

A guide to professional philosophy

Neil Mehta, Professor of Philosophy, National University of Singapore

October 2025

Table of Contents

1. Reading: process.....	2
2. Reading: content.....	3
3. Writing: process.....	4
4. Writing: content and form.....	5
5. Writing: heuristics.....	9
6. Writing: publishing strategies.....	12
7. Cultivating yourself as a philosopher.....	14

This is a guide on how to read and write like a professional philosopher.¹ It is intended primarily for *graduate students* working in *contemporary philosophy*. Its value will drop off for philosophy professors, who will already have their own methods, and historians of philosophy, who will operate under somewhat different conventions. I have written it because little information on these topics is in print, because even excellent philosophy graduate programs rarely address the topic systematically, and because excellence in *general* reading and writing is not nearly enough for excellence in *philosophical* reading and writing.²

It is often thought that philosophical excellence is the sole province of the genius. I by contrast believe that it can be achieved through countless small techniques that can be taught and learned; in any case, if there is some special talent here, then I do not have it. When, near the end of my graduate career, I began submitting my work to journals, I received *18 consecutive rejections* distributed across *5* different manuscripts, with not even an invitation to revise-and-resubmit to cushion the blows. At last, on the 19th try, I secured my first publication in *Philosophical Studies*, and even that was merely a reply piece. I then undertook a disciplined program to improve my writing, and over the next several years my work climbed steadily in quality. I have now published a book with MIT Press and 17 articles in venues including *Noûs*, *Journal of Philosophy*, and *Philosophy and Phenomenological Research*. That does not make me a great philosopher: plenty of my colleagues have better records than this. It does, however, make me someone who has improved enormously, and here I will share precisely how that happened.

My advice will be contentious, so questions, suggestions, and criticisms are welcome. (Praise would be not just welcome but surprising.) You can contact me at neil.jag.mehta@gmail.com. I also offer detailed explanations of foundational philosophical skills for students of all levels: see <http://www.profneilmehta.com/resources.html>.

¹ It is a much-revised version of my 2016 piece, “A writing guide for professional philosophers.”

² There are many other resources for philosophy graduate students. Mercedes Corredor has an excellent compilation of them here:

https://docs.google.com/document/d/1qIcU5T_zbEAOAGAn7svux7etj3ELtOZZiTTwuINsubis/edit?tab=t.0.

1. Reading: process

In philosophy, good reading is active reading: it is less like scooping up water and more like wrestling a bear. I therefore encourage you to use these principles as you take notes.

Use premise-conclusion form. It is hard to answer the question, “What is LXXXII + XVIII?” It is easy to answer the question, “What is $82 + 18$?” If you want to do good arithmetic, use Arabic numerals.

It is also hard to evaluate arguments expressed in long, winding stretches of prose, and, if not *easy*, at least much *easier* to evaluate arguments expressed in premise-conclusion form. If you want to do good philosophy, regiment the key arguments – both the global argument of the entire text and the more local arguments of its sections. Take as your compass your knowledge of good argument forms, including forms that are deductively valid (*modus ponens*, *modus tollens*, argument by elimination) and forms that are not (inference to the best explanation, argument by induction, argument by analogy).

Identify evidence and theories. Typically, philosophy begins with (apparent) *evidence* – roughly, the things we know. Express, as cleanly as possible, whatever evidence a text introduces.

For instance, in the philosophy of perception, it has been observed that we can learn many facts about (say) the color red through methods like testimony. However, when we actually *see* something red, something more happens: we learn *what redness is, in and of itself*.³ You might cleanly express the contrast like this:

- Perceiving something enables us to learn what that thing is, in and of itself.
- Merely thinking about something does not enable this.

Then there are *theories* – attempts to explain or accommodate the evidence (and to meet other constraints). These, too, should be expressed as cleanly as possible. For instance, here is how we might express one *naïve realist* theory that can explain the evidence above:

- To perceive something is to stand in a primitive, non-representational relation of awareness – *acquaintance* – to it.
- Being acquainted with something reveals what that thing is, in and of itself.
- To think about something is to stand in a merely representational relation of awareness to it.
- Merely representing something does not reveal what that thing is, in and of itself.

Take notes on the virtues and vices of each theory: what (apparent) evidence it can explain and what it must explain away, how simple and fruitful it is, how independently plausible its posits are, etc.

Use concrete cases to understand abstract ideas. It is hard to grasp the idea that for virtually any empirical belief that apparently amounts to knowledge, there is some scenario *S* such that *S* is compatible with all of the subject’s evidence but incompatible with the truth of that belief. It is easy to grasp the idea that given all of our evidence, we might be the epistemic

³ Russell (1911); Campbell (2002, chs. 6-7); Johnston (2006, pp. 264-265).

playthings of a deceiving demon. To understand what abstract ideas mean, just keep applying them to simple cases.

Develop productive taxonomies. A productive taxonomy is one that divides possible views into categories that are mutually exclusive, jointly exhaustive, and separated by important fault-lines. Generate such taxonomies as you read about any topic. If you later encounter a view that rejects some foundational assumption of your taxonomy, broaden the taxonomy: include the question of whether or not the assumption holds.

Consider, for example, views about the relationship between facts about reasons and facts about rationality. We might initially sort such views into two categories: as holding either that the former always explain the latter or that the latter always explain the former. We might later realize that this presupposes that there is a uniform direction of explanation here, at which point we can add a third category: that there is no uniform direction of explanation.

Use fruitful categories. Identify the central *metaphysical*, *epistemic*, and *semantic* questions in an area, as well as the connections among possible answers to those questions. Also identify what a theory takes as *metaphysically* or *conceptually fundamental* and how it uses those things to explain what is *non-fundamental*.

