

Writing Philosophy for Publication

I was asked to write this for a 2021 seminar at the ANU run by Al Hájek and Brian Hedden. It's just one guy's opinion—or rather, a bunch of opinions collected by one guy—but hey, maybe you'll find it useful.

So you want to publish some philosophy—preferably, good philosophy in a nice journal. How do you do it?

This document has some tips for each stage of the process, from choosing ideas to appeasing Referee #2. But it's not a full guide. I am not telling you how to do philosophy, or how much you should publish. I'm just offering advice on how to pick and polish ideas so that, *if you want publications*, you can maximize your odds while minimizing stress.

Part 1: Publishability

Before we get to the process, let's talk about the telos. Some papers have “PUBLISH ME!” written all over them. Others, however brilliant, do not. What makes the difference? If you want to publish, what should you be aiming for, and what should you avoid?

You should aim for:

(1) A paper whose contribution is immediately obvious.

I don't just mean you immediately *tell* the reader how important your thesis is. (“In section 2, I prove that everyone else is dumb and wrong.”) Nor do I mean *showing* the reader all of your cards in the first section. That's what the rest of the paper is for! A good intro is a mix. You want to tell the reader what you're up to in a way that gives them a tantalizing glimpse of what's to come—like a movie trailer, but with spoilers.

A snappy descriptive title helps. But to my mind, the really crucial thing is to start your introduction by **framing the question**. Don't just launch into your thesis, bury us in citations, or bust out your “road map.” First lay down the basic concepts that your reader needs to understand the question at issue, then give them a sense of why it matters, and *only then* give a peek at your answer, making sure that your point ends up looking new and interesting.

An easy way to do this is by explaining the main issue in a way that creates tension, then telling the reader (not laboriously showing!) how you're going to resolve that tension in a novel way. I think I do a decent job of this in [my paper with Jack Spencer](#). Our question is: are objective ‘oughts’ authoritative? Our intro explains objectivity and authority, and we state our answer (no!), but we don't nutshell the argument. We just give a glimpse: we'll be using a case (the miners puzzle) and a certain principle (called “monotonicity”). That's enough to build tension. Then we cite around to assure the reader no one's said this before.

Here are some A+ question framers: Judy Thomson, David Lewis, Brian Hedden, [Jake Nebel](#), [Tamar Schapiro](#), Richard Chappell, Al Hájek, [David Builes](#), [Zach Barnett](#).

Here are two classics with less publishable framings: [“Singling out Properties”](#) (one of my all-time favorites, but the style is risky to emulate), [“Reference and Definite Descriptions”](#) (no early glimpse of the main idea—and for no good reason).

(2) A goldilocks thesis.

By which I mean: not too epic, not too tiny—just right. You want a claim that’s grand enough to be interesting but not so grand that you lack room for a proper defense.

If you’re struggling to publish an ambitious project, the problem might not be quality but quantity: your thesis may be overflowing your paper. (You can’t frame the question right if the question doesn’t fit the frame!) If so, consider some ways of cutting down on grandiosity. Can you split the paper in two? Can you weaken the thesis?

One nice way to weaken theses is by changing an argument for a specific view into a paradox for everyone. That way, you aren’t on the hook for giving a comprehensive defense of your view. You can just give a partial defense, be honest about the view’s shortcomings—and that’s enough. The full defense can wait until later papers.

Example. I had a paper that I thought was good—it had a Revise & Resubmit at *Ethics*, and my advisors and friends liked it. But the R&R fell through, and the paper kept getting rejected—even desk rejected!—by journal after journal. What was wrong? Eventually I realized the conclusion was probably too grand. I had been arguing that my own wacky form of deontology was true; I then had to assure the reader that, even though the view invites tons of objections, I had tons of answers, which I’d tell them about in later papers. Cold comfort!

In the end, I changed the conclusion into a dilemma. I said that this wacky deontology is the only kind that’s unified, so deontologists have a hard choice: be disunified, or be wacky. [The paper](#) was then accepted at the next journal without revisions.

The flipside of this problem: if you are struggling to publish a paper that is very long, it might be that you are overestimating how much you have to say. A medium-sized paper that makes a medium-sized point is vastly more publishable than a long paper with a medium-sized point. Example: I have a paper in *Philosophical Studies* that is less important than the journal’s average, but it got in because it was accordingly short (6,500 words). If I’d left it at 8,000 words, it would’ve been rejected. Same if I’d sent it to fancier journals like *Mind*. This brings me to:

(3) The right venue.

