

A writing guide for professional philosophers¹

by Neil Mehta, Yale-NUS College

TABLE OF CONTENTS

0. Introduction	1
1. Reading and thinking: process	2
2. Reading: content	3
3. Writing: content and form	4
4. Writing: process	9
5. Writing: publication strategies	10
6. Cultivating yourself as a philosopher	12

This is a guide to writing for professional philosophers – in particular, a guide to writing *articles* in *contemporary* philosophy. I plead ignorance of writing books and of writing as a historian of philosophy, though I hope that some of my suggestions will generalize to these cases. While the guide is pitched primarily for graduate students, perhaps it will be of some value to philosophy professors as well. I provide it because little information on this topic is widely available, because many philosophy graduate programs, even excellent ones, do not systematically address this topic, and because excellence in general writing is woefully insufficient for excellence in professional philosophical writing.²

It is often thought that excellent philosophical writing is the sole province of the genius. I by contrast believe that writing is a skill and can therefore be taught and learned; in any case, if there is some special talent for writing then I do not have it. When, near the end of my graduate career, I first 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*. Even that paper was merely a reply, and not a particularly ambitious one. I then undertook a disciplined program to improve my writing, as described in §5, and over the next several years my work climbed steadily in quality. I have now published 8 articles in venues including *Journal of Philosophy*, *Philosophers' Imprint*, and *Philosophy and Phenomenological Research*. I make no claim to be a great philosophical writer: I do not take my résumé to be especially impressive, and a résumé rarely tells much anyway given the vagaries of the review process. I claim only that I have come far in my writing, that I can articulate precisely how I have done so, and, most importantly, that my methods are highly replicable.

I begin by discussing the central prerequisites of excellent writing: reading and thinking (§1, §2). The core of the guide is dedicated to the content and form of excellent writing (§3), though I also briefly discuss the process of writing (§4) and the two major poles of publication strategies (§5). I end with comments on how to cultivate oneself as a philosopher (§6).

I welcome questions, suggestions, and criticisms. (Praise would be not just welcome but surprising.) Please feel free to contact me at neil.jag.mehta@gmail.com.

¹ Version 1.1; last revised September 23, 2016. Updated versions will be posted periodically at www.profneilmehta.com.

² For further advice, see Brooks (ms) and Lin (ms).

1. Reading and thinking: process

Students sometimes think that they must first read extensively before they may begin forming their own views. Not so: proper reading is continuous with original thinking.

Hence, *from the very first article that I encounter in a new area*, I begin the critical work of gathering evidence, which I treat as candidate *explananda*; identifying hypotheses, which I treat as candidate *explanantia*; uncovering assumptions; and sketching out my own tentative ideas and arguments. I also attempt to identify the central questions and influential answers, which I then try to subsume into a taxonomy of all possible views, as per the discussion of systematicity in §3. I have no fear of trying my hand at these ambitious tasks from the start because even failure is handsomely rewarded. For I constantly revise my understanding as I read new work, and without a tentative understanding I would have nothing to revise.

I approach any new area with a set of general categories under which I can fruitfully subsume questions, theses, argumentative strategies, etc. For example, I often begin by trying to identify the central **metaphysical**, **epistemic**, and **semantic** questions in an area, as well as the connections among possible answers to those questions. I also try to classify views in terms of what **explanatory relations**, such as relations of explanatory priority, they posit.³ Relatedly, I find it useful to consider what a theory takes as metaphysically or conceptually **fundamental** and **non-fundamental**.

Regarding the foundations of normativity, for example, consider the relationship between facts about reasons and facts about rationality. One of the following views must be true: either the former always explain the latter, or the latter always explain the former, or there is no uniform explanatory relation between the facts of these two types. When I encounter a new view about the relationship between reasons and rationality, I locate it in this taxonomy – a taxonomy that appears in the literature, but that I independently developed as an aid to my thinking. Similarly, I consider which normative entities, if any, are treated as fundamental by which theories. The anti-reductivist holds roughly that some normative entity is metaphysically fundamental, and anti-reductivists have variously identified goodness, rightness, reasons, obligation, or virtue as fundamental. Though reductivists hold in contrast that no normative entity is fundamental, it is still useful to consider which normative entities they take as **relatively fundamental**, i.e., more fundamental than any other normative entity.