For instance, in meta-ethics, the anti-reductivist holds that some normative entities – perhaps goodness, rightness, reasons, obligation, or virtue – are metaphysically fundamental. Reductivists hold in contrast that nothing normative is metaphysically fundamental; they attempt to explain the normative in terms of the non-normative.

Play with new ideas immediately. Students often think that they must read extensively before they can begin forming their own views. Not so: proper reading is continuous with original thinking. Hence, *from the very first text that you encounter*, try out your own ideas. Experiment with novel objections, arguments, theories, and taxonomies.

2. Reading: content

What to read depends very much on your career stage. In the pre-dissertation stage of your graduate career, *build from the foundations*. Start with the classics of the 20th century before reading the seminal works of the past few decades; then, if you are particularly interested in the topic, you might delve into contemporary articles. In philosophy of language, for example, I might start with Frege, Russell, Quine, and Kripke before moving on to more recent work. Focus on telling the forests apart; the names of the trees can wait.⁴

Change tack a year or so before you begin your dissertation. Specialize. Within your chosen area, find some narrow topic that is under discussion *right now* – one about which much has been published within the last 5 years – and read a dozen or two papers on it. Make sure that you have read *every* major paper on the topic, including every paper placed in a top journal or written by a major figure in the field.⁵ Also read some minor papers. Rinse. Repeat with another topic, which should still be within your area of specialization but which need not be obviously

⁴ I would offer very different advice to Masters students intending to apply to Ph.D. programs, since they need to publish more quickly. Such students should instead follow the advice below on specialized reading and can return to building from the foundations after being admitted to a Ph.D. program.

⁵ For a list of the best generalist journals, see <https://leiterreports.typepad.com/blog/2022/07/best-general-philosophy-journals-2022.html>. To identify the major figures in the field, look at who gets cited the most.

related to the first topic. Rinse and repeat again. By now you should be within 3-6 months of starting your dissertation, and you should have a clear idea for one article and a tentative thought about a second.

Once your dissertation is underway, you will likely need to read narrowly most of the time. But it pays to reserve some time, perhaps a month or two each summer, for reading in new areas. Still read with focus – read a dozen papers on parsimony or on Kant’s second *Critique* – but go far beyond the confines of your specialization.

For this practice of broad reading, I offer a pure and an impure rationale. Speaking purely, philosophy is a thoroughly holistic discipline. If you want to construct a powerful theory of the nature of reasons, you must understand philosophy of language, so that you can account for reasons-talk; metaphysics, so that you can account for the ontic status of reasons; and ethics and epistemology, so that you can account for practical and epistemic reasons.

Speaking impurely, if you read only what everyone else in your field is reading, then it will be hard to think what no one else is thinking. While I was writing my dissertation, I was scooped three separate times. These days, that rarely happens, for I have an unusual pattern of familiarity with the literature. I have read deeply about perception, phenomenal character, metaphysical grounding, knowledge-first epistemology, and internalism about practical reasons, so my ideas are unlikely to occur to others – not because I am cleverer or better-read, but because I am *differently* read.

So drill in far-flung places, but drill deep.

3. Writing: process

Writing a paper is like building a city: it takes time, and vast quantities of it. The occasional marathon session is not enough; good ideas must simmer in the subconscious over many months. Thus, I reserve a daily block of time for my writing. I treat that time as sacred: I do not let it get squeezed out by teaching, grading, reading, or emails. (It helps that I schedule my writing block as early as possible, usually starting at 9:00 AM.) I let this time go only during periods of vacation, illness, or emergency.

I write for about 10 hours per week: 1 hour per teaching day and 2.5 hours per non-teaching day. I have informally polled a dozen or so philosophers at top research institutions, and they consistently say that, during periods of peak productivity, they write for about 20 hours per week. I know one person who racks up 40 hours of writing per week. That person’s publication record is spectacular.

Writing is hard, and I often get stuck. When that happens, I do not try to solve the problem in my head; that, I have learned, is like waiting for Godot. For working memory is *extremely limited* – our minds can hold only 5-9 ideas at a time – so, if the solution is 10 steps away, it can never be found. Luckily, computer memory is effectively unlimited. Thus, I like to *write* my way through my problems. I start by describing the trouble: “My argument relies on premise *p*, but I don’t have any evidence for *p*.” Then I brainstorm solutions, jotting down any idea that occurs to me: “Maybe *q* is evidence for *p*.” If I realize that this solution doesn’t work, I can write that down, too: “No, *q* isn’t evidence for *p* because ...” I might then look for other evidence for *p* or try out a different argument that relies on some other premise. All of this, too, goes on the page. As a result, all of my work has been overhauled multiple times, and for every 10,000 words I publish, there are about 100,000 words that get scrapped. (Really. I have kept track.)

All of this is like navigating a labyrinth: I make many wrong turns, and even the right turn often leads away from the exit. Ample patience is what leads me to the light.

4. Writing: content and form

What constitutes excellence in philosophical writing? I say that it is the qualities listed below.

The list will be controversial, for it is the product of a controversial conception of philosophical excellence.⁶ I urge those with a different conception of philosophical excellence – most readers, I imagine – to articulate their own conceptions. I mean this seriously. Such an exercise forces a kind of reflection and precision that will be a great help to any philosopher.

The qualities are ordered by pedagogical importance (which is not the same as contribution to philosophical excellence) from highest to lowest. In the footnotes, I also identify works that model each quality. Study those models as a general might study the battles of Alexander and Napoleon.