Even if your paper is exactly what it's supposed to be, that doesn't mean that it will get in wherever. Ask your advisors and savvy friends what journals you should be sending which papers to. Journals differ in their preferred (a) topics, (b) word counts, and (c) level of grandeur. (*Inquiry* will publish singles; *AJP* wants doubles; *Phil Review* is after grand slams.) There isn't much else to say about this, since it's going to depend on what kind of papers you're trying to publish, and I haven't read your papers. Though I suppose you could always email me.

Preparing the manuscript

These might go without saying, but...

- Pick a nice font, decent margin sizes, and good spacing. For example: 12pt Times New Roman double-spaced with one-inch margins. (And check what the journal likes here; *Phil Review* wants double-spaced footnotes, for example.)
- Anonymize your paper carefully. You don't want to waste your time, as I have before, following up with managing editors about self-citations and what not.
- Cite carefully. Good citations can (a) signal that you know your stuff, (b) help your reader see what makes your paper new, and (c) help the editors find a referee—who will be pleased to see that you've acknowledged their work.

So much for what to aim for. Now, **things to avoid:**

(1) Clutter.

I don't just mean pruning a few words. In the final stages of writing, you should at least consider scrapping whole paragraphs or even pages. If your paper is stodgy, messy, or busy, a helpful exercise is to delete a section and just see what happens. If that ruins the paper—great! Hit undo and send it out. If the paper suddenly flows—even better! Just edit for consistency, and you have a sleeker paper, which means it's easier to read, which means your peer reviewers will be happy.

Of course, it hurts to delete something you've poured your heart into. But that's publishing: one heartbreak after another. You just can't be sentimental about these things. (If it helps, you can keep your deleted paragraphs in a graveyard document like I do to reassure yourself that nothing was really "deleted." Whatever works!)

(2) Dismissiveness.

The best advice I ever got about publishing philosophy was from Agustín Rayo. He told me that, as a young philosopher, it's natural to see the people you disagree with as the enemy, and to see your papers as an attack on them for the benefit of your audience (namely, all other philosophers, who are of course eagerly awaiting your every word). But this is totally backwards. The people you're "attacking" typically *are* your audience—the ones who actually care enough to read your work, think about it deeply, write replies—and they are your colleagues for life.

They are also, for what it's worth, your likely referees. They are the peers who will be doing the reviewing.

So why annoy them? It's in your interest, and intrinsically good, to write with grace. Emphasize agreements. Acknowledge others' contributions. Concede weaknesses in your position. Don't be dismissive, and definitely don't nitpick—especially not over mere words. It's good to clarify that you use terms one way whereas they use them in another; maybe your way is better. But be gracious. Remember who your readers are.

I have never once regretted being generous in a paper, and I have never once been glad I took a cheap shot.

(3) Side Quests.

In a role-playing game, there is a main quest (slay the dragon!), and there are side-quests (collect the relics!). The point of side-quests is to add depth to the game for people who want to get their money's worth exploring the world, having some fun, taking their time—in other words, the exact *opposite* mindset of a busy peer-reviewer, who just wants to finish your paper and learn what you have to say.

Rule of thumb: **no side-quests**. Referees and editors want a straight shot to the dragon. Most of your paper should be preparing for, traveling to, and carrying out the main quest.

This rule can also help you choose which ideas to focus on. I once had a paper that ended with "boring" side-quests: a series of replies to objections. Several people said I was taking too much time away from developing my own views. Eventually I took the hint. For about a year, I kept trimming the section, and I decided to focus on the objections that were most essential to clarifying the paper's positive view, using them to further the main quest. When [the paper](#) was finally accepted, the referee seemed to think the replies were the best part—what a turnaround!

Part 2: The Process

The path to getting your paper accepted starts with brainstorming and ends, if you're lucky, with desperate pleas to an anonymous referee. Let me just give one tip per stage.

Step 1: Choosing Ideas

The first step may be the most underrated. You can pour your heart and soul into a paper—not to mention months of your life—and still struggle to publish it, if you simply choose the wrong idea. I don't mean choosing a *bad* idea. I mean choosing one that isn't publishable.

Partly, this is a matter of whether the thesis is the right size. But like I said, if it's too big, there are ways to fix this: split the paper, reframe the conclusion. If it's too small, you could try combining it with ideas you have left over from other papers, or discussing objections.

Here's my main tip. Don't worry too much about choosing the perfect-sized idea from the get-go. It's important to choose a *topic* that's publishable (best evidence: people you admire are publishing on it). But for the first draft, don't get hung up on framing. Just frame the paper in a way that makes it *easy to write*. If you find it easy to write the giant manifesto first, go ahead—you can chop it up later. I usually prefer to start small, but that's just me.