While reading, I rely on certain heuristics to generate hypotheses worth exploring. My most general heuristic is **replacement**: replacing a salient term in some philosophical thesis with a related term. Three further heuristics, though they usually stand under the umbrella of the replacement heuristic, are noteworthy in their own right: **extension** – applying a similar thesis to a nearby topic; **generalization** – applying a broader thesis to a topic that includes the topic of the original thesis as a special case; and **unification** – identifying a single thesis that entails several existing theses.

Take the thesis that an assertion that p is epistemically permissible just in case one knows that p . I might consider several extensions of this thesis: e.g., that a *belief* that p , or a *treatment of the proposition that p as a reason for action*, is epistemically permissible just in case one knows that p . I might also try extending that last thesis: perhaps a treatment of the proposition that p as a reason for *belief* is epistemically permissible just in case one knows that p . I might then attempt to unify the last two theses: perhaps a treatment of the proposition that p as a *reason, full stop*, is epistemically permissible just in case one knows that p . All the while, I might consider various options for replacement. Perhaps the relevant normative notion is not *epistemic permissibility*, but

³ I suspect that David Chalmers uses a similar heuristic given the remarks in his (2011, pp. 538-539).

*epistemic goodness, epistemic success, or epistemic virtuousness.*⁴ Once I have generated a rich set of possible hypotheses, I go on to test their predictions against the evidence.

On to brass tacks about the reading process, then. I use a detailed system of notation whenever I read: I double-underline any technical term, I mark key examples with an asterisk, I underline any reasonably significant claim once, and I mark especially central claims with an arrow, using triple arrows to mark the major theses of the entire article or book chapter. On top of that, after completing each reading I summarize it, sometimes in great detail, in a separate Word document. Because of this documentation, I know that I have read more than 1000 articles and book chapters in the past 8 years, though I am aware of some very successful philosophers who read much less and others who read much more. I prefer to read fewer texts deeply than to read many texts shallowly. Hence my default is to read any article or book chapter twice, typically spending 45-60 minutes per reading, though if I find a text disappointing I read it only once.

Obviously, what is important is not to adopt this idiosyncratic system. You should just develop *some* method for engaging deeply with texts. As I see it, reading a philosophical text is less like pouring water into a bucket and more like wrestling an ill-tempered bear. I encourage anyone to read in a similar spirit.

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 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 work published in the last few decades. Prioritize learning to tell the forests apart; the names of the trees can wait.⁵

Change tack a year or so before you begin to write your dissertation. Specialize. Within your chosen area (metaphysics, political philosophy, aesthetics), find some narrow topic that is under discussion *right now* – ideally, a topic on which much has been published within the last 5 years, but on which little was published 10-20 years ago – and read one or two dozen papers on it. Make sure that you have read *every* major paper on that topic, including every paper that appears in a top-quality journal and every paper written by a major figure in the field; also read many minor papers.⁶ Rinse. Repeat with another topic, which should still be within your area of specialization but which need not be obviously related to the first topic. Rinse and repeat one last time. 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. At this stage, that is plenty.

⁴ It was by using these heuristics that I developed the ideas for my (2015). I also regard John Turri as a master of the replacement heuristic, e.g. in his (2010), and Schaffer has run the idea of contrastivity through the extension heuristic many times, e.g. in his (2005a), (2005b), and (2012). For a discussion of other heuristics, see Hájek (2014), who 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.

⁵ I would offer very different advice to Masters students intending to apply to Ph.D. programs, who should aim to publish much 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. My advice is also intended primarily for students in North America, as students in the U.K. must specialize much earlier.

⁶ For a poll that illuminates which generalist journals are most highly regarded, see <http://leiterreports.typepad.com/blog/2015/09/the-top-20-general-philosophy-journals-2015.html>. For information about the most influential figures in the field, pay attention to citations, and ask experts – especially your dissertation adviser.

