

# Why I'm going to self-publish my research (and it's not because I think journals are bad)

JEFF HOULAHAN

Look, let's face it. The main reason I'm going to stop publishing in journals and just post my work on my own website is because I can. I am nearing the end of my salaried and funded research career, but I'm hoping I'm mid-late career as a scientist. My pension means that I won't have to worry about being employed to pay my rent and buy groceries. The kind of research I do doesn't require funding for graduate students or fieldwork or equipment. So, I don't need to get 'credit' for my research. That is, I don't need to maintain a cv that tells my bosses I'm earning my salary and funding agencies where I stand among other scientists requesting their money. I no longer need to 'keep score'. But is that the only reason I published in journals? Are there good reasons to publish in journals other than as a way of 'keeping score'? Are journals simply a way of providing accreditation for the research you've done in a way that's analogous to universities providing accreditation for the things you know and the skills you have?

There seem to be three important claims that journals make for their value – first, peer review, second, improved quality of presentation and third, distribution. Let's work through these in reverse order. There seems little doubt to me that if you publish in a high-profile journal, more people will hear about your work, read your work and ultimately cite your work. That is undoubtedly true, if you publish in Science or Nature. Maybe PNAS, as well. As I move into my discipline-specific journals – ecology – I'm less certain. It's not clear to me what the

'eyes-on' effect of publishing in Ecology Letters versus Am Nat versus Oikos versus Oecologia is. I suspect there is a 'fine-grain' distribution effect that depends to some extent on where you publish. That is, some journals have a higher readership than others. And even more likely a 'coarse-grain' distribution effect of publishing in an established journal versus on my own website – even if it's just having your article in an the Abstract databases. But I'm not interested in marketing my research. I'm not interested in convincing people that my research is worth reading...because I have no idea if it is worth reading. The vast majority of us do research that will have at best, very small effects on our general state of understanding and knowledge. In a world crammed with claims to knowledge and understanding, my research is yet another claim. I trust potential readers to decide if it is worth their time...and make no claim that it is. And I have little doubt that self-publishing will mean fewer readers than if I published in an established journal. I just don't care...very much.

Second, improved quality of presentation. This is a claim that I wouldn't care to dispute – the difference in quality between the Word documents my collaborators and I submit to journals and the product that appears on the page and/or online is generally substantial. Most (all?) reputable journals have a technical staff whose job is to make scientific manuscripts look good. So, research I self-publish will not be as aesthetically pleasing, may have more typos and may be less readable. This reduced quality

may prevent readers from reading my work. I just don't care...see above.

Third, peer review. This is the substantive claim that we make for publishing in journals and the claim that journals make for their existence – that published work has been vetted by experts in the field. Peer review provides some assurance that the research is...what? Right? True? Reliable? Well, we know it can't provide any assurances about the truth content of research because that's not how science works. The world is complex and data can mislead so, the truth content of published research is generally uncertain. Most (all?) phenomena we are trying to understand in ecology have multiple, interacting and potentially nonlinear causal variables and are thus going to be difficult to understand even if there are general principles that transfer across space, time and taxonomy. Thus, some and perhaps lots of ecological research is going to be 'good faith' attempts but wrong. Reliability seems like a more realistic target, but reliability can only be based on the process for reaching the conclusion rather than the conclusion itself. So, peer review – if the authors have been transparent and truthful about their methodology – can assess the process by which the conclusions were reached. And I suspect, it does this. However, two or three potentially conflicted reviewers providing reviews of widely varying quality to a busy editor who may be doing little or no evaluation of either the paper or the reviews themselves is, at best, a porous membrane for filtering out 'unreliable' research. Luckily, the large majority of published research has no impact on a discipline and so, published unreliable research is of little consequence. And the published research that has a large effect on a discipline generally goes through lots of post-publication peer review as the research is debated and revisited over time. So, I believe that peer review does improve the reliability of published research relative to what would be

published without peer review, but my intuition is that the improvement is small.

My experience with peer review over 30 years of having my work reviewed is that, occasionally, a reviewer has spotted something that I and my co-authors had missed, and the paper was significantly improved by the reviewer's contribution, though as I try and think of a specific example, I can't find one. I suspect this has more to do with my memory than the lack of real examples, but I'm not sure. The way reviewers helped most often was,

1. Making an irrelevant or misguided criticism that identified where we hadn't been as clear in our writing as we should have been.
2. Asking us to be clearer about something they didn't understand.
3. Cautioning us that we were making inferences that were more strongly stated than were warranted by the results. I generally disagreed with the reviewers, but moderated the inferences because it would be easier to get accepted if I went along with the reviewers. And I tend to be less cautious about inferences than some, so they were probably right.
4. Pointing out a relevant reference that we had missed.

None of these do anything to increase the reliability of the research, though they may make it easier to assess the reliability. As an aside, I have had a very small number (2? 3? 4?) of papers accepted by the journal they were first submitted to. And very few papers that I've submitted weren't ultimately accepted elsewhere. And when those papers were accepted elsewhere it was NEVER after a major revision that led to fundamental changes in methodology or conclusions. Papers generally got rejected because they weren't interesting enough, they were better suited for another

journal or the particular reviewers had major concerns about something that different reviewers (and we) didn't feel were a major concern – to-may-to, to-mah-to.