If you are really struggling to get the paper down (or up) to size, that may be a sign that it's not publishable—at least, not for you right now. Don't beat yourself up. It's part of the process. You can't know which ideas are going to be publishable a priori. Thankfully, your guesses will get better over time.

Step 2: Writing the Damn Paper

As I said, you want your paper's contribution to be immediately obvious, even to referees who disagree with your views and are skimming the paper. In light of that, let me give one tip, which I haven't seen anyone give yet. Your paper should be **focused as much as possible on vivid, specific ideas**, like good journalism. Don't heap on hyper-abstract language.

Witness the awful abstractness of Grice's *Studies in the Ways of Words*, page 1:

Prolegomena

There is a familiar and, to many, very natural maneuver which is of frequent occurrence in conceptual inquiries, whether of a philosophical or of a nonphilosophical character. It proceeds as follows: one begins with the observation that a certain range of expressions E , in each of which is embedded a subordinate expression α —let us call this range $E(\alpha)$ —is such that its members would not be used in application to certain specimen situations, that their use would be odd or inappropriate or even would make no sense; one then suggests that the relevant feature of such situations is that they fail to satisfy some condition C (which may be negative in character); and one concludes that it is a characteristic of the concept expressed by α , a feature of the meaning or use of α , that $E(\alpha)$ is applicable only if C is satisfied. Such a conclusion may be associated with one or more of the following more specific claims: that the schema $E(\alpha)$ logically entails C , that it implies or presupposes C , or that C is an applicability/appropriateness-condition (in a specially explained sense) for α and that α is misused unless C obtains.

Before mentioning suspect examples of this type of maneuver, I would like to make two general remarks. First, if it is any part of one's philosophical concern, as it is of mine, to give an accurate general account of the actual meaning of this or that expression in nontechnical discourse, then one simply cannot afford to abandon this kind of maneuver altogether. So there is an obvious need for a method (which may not, of course, be such as to constitute a clear-cut decision procedure) for distinguishing its legitimate from its illegitimate applications. Second, various persons, including myself, have pointed

What is he *talking* about? (I laugh every time at that indispensable qualification: "(which may be negative in character)." Ah, *now* I get it!)

Now read the first paragraphs of Zach Barnett's "Why You Should Vote to Change the Outcome."

When voting comes with a cost, why pay it? Sometimes, there is a simple answer. We pay the cost to *make our preferred outcomes likelier*. When I cast my vote for Class President or Team Captain, there's a certain result I want, and I'm trying to bring it about.

One might worry, though, that this simple rationale is inapplicable to very large elections. After all, when there are millions and millions of voters, the chance that my individual vote will make a difference to the final outcome is utterly miniscule. How could it be rational for me to do something that's virtually certain to have zero impact?

In response to this challenge, one might encourage me not to overlook the magnitude of the *stakes*. If there are millions of voters, there are, presumably, millions of people who will be affected by the result. Yes, the chance that my vote makes the difference is very tiny, but the difference my vote could make is very great. Arguably, the magnitude of the stakes can, at least sometimes, offset the tininess of the chance of affecting the outcome—making it rational to vote solely in virtue of the expected consequences of doing so. This is the *consequentialist defense of voting*.

And off he goes. He hasn't gotten to his contribution yet, but already, we know *exactly* what the question is—why vote?—and we are getting a vivid look at the specific idea that Barnett will argue for in greater depth: the consequentialist defense.

Now, I don't mean to imply that abstraction is always bad. It's good when you really need to be precise about the scope of what you're saying. But publishable papers are more like Barnett's than Grice's, especially in their opening pages. Papers in high places tend to lead not with gratuitous formalism but with intuitive examples. Barnett gets to his examples in sentence four. In Grice's fourth sentence, we're told that if we want examples we'll just have to wait!

Step 3: Rejection

Expect it. You can write a great paper, send it to a plausible venue, and end up with what feels like nothing to show for it—a mean set of comments or an aloof desk rejection. Again: you can't get sentimental about this stuff. It gets easier. I promise.

One common piece of advice is that you should ignore rejections. The journals are a "crap shoot." If you get snakes eyes, just reroll, and eventually, you'll win.

I'm not sure this is good advice. In my experience, **journals often agree**, and the system is less random than it looks. A single rejection can be a fluke. But if your paper gets rejected multiple times, then unless you're sure that your referees were anomalies, **edit the paper**. Especially if you got multiple complaints about the same thing. That's always a blessing: you know exactly what to work on.

Think of yourself as a movie producer doing test-screenings. You'd rather not have to do any big edits—and would hate to resort to reshoots—but you want everyone to like the final cut. Of course, you can't make *everyone* like it, and you can't just mangle the film. The hope is to find simple fixes that will win back big chunks of the audience.

