneficial for a journal to adopt a lower quality threshold with more accurate enforcement.

3. Strategic authors

In the first version of the model, we assumed that all authors submitted their papers regardless of the probability of acceptance. In reality, an author who knows the quality of her own paper might choose not to submit to a journal where she thinks her chances of being accepted are low. In order to accommodate this possibility, we now allow authors to decide whether they will submit to the journal.

In reality, authors may not know ahead of time the precise threshold or quality of peer review of the journal; however, for simplicity, we assume this knowledge is

public. Given widespread discussion among academics about journal reputation, as well as citation metrics and published rankings in some fields, authors clearly have some knowledge of journal quality, and that knowledge should be widely shared. We will assume an order of operations as follows: The journal announces its peer-review policy (ϵ) and its threshold (Q_T). This is known by the authors. The authors produce a paper and observe its quality, q . For the moment, we will treat a successful publication as worth utility 1 for all authors. This will be revisited later. Because better papers are more likely to pass peer review, authors of better-quality papers stand to gain more from submitting to a journal, so the expected benefit of submission is increasing with paper quality.

We also assume that rejection comes with a cost (denoted “ c_A ” for *author cost*). That should be no surprise to an academic. Rejection comes with a psychological cost, at least, but also comes with other forms of cost, too. For example, one must write a bespoke cover letter or conform to an arbitrary journal style. There is also opportunity cost. Submitting to one journal forecloses the option to submit to another journal for some period.¹⁰ If the paper is rejected and then takes longer to come out in another venue, the paper might be scooped by someone else, or the topic may no longer be relevant or exciting. Some of these costs, such as the burden of complying with formatting rules, are present whenever an author submits, whereas the psychological cost of rejection and the opportunity cost of having missed out on time to submit elsewhere arise only when the paper is rejected. We ignore this difference as inconsequential for modeling purposes and treat all costs as arising only when the paper is rejected—but for convenience, we will switch between the terminology of *rejection cost* and *submission cost*, depending on the context.

We will assume that the cost depends on (at most) the quality of the underlying paper, and we therefore represent it as a function $c_A : [0, 1] \rightarrow \mathbb{R}$. In this article, we will model cost in two ways. One is a constant cost: regardless of the quality of the paper, the cost of rejection is the same. This cost will be represented by $c_A(q) = c$. Alternatively, the cost might vary with the rejected paper quality; it’s worse to have a paper rejected when the paper is relatively high quality. This cost will be $c_A(q) = c \times q$.

A constant-cost model reflects many of the frustrating difficulties of the submission process. One must format a paper for the journal, write a cover letter, suggest suitable reviewers and/or editors, and so forth. The variable model of costs, on the other hand, reflects opportunity costs. If an author has a low-quality paper, the opportunity cost of having it rejected is relatively low because if they had taken it to a different journal, it probably would have been rejected there, too. Further, having a bad paper rejected is not as bad as having a good paper rejected because having a good paper rejected will mean that citations and other impacts of the paper will be delayed, whereas a sufficiently bad paper will have a negligible impact whenever it is ultimately published. In reality, publishing features some combination of both kinds of costs. For analytic purposes, we treat them separately.

¹⁰ In some publications, such as the typical law review journal, it is less common to require exclusivity of submission (presumably, this is related to the relatively lower status of the people performing the review—usually law students)—thus reducing the opportunity cost for authors. An avenue for future research would be to consider how this is likely to affect author and journal behavior.

Given the announced strategy of the journal and the quality of their own papers, the authors know the probability that their papers will be accepted. The authors are also aware of the cost of submission and will choose to submit their papers if the expected utility of doing so is positive. (We normalize not submitting a paper as utility 0.)