Once your dissertation is underway, you must of course continue to read in this narrow way most of the time. But you *must* 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 as a theoretical virtue 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 itself has a thoroughly holistic character and so cannot be expertly handled one narrow topic at a time. To do truly outstanding work on the nature of reasons, you must understand philosophy of language, so that you can analyze reasons-talk, metaphysics, so that you can discuss the ontic status of reasons, and ethics and epistemology, so that you can identify and explain facts about practical and epistemic reasons.

Speaking impurely, if you read only what everyone else in your field is reading then you will find it very hard to think something that no one else in your field is thinking. In the four years dating from the start of my dissertation, I was scooped no less than 3 times. Once I found a major idea of mine already published in a paper that I had overlooked, and twice more I was beaten to the presses by a forthcoming paper. To be clear: these were not instances of plagiarism. They were instances of independent convergence on a natural idea. But after learning this lesson I have rarely been scooped, and I credit the change to my unusual pattern of familiarity with the literature. Because by now I have read deeply on perception, phenomenal character, metaphysical grounding, knowledge-first epistemology, and internalism about practical reasons, my ideas are very unlikely to occur to others – not because I am cleverer or even better-read, but just because I am *differently* read.

So drill in far-flung places, but drill each hole *deep*.

3. Writing: content and form

In this section, which is the heart of this guide, I list those virtues pertaining to content and form that I regard as most central. Along with a description of these virtues, I include many references to exemplary works, and though my emphasis is on contemporary articles, I could not resist including a few books and historical works. I encourage the reader to study these works as a general might study the battles of Napoleon and Alexander: virtue is best learned from models.

The list will be very controversial, for it is the product of a very controversial conception of philosophical excellence.⁷ I urge those with a different conception of philosophical excellence – that will be most readers, I imagine – to articulate their own conceptions and to identify their own models. I mean this seriously. Such an exercise forces a level of reflection, and a degree of explicitness, that will be a great help to any writer.

I start with those virtues that I take to be most worthy of deliberate cultivation – which are not quite the same as those virtues that most contribute to philosophical excellence – and I descend from there.

⁷ 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.

Ambition. *This, above all.* Consider what it would take for a text to meet the highest standards of philosophical excellence. It would have to be profoundly significant, original, systematic, economical, precise, lucid, accessible, rigorously argued, and more. Measured against this ideal, the works of Plato and Nāgārjuna, Hume and Kant are – not bad. You should aim to do better.⁸

This advice will strike many as arrogant and foolish. It is neither. It would be arrogant to expect to reach or even approach the ideal. But to strive for it is *noble* and will in time vastly improve your writing. To appreciate that you will always fall desperately short and to accept the fact with good cheer is modesty enough.

Authority. Regard yourself as a master of the topic. Do not rely uncritically on the distinctions and definitions provided by other philosophers. Instead make all distinctions in the way you regard as most perspicuous, and define all terms in the way that you regard as most precise.⁹ Do not describe the motivations for various positions as they are described by influential figures. Instead describe them in the most apt way. Do not even rest content with rebutting the arguments of your opponents. State the best version of each argument, the version that your opponents *should* have given, and rebut that.¹⁰ Nor should you waste time in the body of the paper demonstrating that your approach is better than those in the literature. Just *make* your approach better, perhaps briefly discussing the shortcomings of extant approaches in footnotes. In short, draw the definitive map of the philosophical terrain.

So as to maintain focus on the ideas, *do not even cite any contemporary philosophers in the body of the paper.*¹¹ (That is, absent compelling reasons to do otherwise, as when you are writing a reply article.) Relegate all citations to the footnotes. Those footnotes should, however, be extensive; they should meticulously document the relationships between your philosophical map and the philosophical maps of others. See the discussion below of mastery of the literature.

While you are at it, prune needlessly modest language, including phrases like “in my view,” “I believe that,” and “it seems that.” Write with confidence.¹²

Systematicity. To be systematic is to minimize arbitrariness. Thus, when laying out possible views on some topic, do not just list the currently influential views in the literature. Instead develop an illuminating **taxonomy** – a categorization of possibilities that are mutually exclusive and jointly exhaustive given your assumptions, with all divisions between these possibilities marking profound differences – and situate all currently influential views within it.¹³ 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, but all n together

⁸ 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*.