Ambition. *This, above all.* Consider what it would take for a text to meet the highest philosophical standards. It would have to be, to the maximum extent possible, rigorous, novel, dialectically sensitive, economical, systematic, precise, cohesive, clear, and more. Measured against this ideal, the works of Plato and Nāgārjuna, Hume and Kant are – not bad. Aim higher. You will never hit your target, but there is nobility in the shot.⁷

Do not rush. Learn to write a reply, then a short original article, then a longer one, then a book. But always be working on what is just beyond your capacity, and let perfection be the horizon that you are striding towards.

Rigor. Support your thesis with *extremely compelling* evidence. This may take the form of a single very compelling consideration⁸ or an array of considerations that become very compelling when taken together.⁹

Limit yourself to dialectically effective evidence – evidence that even your opponent should accept.¹⁰ Therefore, avoid relying on all but the most robust intuitions.¹¹ Focus instead on theoretical considerations (simplicity, explanatory power, fruitfulness, etc.), scientific data, linguistic facts, and obvious ordinary truths,¹² or show that your theory is strictly better than your opponent's theory.¹³

Do not rest content with presenting a *prima facie* objection. Press your point as far as possible: consider all possible categories of response (see the discussion of systematicity [below](#)) and show that they all fail. Also point out fallback positions – e.g., show that an opponent who rejects a crucial premise of your argument may still accept some weaker premise which supports

⁶ It was Tamar Gendler who shared with me the basic outlines of this conception. I remember one conversation in particular: it occurred in Clare's Corner Copia on May 17, 2013 and amounted to the best instruction I have ever had in the art of philosophical writing. Below Tamar will find much of the advice she gave me then – especially in my descriptions of systematicity, significance, mastery of the literature, and authority – and, in keeping with that last virtue, I say that the conception articulated here is ultimately my own.

⁷ Fine (1994); Williamson (2000); Chalmers and Jackson (2001); Schaffer (2009) and (2016); Skorupski (2010); Greco (2012); Berker (2013); and Lormand (ms). These works perhaps fall short of the great works of history and certainly fall short of the ideal. But at least they *try*.

⁸ Williamson (2000, ch. 4); Bailey (2010).

⁹ Williamson (2000, ch. 11); Shah (2003); Kolodny (2005); Schaffer (2009, pp. 366–373) and (2016); Evans (2013); Moss (2013).

¹⁰ Balog (1999); Levine (2010); Greco (2012, §2, especially p. 350).

¹¹ Fine (1994); Johnston (2004); Schwitzgebel (2008).

¹² Chalmers and Jackson (2001); Moss (2013); Sinhababu (2015).

¹³ Lewis (1979); Sinhababu (2015).

a weaker but still interesting version of your conclusion.¹⁴ And, if there are multiple independent routes to your conclusion, say so.

You will often wish to rely on claims that you cannot rigorously support. In such cases, *jettison any arguments for those claims and introduce them as assumptions*.¹⁵ Do not hesitate to make controversial assumptions as long as you identify them as such; just briefly mention the motivations for them and move on. State all assumptions at the outset, however, as your reader will feel cheated if you help yourself to controversial claims once the argument is underway.

Novelty. Emphasize what matters most: your novel contribution. Review what has already been said only to set the stage for this.¹⁶

Novelty is often suppressed in the work of graduate students, where 80-90% of the word count is spent on a literature review and the slim remainder is novel. This is the reverse of the correct proportion.

To get the ratio right, remember that a complex idea can be stated in a hundred pages or a single sentence. Your summaries should mostly tilt towards the latter point, for your audience usually consists of experts who know the background terminology, positions, and arguments. Details are, however, required for those ideas that you will engage with closely. If your paper will undermine the last premise of a 4-premise argument, you might devote a sentence apiece to explaining why your opponent accepts premises 1-3, leaving plenty of citations in the footnotes, while explaining the case for premise 4 in patient detail.

Dialectical sensitivity. Your voice is the most recent contribution to a millennia-long conversation. Show a keen sensitivity to that fact.

The literature on a topic usually consists of a trickle of seminal works and a flood of subsequent works. Have a deep knowledge of all seminal works and of those subsequent works that you regard as excellent. Have at least a working acquaintance with almost everything else, especially recent works – say, everything published within the last 10 years or forthcoming. Spend extra effort to understand the influential ideas towards which you are least sympathetic.

At an intellectual level, display your dialectical sensitivity by making a new and illuminating map of the terrain – see the discussions of authority and systematicity below.¹⁷ At a more practical level, display your dialectical sensitivity through assiduous citations. Credit whoever developed the definitions, championed the views, and constructed the arguments that you are discussing. To minimize arbitrariness, order your citations by date of publication rather than alphabetically (though your reference list must still be alphabetized). Take special care to cite anyone who might referee your paper.¹⁸ If you do not know an area well, ask others what works are relevant.

None of this should make for clutter. Keep as many citations as possible in the footnotes – see the next virtue.

Authority. Be a master of the topic. Do not rely uncritically on the distinctions and definitions provided by other philosophers; be even more precise.¹⁹ Do not describe the

¹⁴ Moss (2013, §5.5).

¹⁵ Lormand (1996, pp. 52-53 and 61).

¹⁶ See almost anything published almost anywhere.

¹⁷ Lormand (1996) and (ms); Neta (2002, p. 664); Turri (2010).

¹⁸ Schellenberg (2014).