In order to understand behavior in this model, we must consider an equilibrium between the authors and the journal. In the previous section, a journal could count on all authors to submit and would simply maximize its expected quality or rejection rate against that. In the strategic-authors model discussed in this section, the journal must anticipate that some authors may not submit.¹¹

This consequence is most obvious in the context of the rejection-rate model. If the selectivity-incentivized journal sets the threshold very high with very high-quality peer review, authors of bad papers may choose not to submit. In such a case, the rejection rate will go down because the journal does not have the opportunity to reject the bad papers. So, the journal might do better by increasing the chance that bad papers are accepted in order to entice the authors to submit. So far, this is just a claim, but when we analyze the model, this is what we find.

3.1. Constant cost

Suppose the setting where $c_A(q) = c$ for all qualities q , so all authors face the same rejection costs. Many of the same basic facts remain from the nonstrategic model. Journals remain incentivized to set a quality threshold of 1, and then they tweak the peer-review quality to alter the results of the process of submission and review. In the nonstrategic model, this was a simple one-party optimization problem: the journal wanted to set the error in peer review to balance the benefits of a more reliable review process, in terms of either quality or rejection rate, against the attendant costs. In the strategic model, the journal must now anticipate the prospect that some authors will not submit.

To understand the behavior of the authors and the journal in equilibrium, we first observe that if any author does not submit, then all authors of worse papers will *also* not submit because authors of worse papers have even less chance of being accepted and the cost of rejection is constant. This means that self-selection works from the “bottom up” (see the [appendix](#) for a proof of this claim; it is illustrated in [fig. 3](#)).

Bottom-up selection has very different effects on the two differently incentivized journals. From the perspective of the quality-incentivized journal, this self-selection is always good. Unless peer review is perfect, there is always a small probability that a bad paper passes peer review. If the bad paper never submits, the probability of that paper being published drops to zero—thus increasing the expected quality of published papers at no cost to the journal.¹²

Although self-selection is good from the journal’s perspective, it does come with a consequence for the authors. When authors of bad papers are self-selecting out, the journal has less of an incentive to maintain high-quality peer review. The journal no

¹¹ Strictly speaking, we infer the subgame perfect Nash equilibria of this game.

¹² We are not the first to notice that self-selection is beneficial from the journal’s perspective (Azar 2005; Tiokhin et al. 2021). However, these models do not evaluate either the effect on the quality of peer review or the desirability of self-selection for a journal incentivized by selectivity.

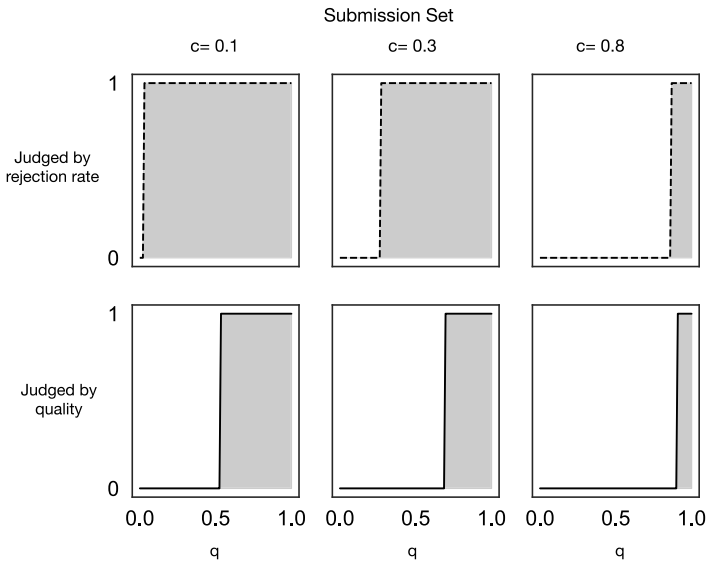


Figure 3. Submission set for a version of the constant-cost/constant-benefit model.

longer has to worry about ensuring that bad papers are not published because the bad papers are not being submitted. As a result, in equilibrium, a journal will have worse peer review than in the situation of the previous section where all authors are submitting. This is illustrated in figure 4 (left plot) by noting that as c increases, the journal chooses a higher ε , thus resulting in lower-quality peer review.