⁹ Lormand (1996, introduction and §1) and (ms, introduction and §1); Schellenberg (2010); Berker (2013, pp. 337-338 and 344-348); Millar (2015).

¹⁰ Lewis (1984).

¹¹ Why the distinction between contemporary and historical philosophers? Not because you should defer to Plato or Kant – see the discussion of ambition – but because referring to them helps locate your view against others at the highest level of generality and has all the literary value of allusion to boot. See Street (2010); Millar (2015).

¹² Williamson (2000); Chalmers and Jackson (2001); Street (2010, pp. 369-370).

¹³ 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).

are inconsistent.¹⁴ 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.¹⁵ 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.¹⁶ Even when just transitioning between paragraphs or sections, make it implicitly clear why the new topic is next on the agenda.

Rigor. Support your thesis with *conclusive* evidence. You may use a single decisive argument¹⁷ or an array of considerations that are together overwhelming.¹⁸

Limit yourself to dialectically effective evidence – evidence that even your opponent should accept.¹⁹ Therefore, avoid relying on all but the most robust intuitions.²⁰ Focus instead on theoretical considerations (simplicity, explanatory power, fruitfulness, etc.), scientific data, linguistic facts, and the most universally acknowledged truths,²¹ or just show that your theory is strictly better than your opponent’s theory.²²

Do not rest content with presenting a *prima facie* objection to a view. Instead, press your point as far as possible: consider all possible categories of response (see the above discussion of systematicity) 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 a weaker but still interesting version of your conclusion.²³ And if there are multiple independent or partly independent routes to your conclusion, say so.

Suppose, as will often be the case, that you wish to rely on claims that you cannot rigorously support. *Then jettison any arguments for those claims and introduce them as assumptions.*²⁴ Do not hesitate to make controversial assumptions as long as you identify them as such. State all assumptions in your introduction, however, as your reader will feel cheated if you help yourself to controversial claims once the argument is underway. You may motivate your assumptions if you wish, but be brief lest it seem that you are trying to argue for them.

Significance. Emphasize what matters. Minimize everything else.

Start by articulating, in the most general terms, the central philosophical problem(s) that you will address and the central solution(s) that you will propose. Highlight these. Use the positions of famous historical philosophers as landmarks to orient your reader, give telling metaphors, be generous with examples, and even restate the same point in equivalent ways if this is helpful. Heed this advice especially in beginnings and endings, whether of your entire article or of individual sections. Likewise, though on a smaller scale, when you get into a welter of details. Just keep pausing to explain, as broadly as possible, what you have done and why it is

¹⁴ Egan and John (ms, p. 1).

¹⁵ Millar (2015).

¹⁶ Gendler (2002, §3); Enoch (2006, §5). Turri (2010) carries out a positive analogue of this: he articulates with great precision how a certain class of views arguably avoids the Gettier problem.

¹⁷ Williamson (2000, ch. 4); Bailey (2010).

¹⁸ Williamson (2000, ch. 11); Shah (2003); Kolodny (2005); Schaffer (2009, pp. 366-373) and (2016); Evans (2013); Moss (2013).

¹⁹ Balog (1999); Levine (2010); Greco (2012, §2, especially p. 350).

²⁰ Fine (1994); Johnston (2004); Schwitzgebel (2008).

²¹ Chalmers and Jackson (2001); Moss (2013); Sinhababu (2015).

²² Lewis (1979); Sinhababu (2015).

²³ Moss (2013, §5.5).

²⁴ Lormand (1996, pp. 52-53 and 61).

important. Your efforts here will be reinforced by attention to ambition, authority, and systematicity. Conclude by recasting your solution in a way that the reader is only now, having read the entire paper, in a position to fully appreciate.²⁵

Inversely, prune whatever is not in service of what matters. Introductory remarks, assumptions, definitions, resolutions of local debates, discussions of fine details, formal/technical language – do not hesitate to include any of these if your argument requires them. Ruthlessly excise them otherwise. This is, in effect, a matter of economy (see below).