¹⁹ Lormand (1996, introduction and §1) and (ms, introduction and §1); Schellenberg (2010); Berker (2013, pp. 337-338 and 344-348); Millar (2016).

motivations for various positions as they are described by influential figures; be even more perspicuous. And do not just rebut the arguments of your opponents. Also state the best versions of their arguments, the versions that they *should* have given, and rebut those, too.²⁰

You are a philosopher, not a historian. (Again, this guide is for those working in contemporary philosophy.) So *focus on the ideas rather than the philosophers who birthed them*: in the main text of your paper, ruthlessly minimize quotations and references. You might need to show your reader that you are attacking widely endorsed theses, deflecting popular objections, and solving prominent problems, but do the bare minimum to achieve these purposes: shorten quotations, quarantine related quotations together, and shift all but the most crucial references to footnotes. Then get on with examining the ideas themselves. In fact, whenever possible, your main text should contain no quotations or references whatsoever.²¹

The footnotes are a completely different story: they should be rich in references, as per the discussion of dialectical sensitivity.

Economy. Some rocket scientists supply a spacecraft with the exact amount of fuel it requires to complete its mission. Construct your arguments in the same spirit: adjust your premises and conclusion until they match precisely.

Working from one end, find *the strongest conclusion* supported by your premises (or the broadest class of views targeted by your objection, or ...). A well-constructed argument will not target only the very specific view of one prominent philosopher. It will target a substantial class of views. Delineate that class. Also consider whether the argument can target a much larger class of views with only slightly stronger premises.²²

Working from the other end, find *the weakest premises* which support your conclusion. Whenever possible, replace controversial premises with less controversial ones. Also consider whether the premises can be weakened substantially while weakening the conclusion of the argument only slightly.²³

Continue this process of adjustment until your argument is perfectly tuned.²⁴

Be economical with words, too. When it fits – and it does not always – say “epistemic,” “ontic,” “semantic,” “social,” and “phenomenal” rather than “epistemological,” “ontological,” “semantical,” “societal,” and “phenomenological.”

Systematicity. Minimize arbitrariness. When laying out possible views on some topic, do not just list the currently influential views in the literature. Instead develop a *productive taxonomy* – again, one that divides possible views into categories that are mutually exclusive, jointly exhaustive, and separated by important fault-lines – and situate all currently influential views within it.²⁵ When stating a problem, do not merely list some claims that are hard to reconcile. Instead identify n claims such that each is very plausible, any $n-1$ of them are consistent,

²⁰ Lewis (1984).

²¹ Although main-text references to historical philosophers have more value than main-text references to contemporary philosophers. That is not because you should defer to Plato or Kant – see the discussion of ambition – but because referring to historical figures helps locate your view against others at the highest level of generality and has the literary value of allusion to boot. See for example Yablo (1997); Street (2010); Millar (2016).

²² Berker (2013) leverages a single core idea against a very broad class of theories.

²³ Neta (2002), Sinhababu (2009), and Lormand (ms, §4) do much work with sparse resources.

²⁴ Bailey (2010); Turri (2011); Sinhababu (2015).

²⁵ Street (2006, §2) and (2010, pp. 369-370). Aristotle and Kant are also wonderfully systematic writers who use taxonomies to strong effect. Arguments by dilemma are taxonomies put to a special use: Chalmers (2007, pp. 173-179).

but all n are inconsistent.²⁶ When addressing objections, do not merely list various objections or opposing views that occur to you. Instead divide objections or opposing views into a taxonomy of salient clusters and address each salient cluster of objections.²⁷ When attacking a view, do not merely provide the telling counterexample. Also identify the relevant class of counterexamples, diagnose the problematic feature of the opposing view that leaves it open to counterexamples of this class, and show that your view lacks this problematic feature.²⁸ Even when transitioning between paragraphs or sections, make it implicitly clear why the new topic is next on the agenda.

Precision. A precise claim is specific and expresses the writer's intended idea. A claim may therefore be imprecise in two ways: it may be vague rather than specific, or it may fail to express the writer's intended idea irrespective of its specificity.²⁹ Avoid these mistakes.³⁰

Metaphorical language, though often a great help in other respects, can disguise imprecision, so make sure that you always know what you literally mean.³¹ Formal tools can in contrast help you be very precise, but do not use them when ordinary English will suffice.

Occasionally, imprecision is for the best – for instance, when it makes for easy communication and does not affect your argument.

Cohesiveness. This virtue does not attach to a single view but to a system of views – a **theory**. A cohesive theory consists of mutually supporting views, views whose collective explanatory power is much greater than the sum of their individual explanatory powers.³²

The virtues of economy and cohesiveness might seem to compete. They do not: though a good cohesive theory will have many posits, these will still be the sparsest posits that can meet all relevant constraints.

Clarity. At every point in your article, your reader should know what has happened, what is happening, and what will happen.

It helps to be generous in your use of **signposts** – devices that indicate the relationships among ideas.³³ These include phrases like, “In this section, I will argue that ...” and words like “thus,” “since,” “because,” “however,” “but,” “therefore,” “nevertheless,” “despite,” “further,” and “consequently.” Signposts are not always words; they can include paragraph and section breaks, punctuation marks (colons, dashes, parentheses), and more.

It also helps to abstract away from the details at crucial points. In your paper's introduction, articulate, as starkly as possible, the central philosophical problem(s) that you will discuss and the central solution(s) that you will propose. In your paper's conclusion, recast your solution(s) in a way that the reader can appreciate only now that they have reached the end. And, in your paper's body, descend into details and complications when your argument requires them; otherwise remove them. Likewise for the introductions, conclusions, and bodies of your paper's individual sections.³⁴

²⁶ Egan and John (ms, p. 1).

²⁷ Millar (2016).

²⁸ Gendler (2002, §3); Enoch (2006, §5). Turri (2011) carries out a positive analogue of this: he articulates with great precision how a certain class of views arguably avoids the Gettier problem.

²⁹ Obviously, a claim may be imprecise in both ways at once.

³⁰ Lormand (1996) and (ms); Williamson (2000, ch.4); Bailey (2010); Schaffer (2016).