Although we do not model this choice, if the journals were capable of influencing the cost of submission (c in our model), they might prefer to increase it in order to encourage self-selection. This illustrates a point made by Tiokhin et al. (2021): that high costs for submissions might be used as a filtering device. Unlike the conclusion discussed in that article, and as we show in figure 4, this then leads to lower-quality reviewing. In turn, the lower-quality review has a negative epistemic consequence, in that the chance of a better paper being rejected while a worse one is accepted increases. Difficult submission processes come to substitute for reviewing quality. However, evaluated altogether, the quality of published papers goes up, indicating that the epistemic benefit of self-selection outweighs the loss from lower-quality peer review (see the right-hand plot in fig. 4, where \bar{q} increases as c increases for a journal incentivized by quality).¹³

All this is different for the journal incentivized by rejection rates. For a journal incentivized by selectivity, self-selection is a *problem*. If a paper is not submitted, it cannot be rejected, and if it cannot be rejected, it doesn't count toward the journal's

¹³ Azar (2015) includes a constraint that the journal must publish a certain minimum number of papers. He shows that when c increases, the journal must lower its acceptance threshold in order to ensure that the right number of papers are published. This occurs because the journals in his model cannot control the quality of peer review. In our model, the journals keep the same threshold but can manipulate who submits by altering the quality of peer review.




Figure 4. Journal strategy and outcomes in equilibrium for a version of the constant-cost/constant-benefit model. (Left) Error rate adopted by the journal as a function of the author’s rejection cost. Higher rejection costs lead to worse peer review (higher error rates). The small dips that occur near $\varepsilon=1$ are the result of small rounding errors and should not be taken to be meaningful. (Right) Average quality of papers published and journal payoffs as a function of the cost to authors of rejection. For the quality-incentivized journal, published quality and journal payoff are both monotonically increasing with the author’s submission cost. For the selectivity-incentivized journal, quality of published papers exhibits a nonmonotonic relationship with the author’s costs, whereas the journal’s utility is always decreasing in the author’s costs. Observe that the quality of papers published in the selectivity-incentivized journal is never higher than that of the quality-incentivized journal. ($k = 8$).

rejection rate. Thus, there is no reason for such journals to prefer a higher submission cost. Indeed, in this variant of our model, where authors self-select from the bottom up, a journal motivated by selectivity would ideally prefer to set the quality of peer review low enough that even authors of the worst papers think it is worth it to submit. Even with very poor reviewing, however, as the cost of submission gets sufficiently high, authors of the worst papers calculate that the chances of acceptance are too low and won’t submit. Hence, the journal’s utility declines with increasing c .

This illustrates two important incentives that lead journals to choose low-quality reviewing. Journals incentivized by quality can use self-selection to take the place of high-quality peer review. Journals incentivized by rejection rate want to keep reviewing sufficiently poor to combat self-selection. Both are epistemically undesirable, but as we show, the latter leads to worse reviewing than the former.

These conclusions can all be seen in figure 4: see the dotted lines (where journals are incentivized by rejection rate). In the left-hand plot, the journal judged by rejection rate maintains a lower quality of review (higher ε) than one incentivized by quality. In the right-hand plot, the payoff of the journal incentivized by its rejection rate goes down as the submission cost parameter, c , goes up. This occurs because as cost increases, fewer bad papers are submitted (see fig. 3), and the rejection rate is lower.

What are the epistemic consequences of this state of affairs? In particular, what might the effect be on people who are neither authors nor journal editors and rely on journals as a source of information? If a journal is incentivized by rejection rate and could eliminate the cost to authors, it would lead to a journal that was (a) inferior in quality to a quality-incentivized journal with the same costs (compare the solid and

the dashed red lines at $c = 0$) and (b) inferior in quality to a journal incentivized by rejection rate but where the cost to authors is significantly higher. This is illustrated by the dashed red line in figure 4, which represents the average quality of papers published in a journal that is incentivized by its rejection rate.¹⁴

3.2. Variable costs

The previous model reflects the fixed costs that come with journal rejection. There is a cost to filling out forms and formatting the paper to follow submission guidelines. Beyond the time, there are psychological costs that come with receiving a rejection. It is not a bad first approximation to treat these as constant across paper quality.