If you cannot find anything of significance in your idea for an article, scrap it and write something else.

Economy. I am told that rocket scientists supply a spacecraft with the exact amount of fuel it requires to reach its destination. Perhaps this is a myth. But you should construct your arguments in the same spirit of economy. In particular, adjust your premises and conclusion until they precisely match.

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 some very specific view of a prominent philosopher. It will target a substantial class of views, and it is your job to delineate that class. Also consider whether the argument can target a much larger class of views with only slightly stronger premises.²⁶

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.²⁷

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

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.²⁸ Avoid imprecision in either form.²⁹

Metaphorical language, though often a great help in other respects, can also disguise imprecision, so be sure that you can rephrase your metaphorical claims in literal terms. Formal tools can in contrast help you be very precise, but be careful not to use them to make claims that are needlessly specific. Though it is not ideal, it is acceptable to formulate a claim roughly as long as you explicitly state that you are doing so and the imprecision is irrelevant for your argument.

Focus. An article ought to have very few central argumentative tasks – sometimes just one. Delineate your task(s) precisely: by the end of your introduction, articulate your conclusion in detail as well as all of your assumptions; also sketch your central argumentative moves. Then strip away all material not required to accomplish your argumentative task(s).³⁰

²⁵ Fine (1994); Lewis (1979); Lormand (1996) and (ms); Schaffer (2009) and (2016); Berker (2013); Evans (2013); Lord (forthcoming).

²⁶ Berker (2013) leverages a single core idea against a very broad class of theories.

²⁷ Neta (2002), Sinhababu (2009), and Lormand (ms, §4) do much work with sparse resources.

²⁸ Obviously, a claim may also be imprecise in both ways at once.

²⁹ Lormand (1996) and (ms); Williamson (2000, ch.4); Bailey (2010); Schaffer (2016).

³⁰ Bailey (2010); Turri (2011); Sinhababu (2015).

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.³¹

Focus and cohesiveness tend to compete, since one typically achieves the former by sparsity of posits and the latter by richness of posits. Which virtue to prefer can then be a hard call; see the earlier discussion of economy.

Mastery of the literature. The literature on a topic usually comprises a few seminal texts, whether historical and contemporary, and a flood of subsequent texts. Have a deep knowledge of all seminal texts, as well as of those subsequent texts that you regard as excellent. Have at least a working acquaintance with almost everything else, *especially those texts published in the past 5-10 years, as well as forthcoming texts*. Keep up with work defending *every* major position on the topic, even those towards which you are least sympathetic.

At an intellectual level, display your mastery of the literature by using your knowledge to draw a new and illuminating map of the terrain. Characterize the problems and define all terms in the way you regard as most perspicuous, and locate your position within a taxonomy, of your own devising, of all possible positions. In other words, use your knowledge of the literature to be authoritative and systematic, as discussed above.³²

At a more practical level, display your mastery of the literature through excellence in citations. Include an early footnote citing all parties to the central debate of the paper. This footnote should first cite the seminal literature, flagging it as such, and should then cite the rest of the relevant literature, relating it as appropriate to the seminal literature.³³ Take special care to cite every author who might reasonably be asked to referee your paper. I repeat, however: absent some *very* special reason, *do not cite contemporary philosophers in the body of your article*. Relegate all such citations to the footnotes.

Clarity. At every point in your article, your reader should know what has happened, what will happen, and what is currently happening. Accordingly, the ideal article is liberally marked with signposts. Some helpful signposting devices include: providing an overview of the article in the introduction; clearly stating your aim and your assumptions; using informative names, rather than numbers or unmemorable abbreviations, for central theses; using informative titles for sections; periodically recapitulating central results; and, most importantly, making generous use of transition words like *thus, since, because, however, but, therefore, nevertheless, despite, further, and consequently*.³⁴

Beware of sounding formulaic, however. For example, do not begin every section by stating, “In this section, I will show that” Excellent signposting is not so heavy-handed and involves artful variation.