³¹ Schaffer is a master of the precise metaphor. See his (2009) and (2010).

³² Williamson (2000); Neta (2002); Skorupski (2010); Graham (2012); Moss (2013); Schaffer (2016); Lord (2017).

³³ Lewis (1979); Schaffer (2009), (2012), and (2016).

³⁴ Fine (1994); Lewis (1979); Lormand (1996) and (ms); Schaffer (2009) and (2016); Berker (2013); Evans (2013); Lord (2017).

Concreteness. It is not in our nature to think in the abstract: see the [earlier](#) discussion of using concrete cases to understand abstract ideas. Thus, as an aid to the reader’s understanding, make generous use of concrete examples to supplement (but not to replace) abstract thought.³⁵

Flair. Seize opportunities for the wry remark, the keen observation, the striking metaphor.³⁶ And make your examples memorable, perhaps with humor, allusion, or picturesque description. Stock examples are a waste.

Do not, however, use flair at the expense of clarity or precision. And be sparing. Flair is like salt: pleasing in moderation and noxious in excess.

5. Writing: heuristics

Brilliant ideas do not come from mystical chanting, but from simple heuristics. Here are some productive ones. The last heuristic is by far the most important one and thus receives by far the most space.³⁷

The regimentation heuristic: *Put all arguments, whether yours or your opponents’, in premise–conclusion form.*

For reasons given in §1, practice this heuristic religiously. Apply it even to objections and replies: these are just arguments to the effect that another argument is faulty (that it has a false or unsupported premise, involves a bad inference, etc.).

The web heuristic: *Start with something important. Then trace the web of other important things that (arguably) relate to it and use that to generate productive questions about the original thing.*

Example. Start with knowledge. This is (arguably) related to justification, evidence, defeat, Gettier cases, lottery cases, norms of assertion, and causal explanations of action. That raises a series of productive questions about any theory of knowledge: what does it imply about all of these other things, and are those implications plausible or not?

The new work heuristic. *Seek new work for old theories.* This is perhaps a special case of the web heuristic.

Example. It has been argued that knowledge helps to causally explain action, but this view has primarily been defended by *non-skeptical invariantists* about knowledge. Thus, you might ask whether other theories can also use knowledge to provide causal explanations of action, and if so, how those explanations compare to those given by non-skeptical invariantists. This question is what prompted the main idea of my 2024 paper, “Invariantism, contextualism, and the explanatory power of knowledge,” where I argue that contextualists do even better than non-skeptical invariantists on this front.

³⁵ Lewis (1979); Gendler (2002, p. 47) and (2008); Sinhababu (2015).

³⁶ Williamson (1995, first sentence); Yablo (1997, fn. 17); Schaffer (2009, p. 355), (2016, pp. 92-93), and (2017, penultimate paragraph of §3.1); Schwitzgebel (2012, pp. 39-40).

³⁷ More heuristics are discussed in Hájek (2014). He recommends, among other things, checking whether a definite description has a unique referent, examining extreme or near-extreme cases, and considering cases of self-reference.

The replacement heuristic: *State an important idea in one short sentence, underline key expressions, and try replacing those with related expressions.*³⁸

I cannot overstate the power of this heuristic. The best thing I have written is my book *A Pluralist Theory of Perception*, and that came largely from systematic iterations of the replacement heuristic.

This heuristic can generate new evidence from old. For instance, recall this earlier piece of evidence:

Perceptions enable us to learn what certain things are, in and of themselves.

Let us apply the replacement heuristic here. First, we underline some key expressions:

Perceptions enable us to learn what certain things are, in and of themselves.

Then we generate new claims by replacing one of the underlined expressions – say, “perception” – with a similar expression. Three easy moves are to step up (from “perception” to “sensory state”), down (from “perception” to “conscious perception”) or sideways, to:

Hallucinations enable us to learn what certain things are, in and of themselves.

Now we reflect on whether this claim is true or false. Either way, we get new evidence: we discover either a symmetry, or an asymmetry, between perception and hallucination. In this case, I believe that the claim is true: if you have never experienced anything red, and you hallucinate a red thing, then this experience, too, would seem to enable you to learn *what redness is, in and of itself*.

Evidence is valuable. Thus, it is worth iterating the heuristic many times. For instance, we can also generate these claims:

Perceptions enable us to learn what certain objects are, in and of themselves.

Perceptions enable us to learn what certain properties are, in and of themselves.

Hallucinations enable us to learn what certain objects are, in and of themselves.

Hallucinations enable us to learn what certain properties are, in and of themselves.

I will glide over the details here, but I believe that reflection supports the first, second, and fourth of these claims while undermining the third one. If so, then there is both an interesting *symmetry* between perceptions and hallucinations – both enable us to learn what certain properties are, in and of themselves – and an interesting *asymmetry* between them – perceptions, but not hallucinations, enable us to learn what certain objects are, in and of themselves.

Once we have discovered new evidence, we can use that to test existing theories. And, if those theories prove inadequate, we can use the replacement heuristic to generate better ones. Consider, for instance, these popular naïve realist views:

To perceive an object and its properties is to stand in a primitive, non-representational relation of awareness – *acquaintance* – to that object and those properties.

³⁸ John Turri and Jonathan Schaffer are especially proficient with the replacement heuristic. See for example Turri (2010); Schaffer (2005a), (2005b), and (2012).

To hallucinate is not to stand in any relation of awareness to anything.

The second claim is arguably inconsistent with the evidence: if hallucinations did not make us aware of anything, then they could not enable us to learn what any properties were, in and of themselves. Let us try to use the replacement heuristic to improve the theory. We first underline some key expressions:

To hallucinate is not to stand in any relation of awareness to anything.