Some costs, however, vary with respect to the quality of the paper. A low-quality paper may make little impact even when published, so the cost of a delay is small. A high-quality paper may have a significant risk of being scooped or might lose out on early uptake if publication is delayed. In order to model this, we will consider a second model where $c_A(q) = cq$ for some c .

This immediately changes the dynamics of author behavior. Authors of the lowest-quality papers have no reason not to submit because the cost of rejection is zero.¹⁵ Consequently, in this model, self-selection proceeds either from the middle out or from the top down (as illustrated in fig. 5).

It is no longer obvious whether the quality-incentivized journal does better by discouraging submissions. It does eliminate some papers the journal would want to reject, but it also does so by increasing the *proportion* of submitted papers that are quite bad. The journal now faces a more complicated trade-off.

It remains the case that the quality-incentivized journal will rely on some self-selection. As c increases, all journals will, for the most part, reduce their quality of review (increase their ε), relying on self-selection to take the place of quality reviewing. This is illustrated in figure 6. As was the case in the constant-cost/constant-benefit model, the rejection-incentivized journal will choose to have much worse peer review because it wants to incentivize as many paper submissions as possible.

There is one important difference between this model and the previous one that requires discussion. In the previous model, we showed that the selectivity-incentivized journal always preferred lower submission costs to encourage more submissions. In this version of the model, things are somewhat more complicated. In this version of the model, self-selection does not proceed “from the bottom up,” so a journal incentivized by its rejection rate might not always want to encourage submissions.

If a journal can discourage the *best* papers from submitting while holding its quality threshold for acceptance constant, it improves the journal’s rejection rate. Consider a journal where only the very worst papers are being submitted: in this case,

¹⁴ It is worth noting that the case where a journal is incentivized by its rejection rate would choose $c = 0$, and this is not the epistemically worst outcome. Instead, the epistemically worse outcome occurs around $c = 0.2$.

¹⁵ For analytic convenience, we assume that an author who is indifferent between submitting and not submitting will choose to submit.

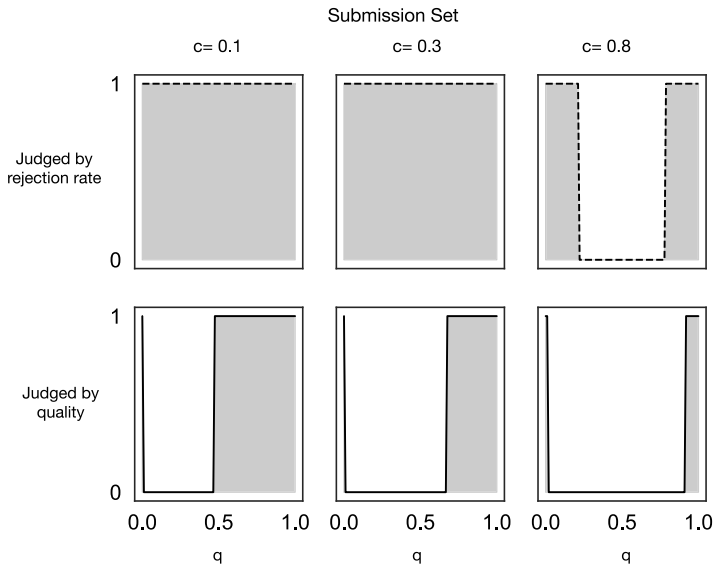


Figure 5. Submission set for a version of the variable-cost/constant-benefit model.

even with relatively poor peer review, the journal can have a rejection rate near 100%.¹⁶ If the *best* paper were then submitted, the journal’s rejection rate might actually go down. So, unlike in the previous version of the model, the journal does not always benefit from encouraging more submissions.