Concreteness. It is not in our nature to think solely via abstract principles. We find it hard to grasp the principle 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. We find it easy to grasp the claim that given all of our evidence, we might be the epistemic playthings of a deceiving demon. So make frequent use

³¹ Williamson (2000); Neta (2002); Skorupski (2010); Graham (2012); Moss (2013); Schaffer (2016); Lord (forthcoming).

³² Lormand (1996) and (ms); Neta (2002, p. 664); Turri (2010).

³³ Schellenberg (2014).

³⁴ Lewis (1979); Schaffer (2009), (2012), and (2016).

of concrete examples to supplement, though of course not to replace, abstract thought. Well-chosen examples will clarify principles, make definitions easily digestible, and illustrate structural points.³⁵

Flair. Be alert to opportunities for the wry remark, the keen observation, the striking metaphor.³⁶ And make each example memorable, perhaps with humor, allusion, or picturesque description. To rely on stock examples is a waste.

Do not, however, use flair at the expense of clarity or precision. Use metaphorical language as a supplement to literal language, not as a substitute for it. And be sparing. As with cooking spice, flair is pleasing in moderation and noxious in excess.

Originality. Originality, though it is of great importance for excellent writing, requires no separate attention. Just focus on the virtues already discussed and let originality emerge naturally.

4. Writing: process

The production of a body of excellent writing, like the building of a great city, requires *time*, and vast quantities of it. It is not enough to put in the occasional marathon session. Good ideas must simmer in the subconscious over many nights, and sometimes many months, before they mature to excellence. I therefore recommend that you dedicate a regular block of time to writing.

I reserve at least one hour for writing per weekday during the teaching term and two hours per weekday otherwise. Though I may or may not write for longer, as time permits, I treat that reserved time as sacred. I do not let it be squeezed out by other professional obligations – not by teaching, not by grading, not even by reading and thinking. I make an effort to preserve my writing time even when attending conferences, though I let it go during periods of illness, vacation, or emergency. To ensure that other commitments lapse before this one, I schedule my writing for my *very first* block of open time each day, usually from 10–11 AM.

Mistakes are a necessary and even fruitful part of my writing process, and keeping to this regimen gives me plenty of time to err. I have never expressed any significant idea correctly on my first try, and usually not on my second or third, either. I find that writing is like picking my way through a labyrinth: I *always* make many wrong turns, and even the right turn often leads me away from the exit. So I reward myself not for the number of polished pages that I write but for the amount of time that I spend.

Because my mistakes are often radical, my willingness to revise is radical, too. Without exception, all of my best papers have been completely overhauled at least once and sometimes twice. I will mention a recent case about which I have kept exact records. I spent 80 hours writing the first draft of the paper; after receiving an invitation to revise and resubmit, I spent an additional 100 hours scrapping and then replacing the basic framework of the paper; and upon receiving a rejection letter with further comments, I spent 80 hours more revising crucial sections before sending out the paper once more. It is only because I treat my daily writing time as sacred that I can afford the luxury of proper revisions.

That is my way through the labyrinth; you will have to find your own.

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

³⁶ Schaffer (2009, p. 355), (2016, pp. 92–93), and (forthcoming, penultimate paragraph of §3.1); Schwitzgebel (2012, pp. 39–40); Williamson (1995, first sentence).

5. Writing: publication strategies

A professional philosopher's body of work is standardly judged along two dimensions, quantity and quality, which correspond to two extreme strategies for publishing. To take the **extreme quantity strategy** is to submit an article as soon as you believe it to meet the minimum quality standards of a solid journal, the idea being to produce an enormous quantity of solid work. To take the **extreme quality strategy** is to submit an article only once you believe it to be of the highest quality that you can achieve.

Few go to such extremes. But there are very successful professionals near each of them, and you will find it helpful to consider roughly what point between these poles you wish to occupy. I prefer to be near the extreme quality pole, and my instructions in this guide have been formulated with that end in clear sight; those with different strategies will need to modify much of my advice. I am not, however, endorsing any judgments about the *philosophical value* of quality relative to quantity. I simply enjoy myself most when I attempt to write articles of the highest quality that I can achieve, so that is what I try to do and what I have the experience to discuss.