We then replace these with similar expressions with an eye towards explaining our new evidence. This might eventually lead us to this new naïve realist theory:

To perceive an object and its properties is to stand in a primitive, non-representational relation of awareness – *acquaintance* – to that object and those properties.

To hallucinate is to stand in a primitive, non-representational relation of awareness – acquaintance – to certain properties (but not to any objects).

This arguably explains the evidence. Both perceptions and hallucinations enable us to learn what certain properties are, in and of themselves, because both perceptions and hallucinations acquaint us with properties; but only perceptions enable us to learn what some objects are, in and of themselves, because only perceptions acquaint us with any objects.³⁹

This is just the beginning of the story. We have much more evidence about perceptions and hallucinations, and we can run that through the replacement heuristic to generate far more evidence. In addition, there are plenty of other theories besides the naïve realist theories above, and we can also test them against our evidence, old and new. If these theories prove inadequate, then we can use the replacement heuristic to generate alternatives.

Here is one last example of the replacement heuristic in action. Consider this influential view in epistemology:

Asserting p is epistemically permissible just in case one knows p .

We can generate new views from this one. First, we underline some key expressions:

Asserting p is epistemically permissible just in case one knows p .

Then we replace the underlined expressions with related ones. For instance, we might replace the first expression:

Acting on p is epistemically permissible just in case one knows p .

Believing p is epistemically permissible just in case one knows p .

Or the second expression:

Asserting p is epistemically good just in case one knows p .

Asserting p is epistemically virtuous just in case one knows p .

³⁹ This is the main line of thought in Johnston (2004).

Asserting p is epistemically successful just in case one knows p .

Or the third expression:

Asserting p is epistemically permissible just in case one justifiably believes p .
Asserting p is epistemically permissible just in case one is certain of p .

Lo and behold, most of these views have been defended in the literature.⁴⁰

We can also iterate the heuristic to generate more distant views – for example:

Asserting p is epistemically successful just in case one knows p .
Acting on p is epistemically successful just in case one knows p .
Believing p is epistemically successful just in case one knows p .

I defend these last three views in my 2016 paper, “Knowledge and other norms for assertion, action, and belief,” and much of the evidence that I use to support them also came from the replacement heuristic.

Three final notes.

First, do not just use these heuristics in isolation. The real power comes when you iterate each heuristic multiple times and chain different heuristics together.

Second, even the best heuristics mostly generate junk, so do not get attached to whatever ideas come first. Generate *lots* of ideas. Then sift through the dirt for those rare flecks of gold.

Finally, *evidence comes first*. Do not generate a new view and then try to invent new evidence to support it. Start by generating new evidence and, if the old views do not fit the total evidence, generate new views that do. (It is also acceptable to generate a new view and then, with an open mind, seek evidence that would confirm or disconfirm it.)

6. Writing: publishing strategies

Philosophy journals tend to have acceptance rates in the 5%-15% range – call it 10% on average. That means that when you submit a paper, you are playing poker. So *think like a poker player*.

Do not look to get any particular paper accepted at any particular journal; rejection is always the most likely result. Have a heart of stone about that. You are aiming to *cultivate an edge* – an advantage in your acceptance rate versus the average – *and apply it repeatedly*.

Even a small edge is massive in the long run. My acceptance rate, across all of my 119 submissions, is 15%, so my edge is just 5%. But that is *huge*: it means a 50% increase in publications over the long term. You should aim higher. My guess is that among top-tier philosophers, an acceptance rate of 30% is common enough, and perhaps the very best get to 40%. Every paper you write can be clearer, more rigorous, more systematic. Always look for the next 1% bump.

Does this mean that you should flood journals with submissions? No. Look at my lifetime acceptance rates for original articles:

At the first venue they're submitted to: 0%.

At the second venue they're submitted to: 12.5%.

⁴⁰ See the extensive literature that grew from Williamson (2000, ch. 11).

At the third venue they're submitted to: 11.1%.
At the fourth venue they're submitted to: 14.3%.
At the fifth or higher venue they're submitted: 21.9%.

Why is the first number at rock bottom? Not for want of effort; my first submissions are as strong as I know how to make them. It is because my papers always have *some* serious weaknesses that I miss and that referees notice. (Philosophy is hard; no one can think of everything.) Why do the numbers tick up so much? Because, when those weaknesses are pointed out to me, I make deep revisions to shore them up, so the more rejections my papers accumulate, the better they get. This makes a massive difference to the odds – at least 21.9%. So focus on quality, not quantity. To publish in high quantities, you must write *many* papers of high quality.

The path to a quality paper is not easy: I have to sit with negative referee reports, and that kind of blunt criticism is not pleasantly taken. But take it I do, for although I have received terrific comments from mentors and friends, referees have been my *single most helpful* source of comments: ordinarily, they are experts on the topic, they read with great care, and they offer extensive and sharply critical feedback.

Still, when I first receive negative referee reports, I am too upset to begin revisions. I must wait a day or two before I can read them with clear eyes. Then I give each comment my most careful consideration. Rarely, I decide that a comment is just misguided and should be ignored. More commonly, I think that a referee has missed the point, but that I can ensure that it is not missed again. Most commonly of all, however, I conclude that the referee has found an important weakness in my argument, and I pop the hood for major repairs.

I can recommend some strategies for increasing the quantity of what you write while maintaining its quality. It takes a lot of time to read and think deeply on any topic, so maximize your investment: write a series of papers on a single topic rather than flitting from one to the next. Better yet, have a **project**: a philosophical idea large enough that many articles are required for its proper development. Working on a project will, as a happy byproduct, improve your thinking as you settle into it. In addition, when writing articles, whether or not these are part of a project, limit each to a single sharply demarcated topic, as per the earlier discussions of focus and economy. Similarly, see a single paper to completion rather than working on several at once. This will let you submit papers much more rapidly.