One can see this complicated relationship in the right panel of figure 6, where the cost is related to the journal’s payoff. A journal incentivized by its rejection rate may want to set c either very low or very high. In this particular case, choosing a low c is epistemically superior to choosing a high one because the average quality of published papers is higher in the former case than the latter.

What should we make of these divergent results? As mentioned earlier, we think the real situation features costs of both types: those that are insensitive to the quality of the paper and those that vary with the quality of the underlying paper. However, we think that the journal would, for the most part, have control over those costs that are constant—things like submission charges and formatting requirements.

3.3. Nonconstant benefit

So far, we have modeled the authors as receiving a constant benefit from submitting. Whether the journal is good or bad, the authors are happy to be published there. In the short term, this might be an appropriate model: journal reputations change very slowly, and even if a journal is going “downhill,” an author may gain from the past

¹⁶ Lest the reader think it implausible that a journal would have such a high rejection rate, we recommend Hörner (2019), the annual report of an editor of a well-regarded economics journal that accepted precisely zero of the papers submitted to it in 2018! (Some papers did receive “revise and resubmit” decisions.)




Figure 6. Journal strategy and outcomes in equilibrium for a version of the variable-cost/constant-benefit model. ($k = 8$).

reputation of the journal. However, we should ensure that our results do not depend on this and look at versions of the model where the current quality of the journal determines the benefit to the authors of publishing there.

For this new model, we assume that the authors receive a positive payoff of \bar{q} , the average quality of published papers in the journal. They then must pay either a constant cost (as in section 3.1) or a variable cost (as in section 3.2).

Results for these models are presented in the appendix. Our earlier conclusions remain largely robust to this modification. The version with constant cost and variable benefits looks much like the model with constant cost and constant benefits. The version with both variable cost and variable benefit looks like the model with variable cost and constant benefit. We therefore conclude that in this model, the structure of costs (both costs of reviewing for journals and costs for authors) drives the interesting results.

4. Discussion of assumptions

We have developed a series of models to analyze how well a laissez-faire system will incentivize high-quality peer review. We have identified several impediments that will be summarized in the next section. Before we do that, however, it is important to discuss the ways our model is limited by its assumptions.

First, we have chosen two extreme ways to model the cost of submission to authors: constant across paper quality or variable as an increasing function of paper quality. In reality, there is a complex web of costs in between those two extremes. We believe that our results would be robust to more complicated hybrid cost functions, but this was not tested.

On the journal end, we have treated the journal as incentivized by either the quality of its papers or its rejection rate. It is unlikely that any journal is exclusively incentivized by its rejection rate. However, as noted earlier, there is considerable anecdotal evidence that a high rejection rate serves as a proxy of quality and may sometimes be adopted as desirable in itself.

We have also assumed away any hard constraints, such as page constraints, that some journals face. Many journals could not reject all submitted papers, even if they wanted to, because publishers expect them to publish a certain number of issues each year. Journals also face constraints regarding the maximum number of issues they can publish. We see no reason to believe that our results would be qualitatively different if we introduced such constraints, but this remains untested.

Regarding the process of peer review, our model treats this as a black-box process. We do not model the peer reviewers as influenced by the population of papers that are submitted. Were those agents more sophisticated, their judgments might be more informed, and the model might yield very different conclusions. We have made this assumption quite intentionally: we think the black-box model is a more accurate model of how reviewers typically work rather than modeling them as ideal Bayesian agents. Reviewers, of course, are aware of the general quality of the journal, and they are attempting to make a judgment as to whether a particular paper is good enough for that particular journal—our model is consistent with this. But reviewers generally know very little about the overall population of papers submitted to the journal and very little about other inputs to the editorial decision process.