There may seem to be a third extreme strategy, the strategy of producing work that is extreme in quantity and quality alike. In fact, this strategy is a version of the extreme quality strategy – philosophers who follow it just spend an astonishing amount of time on writing and therefore produce in great quantities. Such philosophers usually rank among the most respected in the profession, but do not underestimate their labor.

You may as well take measures to increase quantity with no sacrifice in quality. For example, it requires a great deal 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. In this way you will submit articles for review much more rapidly.

Speaking of which: know that it is a long road between putting finger to keyboard and having your article accepted. Though I lack precise general data, I can at least share my own statistics. (I welcome information from others willing to share.) Across my 8 publications, between my first submitting an article and its being accepted there has been a median gap of 15 months and a mean gap of 16.5 months. I estimate also that the gap between the moment that I begin a new article and the moment that I first submit it is an additional 4–8 months. In total, then, it typically takes about 19–23 months from when I start an article to when it is accepted. That's a long time!

Hence, you must take joy in the process of writing itself rather than in the distant prospect of publication, and for the sake of your career you must also start writing *well in advance* of milestones. You should have at least one publication, preferably more, 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 work on your first article when you start your dissertation, if not earlier. (These two activities should be largely co-extensive.) Submit that article for review within the year. Also, remember that you may need to submit your tenure file at the start of an academic year, and that by this point you must have a substantial body of work. So you have no 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*.

If you have an abundance of talent, luck, and connections, then your articles will never be rejected, and you may skip to the last two paragraphs of this section. I, on the other hand, have amassed rejections by the truckload, and I have two pieces of advice on how to handle them. The first is psychological: maintain your confidence! Even if you receive one rejection after another, with no end in sight, treat these as signs of inadequate proficiency and not as signs of inadequate intellect. Easier said than done, I know. I have mentioned that my career began with a streak of 18 consecutive rejections, and by the end of it my confidence was below sea-level. Now that I have made it to the other side, however, I can assure you that the main requirement for success in publishing is just *toil* – reflective, disciplined toil. Knowing how to write a publishable article is a skill. It is a skill that you will learn if you put in consistent effort and reflection.

My second piece of advice is to take the content of referee comments very seriously, while ignoring the tone. Referees are often sarcastic, cruel, and dismissive, but at the same time their comments can almost always help you revise. In fact, though long, critical referee reports are invariably painful to receive, these have been my *single most helpful* source of comments. I have of course received terrific comments from advisors, mentors, friends, correspondents, and audiences at talks. But good referees are experts on the topic and have scrutinized the paper with unusual care. Many referees have given me 4 pages or more of thoughtful and sharply critical feedback. Only from a handful of others have I gotten anything comparable. Thus, when I receive negative comments from a referee, I do not attempt to revise immediately. I always feel upset when I first read such comments, and I must set them aside for a day or two so that I can return to them with a level head. Then I consider the comments one at a time, taking each seriously.

Taking a comment seriously is compatible with making no revisions in response, but it is only rarely that I find myself so unmoved. Even when I think that a referee has simply missed the point – and I think this often – I almost always find that I can clarify my idea so that it will not be missed again. On occasion, however, a comment is just misguided and is best ignored. Let your calm, considered judgment be your guide.

Indeed, before I submit an article for review, I find it helpful to examine my work from the perspective of two kinds of referees. I first ask what my article does *best* according to my *most sympathetic* referee. Taking this perspective tells me what to showcase and what to strip away. I then ask what my article does *worst* according to my *least sympathetic* referee. Taking this perspective reminds me to address important objections, add key qualifications, head off potential misunderstandings, and most of all define my dialectical task precisely. To put things another way, when a work has certain virtues, an attentive referee will want to accept it, and when a work has certain vices, an attentive referee will want to reject it, so it is wise to write both so as to *invite acceptance* and so as to *avoid rejection*.