Know that it is a long road between putting finger to keyboard and having your article accepted. It takes me about 4 months to put together a first submission, and, across all of my publications, there has been a median gap of 22 months between first submission and final acceptance. That brings the total to *just over 2 years*.

Hence, for the sake of your career, you must start writing years before key milestones. You should have several publications before you go on the job market, and you will probably send out job applications almost a full academic year before you defend your dissertation, so you should begin to assemble your first article as soon as you start your dissertation, if not earlier. (These two activities should be largely co-extensive.) Submit that article for review within 6 months. Also, for most tenure-track positions, you will need to submit your tenure file around the start of your fifth or sixth year, and by then you must have a substantial body of work. So, if you land a job like that, you will have little time to spare. Continue to submit parts of your dissertation for review as you complete them. Once your dissertation is complete, allow yourself a month to celebrate; then start new work *immediately*.

This is all hard work. But it should be joyful. Find your pleasure not primarily in the distant prospect of publication, but in the proximal act of writing itself – the act of philosophical discovery.

7. Cultivating yourself as a philosopher

This writing guide began with an autobiography, and it is time to finish the tale. After my miserable first outings in the world of publishing, I realized that I lacked crucial writing skills, and I resolved to fill in the gaps. I began by requesting all the feedback that I could get.

I met an obstacle straightaway. I would learn that I should include a particular citation, trim a particular section, and provide more evidence for a particular claim. But, although I could follow these pieces of advice, I had no idea what *general principles* were in operation. Thus I just repeated my mistakes elsewhere.

It was here that my lack of talent was made manifest, for talent, I conjecture, consists largely in the ability to pick up general principles from mere hints and examples. Like a talented dancer who sees a dramatic performance of the tango and *just knows* how to dance like that, a talented philosopher sees the ambitiousness of Plato and the systematicity of Kant and *just knows* how to write like that. The rest of us need to be told to lengthen our strides, square our hips, lift our chins.

Without the benefit of talent, I had to try to discover the *precise but general principles* of philosophical excellence. When advised to include a citation, trim a section, and provide more evidence for a claim, I attempted to articulate the precise general principles that warranted including that citation, trimming that section, and providing that evidence. Similarly, I sought exemplars of philosophical writing – for instance, I dug into the work of several philosophers who were consistently publishing in top journals – and attempted to articulate the precise general principles at play there. I also asked successful philosophers to tell me what works they regarded as best and why, what they saw as the best features of their own work, and what their processes of reading and writing were like.

I recorded what that I found in a document intended for personal use. That eventually expanded into this writing guide.

In these attempts to improve my writing, I was engaging in reflection at three levels. I was reflecting, first, on what makes for excellent philosophical *work*; second, on what makes for an excellent *process* of producing philosophical work; and third, on how a *person* can, without relying on talent, identify and learn those processes. I think of these levels as corresponding to goods, rules, and virtues that are distinctive to philosophy. My most basic advice on self-cultivation, then, is to carry out your own reflections across all three levels.

One practical method for doing that is to *make your own guide to doing philosophy*. (If you have internalized my previous advice about ambition and authority, then you will surely doubt much of what you find here.) Start small: whenever you read a great work of philosophy, write down what makes it great in precise but general terms. There's no need to go it alone; discuss your ideas with friends and mentors. You can reserve your guide for personal use, but articulating your conception of philosophical excellence as precisely as you can will give you much more control over your work.

Cultivate yourself not only as a writer, but also as a thinker. Read broadly as well as deeply, taking time to appreciate great historical works. And be fearless about developing new skills at every stage of your career. It is never too late to acquaint yourself with Sanskrit, modal logic, or vision science.

Make sure to have a life beyond the profession, too. For us, philosophy is a part of the good life, but it is only a proper part. I rarely put in more than 40 hours per week of research, teaching, and service, and I make it my personal ideal to have satisfying personal relationships, practice hobbies, eat well, exercise regularly, meditate occasionally, and sleep plenty.

Do I manage all of that? Never! But I come close enough (and that is not particularly close) to be content.

ACKNOWLEDGMENTS

To Tamar Gendler, Eric Lormand, and Andy Egan – you taught me to think and to write. I remember, and I am grateful.

To Andrew Bailey, Jake Berger, Selim Berker, Ben Blumson, Amber Carpenter, Victor Caston, Jay Garfield, Fabian Geier, Sam Liao, Nico Silins, Neil Sinhababu, and Matt Walker – you might all think that my advice in this guide is mad. But at least it is less mad because I have talked to you.