Perhaps most critically, we assume that there is a single journal that exercises monopoly power. We think this model is appropriate for circumstances where there is a single “top” journal that is regarded as a critical journal for promotion and tenure. Our results might change in a setting where journals compete with one another in order to attract the best papers.¹⁷

5. Conclusion

Although limited by its assumptions, our model has identified important themes. First, and most worrying, is that journals incentivized by selectivity have strong incentives to maintain worse peer review than those incentivized by their quality. This occurs largely because journals incentivized by selectivity want to avoid self-selection. A paper that is never submitted cannot be rejected. As a result, we would anticipate that the quality of published papers will be lower in settings where journals and conferences advertise and are judged by their rejection rates. This has a quite clear implication: science functions better when journals are judged by the quality of the papers published therein as compared to a situation where journals are compared by their rejection rates.

Second, all journals are incentivized to create some appearance of high standards. That is, they’re incentivized to announce that their threshold is “we publish only the best papers.” But they’re also incentivized to imperfectly enforce their own standards. A journal with a more accurately enforced quality threshold—

¹⁷ Heintzelman and Nocetti (2009), Muller-Ippen (2022), and Oster (1980) consider the problem of multiple journals from the author’s point of view. They hold fixed the quality of the journal and ask what strategy an author should use in determining where to submit. Oosterhaven (2015) presents a model suggesting that there may be “too many” journals because in that model, every paper has a probability of acceptance approaching 1. But this result is driven by the idiosyncratic feature that all papers have the same probability of being accepted, regardless of quality.

even if the threshold is lower—might be a better-quality journal, and its quality would be more transparent to outsiders.¹⁸

In addition, should journals be able to affect the cost of submission, we might expect journals incentivized by quality to use that cost as a substitute for peer review. Increasing the cost of submission might, in some contexts, cause self-selection that acts as a substitute for improving the peer-review process. This is bad for the welfare of authors, and it is probably inefficient, given that it is an externality of the journal's submission policy.

As a result, we should not expect that incentivizing a journal by either its quality or its rejection rate will achieve high-quality peer review, nor should we expect it to maximize the efficiency of collective knowledge production. Identifying alternative incentive schemes or social organizations for journals should be an area of ongoing research.

Supplementary material. The supplementary material for this article can be found at <https://doi.org/10.1017/psa.2023.81>.

Acknowledgments. This research was funded by the Australian Research Council (Project DP190100041).

References

- Aarssen, Lonnie W., Tom Tregenza, Amber E. Budden, Christopher J. Lortie, Julia Koricheva, and Roosa Leimu. 2008. "Bang for Your Buck: Rejection Rates and Impact Factors in Ecological Journals." *Open Ecology Journal* 1 (1).
- Abrahams, S. G. 1977. "Framework for Estimating the Quality of Scientific Journals." *IEEE Transactions on Professional Communication* PC-20 (2):133–36.
- Azar, Ofer H. 2005. "The Review Process in Economics: Is It Too Fast?" *Southern Economic Journal* 72 (2):482–91.
- Azar, Ofer H. 2015. "A Model of the Academic Review Process with Informed Authors." *Berkeley Electronic Journal of Economic Analysis & Policy* 15 (2):865–89.
- Björk, Bo-Christer, and David Solomon. 2015. "Article Processing Charges in OA Journals: Relationship between Price and Quality." *Scientometrics* 103 (2):373–85.
- Bornmann, Lutz. 2011. "Scientific Peer Review." *Annual Review of Information Science and Technology* 45 (1):197–245.
- Brembs, Björn. 2018. "Prestigious Science Journals Struggle to Reach Even Average Reliability." *Frontiers in Human Neuroscience* 12 (February):1–7.
- Campanario, Juan Miguel. 1996. "Have Referees Rejected Some of the Most-Cited Articles of All Times?" *Journal of the American Society for Information Science* 47 (4):302–10.
- Campos-Arceiz, Ahimsa, Richard B. Primack, and Lian Pin Koh. 2015. "Reviewer Recommendations and Editors' Decisions for a Conservation Journal: Is It Just a Crapshoot? And Do Chinese Authors Get a Fair Shot?" *Biological Conservation* 186:22–27.
- Cicchetti, Domenic V. 1991. "The Reliability of Peer Review for Manuscript and Grant Submissions: A Cross-Disciplinary Investigation." *Behavioral and Brain Sciences* 14 (1):119–35.