That's why I say it's a blessing when two referees agree. It's like having multiple test audiences say they hate the love triangle. A good producer lives for that kind of feedback!

If I may give my own advice: when you're in the process of (re)submitting your paper, **identify the parts most likely to get it rejected**—the scenes when most people leave the theatre. Your paper won't please everyone. But you want to know which things might confuse the reader, annoy your opponents, and risk further rejection.

Step 3*: Resubmission

Congratulations: you've gotten a Revise and Resubmit...now what?

Obviously, you have to resubmit a revised paper. Less obviously, you should add in a short cover letter, in which you go through the referees' points and explain the changes you did or didn't make. This report is almost as important as the revisions themselves.

How do you write it? Standard advice: **be succinct and professional**. You can start with a quick "thank you," and by the end you want to have covered the referees' major points. No tangents, no flippant slang. (You'd be amazed at the stuff you see as a referee.) I find it helpful to start the letter before doing the actual revisions, but you could go the other way around. Just double-check that the letter is accurate before sending it off.

Now the substance. How do you decide when to listen to the referees?

It depends, of course, on whether you agree with their advice. If you do, life is simple: just make the paper as good as you can in the ways that they want. If you don't agree, life gets complicated. You can grit your teeth and do what they tell you, or you can plead your case in the cover letter. Neither option is ideal. Generally, I grit my teeth on the little things and stand my ground on the big stuff.

But there's a helpful rule of thumb. **Since life is easy when the refs are right, always see their suggestions in the best possible light.** For example, if they say the introduction needs cuts, don't get defensive, and don't think "That sounds like work." Sit down and look for good ways to make cuts. If you *really* can't find any, that's ok; just say so in the cover letter. But look carefully, because if you do find cuts, it's a double win. You've got a better paper *and* a better cover letter.

Repeat?

I've also got a tip about quantity. If you want to publish more papers at the highest quality you manage, **don't just put out singles: put out albums.**

It takes time, effort, and even cash to read up in a new area. If you can use what you learn to write multiple papers, why not at least try? If the band is warmed up, why leave after the first song?

Of course, there are also reasons to branch out into new topics—inspiration, for example. Being narrow can be a false economy. But I’m not saying “be narrow.” I’m just saying: you won’t publish lots of great work if each project starts from scratch. If you’ve never written two papers on the same topic, try it. Or try splitting your biggest papers in two. Every paper needs something new—but not *everything* new!

Conclusion

That’s it from me. Feel free to email me suggestions for this document, since I plan to keep it updated. Most of the advice in here is the sort of thing one has to repeat to oneself. (“You just can’t be sentimental” about journal rejections—if only it were that easy!) The advice is also supposed to be universalizable, not just a way to get ahead of others. I’d rather focus on things that make the process more pleasant for everyone.

On that note, one last thought.

I know I said publishing is suffering, and that you should expect rejection. It is, and you should. But that doesn’t mean that you have to *dread* publishing, or that you can’t have any fun. If you are struggling to motivate yourself to publish, I recommend trying to find things about the process that are intrinsically interesting or rewarding to you. Try on new framings. Instead of trying to please the refs, you’re aiming to make your paper’s ideas accessible to everyone. Instead of racking up publication points, think of yourself as getting to the bottom of a fascinating puzzle.

Because that *is* what you’re doing. Publishing is an imperfect part of a beautiful cooperative system by which we produce and circulate ideas. To succeed in the system, sometimes you have to think strategically (and read dreary guides like this one). But you don’t want to overdo the cynical strategizing. You have to keep in touch with the reasons why you *like* your papers, and the reasons why it can be such a joy to share your ideas with the philosophical world.

Daniel Muñoz
munozdanielb@gmail.com

Acknowledgments

Thanks to everyone who’s given me writing and publishing tips over the years, including:

Al Hájek, Brian Hedden (“You want your paper to be skimmable”), Zach Barnett, Lindsay Brainard, David Builes, Steve Yablo (“That’s what the rest of your paper is for!”), Theron Pummer (speaking of [publishable](#)...), Kerah Gordon-Solmon, Caspar Hare, Kieran Setiya, Brendan de Kenessey, Brad Skow (“Don’t rush your intro”), Tamar Schapiro (“You start out with your ideas on top of you; the goal is to get on top of them!”).

Mea Culpa

Sorry to the people I forgot to thank, and especially to any friends left off my (obviously incomplete) list of good writers. I may make a bigger list later—then I'll be *really* sorry if I leave you off!