REFERENCES

- Bailey, A. (2010). “Warrant is unique.” *Philosophical Studies* 149 (3): 297-304.
- Balog, K. (1999). “Conceivability, possibility, and the mind-body problem.” *Philosophical Review* 108 (4): 497-528.
- Berker, S. (2013). “Epistemic teleology and the separateness of propositions.” *Philosophical Review* 122 (3): 337-393.
- Campbell, J. (2002). *Reference and Consciousness*. Oxford: Oxford University Press.
- Chalmers, D. (2007). “Phenomenal concepts and the explanatory gap.” In *Phenomenal Concepts and Phenomenal Knowledge*, eds. T. Alter and S. Walter (New York: Oxford University Press), pp. 167-194.
- Chalmers, D. and Jackson, F. (2001). “Conceptual analysis and reductive explanation.” *Philosophical Review* 110: 315-361.
- Egan, A. and John, J. (ms). “A puzzle about perception.”
- Enoch, D. (2006). “Agency, schmagency: Why normativity won’t come from what is constitutive of action.” *Philosophical Review* 115 (2): 169-198.
- Evans, I. (2013). “The problem of the basing relation.” *Synthese* 190 (14): 2943-2957.
- Fine, K. (1994). “Essence and modality.” *Philosophical Perspectives* 8: 1-16.
- Gendler, T. (2002). “Personal identity and thought-experiments.” *Philosophical Quarterly* 52 (206): 34-54.
- Gendler, T. (2008). “Alief and belief.” *Journal of Philosophy* 105 (10): 634-663.
- Graham, P. (2012). “Epistemic entitlement.” *Noûs* 46: 449-482.
- Greco, D. (2012). “The impossibility of skepticism.” *Philosophical Review* 121 (3): 317-358.
- Hájek, A. (2014). “Philosophical heuristics and philosophical creativity.” In *The Philosophy of Creativity: New Essays*, eds. E. Paul and S. Kaufman (Oxford: Oxford University Press).
- Johnston, M. (2004). “The obscure object of hallucination.” *Philosophical Studies* 120: 113-83.
- Johnston, M. (2006). “Better than mere knowledge? The function of sensory awareness.” In *Perceptual Experience*, eds. T. Gendler and J. Hawthorne (Oxford University Press), pp. 260-290.
- Kolodny, N. (2005). “Why be rational?” *Mind* 114 (455): 509-563.
- Levine, J. (2010). “The Q factor: modal rationalism vs. modal autonomism.” *Philosophical Review* 119 (3): 365-380.
- Lewis, D. (1979). “Attitudes *de dicto* and *de se*.” *Philosophical Review* 88: 513-543.
- Lewis, D. (1984). “Putnam’s paradox.” *Australasian Journal of Philosophy* 62 (3): 221-236.
- Lord (2017). “What you’re rationally required to do and what you ought to do (are the same thing!).” *Mind* 126 (504): 1109-1154.

- Lormand, E. (1996). "How to be a meaning holist." *Journal of Philosophy* 93: 51-73.
- Lormand, E. (ms). "How to be a meaning atomist." At <http://www-personal.umich.edu/~lormand/phil/meaning/atomism.htm>.
- Mehta, N. (2016). "Knowledge and other norms for assertion, action, and belief." *Philosophy and Phenomenological Research* 91 (1).
- Millar, B. (2016). "Frege's puzzle for perception." *Philosophy and Phenomenological Research* 93 (2): 368-392.
- Moss, S. (2013). "Epistemology formalized." *Philosophical Review* 122 (1): 1-43.
- Neta, R. (2002). "S knows that p." *Noûs* 36 (4): 663-681.
- Russell, B. (1911). "Knowledge by acquaintance and knowledge by description." *Proceedings of the Aristotelian Society* 11 (1): 108-128.
- Schaffer, J. (2005a). "Contrastive causation." *Philosophical Review* 114 (3): 327-358.
- Schaffer, J. (2005b). *Oxford Studies in Epistemology* 1.
- Schaffer, J. (2009). "On what grounds what." In *Metametaphysics: New Essays on the Foundations of Ontology*, eds. Chalmers, D., Manley, D., and Wasserman, R. (Oxford: Oxford University Press).
- Schaffer, J. (2010). "Monism: the priority of the whole." *Philosophical Review* 119 (1): 31-76.
- Schaffer, J. (2012). "Grounding, transitivity, and contrastivity." In *Metaphysical Grounding: Understanding the Structure of Reality*, eds. F. Correia and B. Schnieder (Cambridge: Cambridge University Press).
- Schaffer, J. (2016). "Grounding in the image of causation." *Philosophical Studies* 173 (1): 49-100.
- Schaffer, J. (2017). "Social construction as grounding; or: Fundamentality for feminists, a reply to Barnes and Mikkola." *Philosophical Studies* 174: 2449-2465.
- Schellenberg, S. (2010). "The particularity and phenomenology of perceptual experience." *Philosophical Studies* 149: 19-48.
- Schellenberg, S. (2014). "The relational and representational character of perceptual experience." In *Does Perception Have Content?*, ed. Brogaard, B., Oxford: Oxford University Press.
- Schwitzgebel, E. (2008). "The unreliability of naive introspection." *Philosophical Review* 117 (2): 245-273.
- Schwitzgebel, E. (2012). "Introspection, what?" In *Introspection and Consciousness*, eds. D. Smithies and D. Stoljar (Oxford: Oxford University Press).
- Shah, N. (2003). "How truth governs belief." *Philosophical Review* 112 (4): 447-482.
- Sinhababu, N. (2009). "The Humean theory of motivation reformulated and defended." *Philosophical Review* 118 (4): 465-500.
- Sinhababu, N. (2015). "Advantages of propositionalism." *Pacific Philosophical Quarterly* 96 (1): 165-180.
- Skorupski, J. (2010). *The Domain of Reasons*. Oxford: Oxford University Press.
- Street, S. (2006). "A Darwinian dilemma for realist theories of value." *Philosophical Studies* 127: 109-166.
- Street, S. (2010). "What is constructivism in ethics and meta-ethics?" *Philosophy Compass* 5 (5): 363-384.
- Turri, J. (2010). "Prompting challenges." *Analysis* 70: 456-62.
- Turri, J. (2011). "Manifest failure: The Gettier problem solved." *Philosophers' Imprint* 11 (8): 1-11.
- Williamson, T. (1995). "Is knowing a state of mind?" *Mind* 104 (415): 533-565.
- Williamson, T. (2000). *Knowledge and Its Limits*. Oxford: Oxford University Press.

Yablo, S. (1997). "Wide causation." *Philosophical Perspectives* 11: 251-281.