¹⁸ Journals such as *PLoS ONE* and *Scientific Reports* eschew reference to importance or significance in determining whether a paper is worth publishing. They instead adopt purely methodological standards relating to scientific validity and rigor, with the idea being that any research, provided it is reliably conducted, should be made available to the public. This approach is made feasible by using an open-access publishing model, whereby authors pay for the cost of publishing the research. Journal editorial policies like these suggest a relatively transparent quality threshold for publication, combined with a stark refusal to be motivated by concerns of selectivity. It is relatively early in the history of these journals, but we expect that they will constitute useful cases in the future for measuring the impact of alternative approaches to academic publishing.

- Cullen, David J., and Anne Macaulay. 1992. "Consistency between Peer Reviewers for a Clinical Speciality Journal." *Academic Medicine: Journal of the Association of American Medical Colleges* 67 (12):856–59.
- Deveugele, Myriam, and Jonathan Silverman. 2017. "Peer-Review for Selection of Oral Presentations for Conferences: Are We Reliable?" *Patient Education and Counseling* 100 (11):2147–50.
- Egbert, Joy. 2007. "Quality Analysis of Journals in TESOL and Applied Linguistics." *TESOL Quarterly* 41 (1):157–71.
- Ellison, Glenn. 2002. "Evolving Standards for Academic Publishing: A $q-r$ Theory." *Journal of Political Economy* 110 (5):994–1034.
- Engers, Maxim, and Joshua S. Gans. 1998. "Why Referees Are Not Paid (Enough)." *American Economic Review* 88 (5):1341–49.
- Ernst, Edzard, Tobias Saradeth, and Karl Ludwig Resch. 1993. "Drawbacks of Peer Review." *Nature* 363 (6427):296.
- Geddes, Chris D. 2010. "JoF Rejection Rate Exceeds 55%." *Journal of Fluorescence* 20 (1).
- Heesen, Remco. 2018. "Why the Reward Structure of Science Makes Reproducibility Problems Inevitable." *Journal of Philosophy* 115 (12):661–74.
- Heesen, Remco, and Liam Kofi Bright. 2021. "Is Peer Review a Good Idea?" *British Journal for the Philosophy of Science* 72 (3):635–63.
- Heintzelman, Martin, and Diego Nocetti. 2009. "Where Should We Submit Our Manuscript? An Analysis of Journal Submission Strategies." *Berkeley Electronic Journal of Economic Analysis and Policy* 9 (1).
- Hörner, Johannes. 2019. "Report of the Editor: American Economic Journal: Microeconomics." *AEA Papers and Proceedings* 109:660–65.
- Howard, Louise, and Greg Wilkinson. 1998. "Peer Review and Editorial Decision-Making." *British Journal of Psychiatry* 173 (2):110–15.
- Jackson, Jeffrey L., Malathi Srinivasan, Joanna Rea, Kathlyn E. Fletcher, and Richard L. Kravitz. 2011. "The Validity of Peer Review in a General Medicine Journal." *PLoS ONE* 6 (7):1–8.
- Kravitz, Richard L., Peter Franks, Mitchell D. Feldman, Martha Gerrity, Cindy Byrne, and William M. Tierney. 2010. "Editorial Peer Reviewers' Recommendations at a General Medical Journal: Are They Reliable and Do Editors Care?" *PLoS ONE* 5 (4):2–6.
- Lee, Kirby P., Marieka Schotland, Peter Bacchetti, and Lisa A. Bero. 2002. "Association of Journal Quality Indicators with Methodological Quality of Clinical Research Articles." *JAMA* 287 (21):2805–8.
- Loison, Anne, Jon E. Swenson, and Henrik Andrén. 2006. "Editorial." *Wildlife Biology* 12 (1):1.
- Lowe, Alan, and Joanne Locke. 2005. "Perceptions of Journal Quality and Research Paradigm: Results of a Web-Based Survey of British Accounting Academics." *Accounting, Organizations and Society* 30 (1):81–98.
- Mahoney, Michael J. 1977. "Publication Prejudices: An Experimental Study of Confirmatory Bias in the Peer Review System." *Cognitive Therapy and Research* 1 (2):161–75.
- Muller-Itten, Michele. 2022. "Gatekeeper Competition." SSRN Scholarly Paper. <https://papers.ssrn.com/abstract=4059806>.
- Oosterhaven, Jan. 2015. "Too Many Journals? Towards a Theory of Repeated Rejections and Ultimate Acceptance." *Scientometrics* 103 (1):261–65.
- Oster, Sharon. 1980. "The Optimal Order for Submitting Manuscripts." *American Economic Review* 70 (3):444–48.
- Peters, Douglas P., and Stephen J. Ceci. 1982. "Peer-Review Practices of Psychological Journals: The Fate of Published Articles, Submitted Again." *Behavioral and Brain Sciences* 5 (2):187–95.
- Regazzi, John J., and Selenay Aytac. 2008. "Author Perceptions of Journal Quality." *Learned Publishing* 21 (3):225–35.
- Reinhart, Martin. 2009. "Peer Review of Grant Applications in Biology and Medicine. Reliability, Fairness, and Validity." *Scientometrics* 81 (3):789–809.
- Rothwell, Peter M., and Christopher N. Martyn. 2000. "Reproducibility of Peer Review in Clinical Neuroscience: Is Agreement between Reviewers Any Greater Than Would Be Expected by Chance Alone?" *Brain* 123 (9):1964–69.
- Rubin, Haya R., Donald A. Redelmeier, Albert W. Wu, and Earl P. Steinberg. 1993. "How Reliable Is Peer Review of Scientific Abstracts?" *Journal of General Internal Medicine* 8 (5):255–58.
- Saha, Somnath, Sanjay Saint, and Dimitri A. Christakis. 2003. "Impact Factor: A Valid Measure of Journal Quality?" *Journal of the Medical Library Association* 91 (1):42–46.