Post-publication review never, in my opinion, identified serious flaws in the research. So, while research I've worked on has often been deemed unfit for publication in a particular journal, the research has almost always been published later with minor or no changes. I suspect this is true for most ecologists and perhaps most researchers. My conclusion is that there are minor benefits to peer review, but they have little to do with assuring that 'bad' research doesn't get published. At what cost?

For me the costs have been – (1) publishing less of the research I've done, (2) taking too long to publish research, and (3) following a written format that emphasizes the wrong aspects of scientific research.

I've had a handful of papers that I have given up on. That is, papers that were rejected and I have ultimately decided not to submit elsewhere. I've never given up because peer review revealed, to my satisfaction, a fatal flaw – or even a significant flaw that would involve a large investment of time to fix. It was always that I just lost interest, got distracted by newer and more interesting ideas and wasn't willing to commit the time and effort to pushing a manuscript over the finish line. In fact, I rarely made the explicit decision not to resubmit – they just slid off the side of my desk. Further, there are another 5 – 10 manuscripts that I never submitted at all, simply because the tedium of getting over the 'acceptance' bar prevented me from submitting them even once. What is the nature of that tedium? Well, it is the reading and writing necessary to prepare an acceptable intro, the reading and writing necessary to show an acceptable familiarity with the published literature in the discussion, and the finicky formatting of text and figures to

meet the demands of a particular journal. But it is also knowing that unfinished thoughts are hard to get published. So, interesting but incomplete results that would require a great deal of additional analyses to close off critical gaps go unpublished. Because I know they will be met with resistance by reviewers and editors buried beneath a pile of submitted manuscripts that may be more complete but also more pedestrian. I have a significant amount of unpublished research that is unpublished because it's only partially baked. But to continue the recipe, ingredients, meal analogy – I buy lots of unbaked and half-baked meals and finish them off on my own time. All scientific disciplines are work in progress - to some extent half-baked. Even stories that appear complete, never are. I think the deluge of submitted research has set a standard for coherence that I'm not sure is good for scientific progress and that I'm no longer interested in meeting. I want to be able to publish work that has hanging threads, missing buttons, frayed cuffs and collars, that's loose in the waist and tight across the shoulders, that has one sleeve shorter than the other and a missing pant leg – I know, I've shifted from baker to tailor – with an invitation for others to make the alterations...or not. If there is no interest.

The time to publication problem is tightly linked to the 'less research' problem because as time passes between completing one stage of a piece of research and acceptance my excitement and motivation diminishes, and the research slides closer and closer to the edge of my desk.

The two essential pieces of any scientific research are the methods and the results – and published manuscripts place a lot of emphasis on the introduction and the discussion. The methods are the recipe and the ingredients, and the results are the finished product – the meal. Everything else is garnish and presentation. The introduction, in particular, has little value. The

introduction is where we provide the rationale for a piece of research. Why have we done this work? The rationale usually takes one of two forms – some practical problem that our research will have putative relevance to, or some previous research left a gap to be filled. But who cares? Who cares why we did it? If I (or you) find a result interesting, why would you care why the researcher did the work? Most introductions are really just a way for authors to show they've done the upfront work. And doing the upfront work may prevent you from repeating something that has already been done – although why we don't want people to repeat research that has already been done escapes me. And would it be preferable that the author realizes they are repeating research? Sure. But is it really important? I can't see why. And doing the upfront work may prevent you from repeating previous mistakes. That would be valuable. But I'm not sure why I, as the reader, need to hear about it. Most of us decide to read a paper because the title captures our attention, and a quick skim of the abstract suggests results we would be interested in. I have never used introductions to decide if I will read the rest of a paper and, in fact, I rarely read them. Except for the last paragraph where the authors usually state their objective(s). Any more than a paragraph for the introduction usually exceeds my interest (unless, of course, they cite me). The discussion has more value than the introduction, but not a lot more value. It is where we use the results to make the conclusions. Now, often the conclusions that arise from our results are so obvious as to be trivial – but not always. There are times when the results imply conclusions that are not obvious and may, in fact, be non-intuitive. So, a paragraph or two identifying the inferences and conclusions warranted by the results is an important piece. But the key role of the discussion is “to place your research in the

context of other published research”. What previous research is consistent with your findings? What research is inconsistent? Why might previous research be consistent or inconsistent? Who cares? Why is it relevant? What I really want to know is what new understanding or knowledge your research has provided and that will be measured by improved predictive ability. That does mean I should know what the current ‘best’ model is and how the predictive ability of my modified model compares to the previous gold standard. But discussions rarely, if ever, include those pieces. Sure, identifying weaknesses of your paper and logical next steps are useful, but not essential. Two or three or maybe four paragraphs seem like enough for most discussions.

And each of those two sections, the introduction and discussion, will be packed with citations – references to previous work that supports assertions of fact you're making. However, most of us write the ‘fact’ and then go searching for the supporting citation. If we constrained ourselves to only referencing papers that actually motivated the sentence we wrote, we would dramatically decrease the number of references. For all other statements of ‘fact’, leave them unreferenced...or unwritten. If people want to dispute the unreferenced fact, let them. These statements of ‘fact’ are rarely very important or very relevant to the important findings of a paper. So, who cares? More than 6 – 10 citations is probably overkill.

So, I will no longer be submitting articles to journals (except for co-authoring student papers...and that is almost finished). And not in the name of some higher principle like open and equal access to the scientific literature – though I think that is an important principle. I am doing this because I want to – and I can.