- Schultz, David M. 2010. "Rejection Rates for Journals Publishing in the Atmospheric Sciences." *Bulletin of the American Meteorological Society* 91 (2):231–44.
- Shideler, Geoffrey S., and Rafael J. Araújo. 2016. "Measures of Scholarly Journal Quality Are Not Universally Applicable to Determining Value of Advertised Annual Subscription Price." *Scientometrics* 107 (3):963–73.
- Siegelman, Stanley S. 1991. "Assassins and Zealots: Variations in Peer Review: Special Report." *Radiology* 178 (3):637–42.
- Siler, Kyle, Kirby Lee, and Lisa Bero. 2015. "Measuring the Effectiveness of Scientific Gatekeeping." *Proceedings of the National Academy of Sciences of the United States of America* 112 (2):360–65.
- Snodgrass, Richard. 2006. "Single- versus Double-Blind Reviewing: An Analysis of the Literature." *SIGMOD Record* 35 (3):8–21.
- Strevens, Michael. 2012. "Economic Approaches to Understanding Scientific Norms." *Episteme* 8 (2):184–200.
- Tiokhin, Leonid, Karthik Panchanathan, Daniel Lakens, Simine Vazire, Thomas Morgan, and Kevin Zollman. 2021. "Honest Signaling in Academic Publishing." *PLoS ONE* 16 (2):e0246675.
- Zollman, Kevin J. S. 2018. "The Credit Economy and the Economic Rationality of Science." *Journal of Philosophy* 115 (1):5–33.

Cite this article: Zollman, Kevin J. S., Julian García, and Toby Handfield. 2024. "Academic Journals, Incentives, and the Quality of Peer Review: A Model." *Philosophy of Science* 91 (1):186–203. <https://doi.org/10.1017/psa.2023.81>