By the way, please never let yourself become a sarcastic, cruel, or dismissive referee. A referee is right to have high standards, but it is one thing to state precisely and in detail why an article is not publishable and another to make its author an object of scorn. Indeed, I believe that it is the responsibility of a referee also to note what an article does well, *even when that article falls far short of being publishable*. Such observations, besides being a kindness to the author, help her identify what ideas to preserve, highlight, or develop when she revises.

6. 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 came the realization that I lacked crucial writing skills. I resolved to learn all that I could. I began by requesting all the feedback that I could get on my own work.

But I met an obstacle: I was no good at putting the information I gathered to any general use. I would learn that I needed to include certain citations on p. 29 or that I should assume rather than arguing for the transparency of experience, and I understood how to apply these pieces of advice in the immediate context. But because I did not grasp the general principles underlying the particular advice, I continued to make the mistakes of just the same type.

It was here that my lack of talent was made manifest, for talent – I conjecture – consists largely in the ability to pick up principles of expertise 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 talented person moves smoothly from seeing examples of expertise to internalizing the underlying principles without ever needing to (and usually without being able to) articulate them. The rest of us need to be told to lengthen our strides, square our hips, lift our chins.

Lacking the benefits of talent, I had to practice reformulating any advice that I received into *specific* but *general* principles. (Note that I do not use the italicized terms as contraries; the specific is to be contrasted with the vague, the general with the particular.) When told that I needed to include a certain citation, I attempted to articulate the specific general condition that warranted that citation. When told that I should assume rather than defending the transparency of experience, I attempted to articulate the specific general condition under which I should assume rather than defending a claim. And so on.

Still with the aim of articulating specific yet general principles, I also began to seek exemplars of excellent philosophical writing. I followed the work of several contemporary philosophers who were consistently publishing in top-tier journals. So that I might learn from a wealth of data, I would read a series of papers by one philosopher before moving to the next. Once I gained some confidence in my ability to discern excellence, I expanded my search to include any outstanding work, whether in the form of a contemporary article or a historical book. All the while I asked successful philosophers to tell me, with as much precision as possible, what works they regarded as best and why, what they saw as the best features of their own work, and, ascending a level, what their processes of writing, reading, and thinking were like.

I recorded the principles that I posited in a Word document intended for personal use. That document 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 certain goods, rules, and virtues distinctive to philosophy. My most basic advice on self-cultivation, then, is to carry out your own reflections, which will mutually inform one another, across all three levels.

A practical method for doing so is to *design your own guide to philosophical writing*. (If you have internalized my previous advice about ambition and authority, then you must surely doubt much of what you find here!) Start small: whenever you read a great philosophical article or book, whether it is contemporary or historical, identify and write down its virtues in specific but

general terms. And there's no need to go it alone; learn from others. When I meet philosophers at conferences, I like to ask which contemporary and historical works they most admire and why, and especially which works they most admire while adamantly disagreeing with. Many times such conversations have helped me recognize virtues that I had previously appreciated only implicitly, if at all. Even if you reserve your guide for personal use, just articulating your understanding of philosophical excellence as precisely as you can will give you much more control over your writing.

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 only a part. A reader of this guide might be left with the impression that philosophy is my consuming passion, one that leaves no time to spare. That is not true. I rarely put in more than 50 hours per week of research, teaching, and service, and I make it my personal ideal to have satisfying personal relationships, keep up one or two serious hobbies, eat well, exercise regularly, meditate daily, and sleep plenty.

Do I reach this ideal? 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.
- Brooks, T. (ms). "Publishing advice for graduate students."
- Chalmers, D. (2011). "Verbal disputes." *Philosophical Review* 120 (4): 515-566.
- 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.
- 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.
- Lin, H. (ms). https://docs.google.com/document/d/128nK28AQIH_XJEzLQmKtzCGmOHRjE0ksWxk3Fui98lU/mobilebasic?pli=1#ftnt_ref1
- Lord (forthcoming). "What you're rationally required to do and what you ought to do (are the same thing!)." *Mind*.
- 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. (2015). "Knowledge and other norms for assertion, action, and belief." *Philosophy and Phenomenological Research* 91 (1).
- Millar, B. (2015). "Frege's puzzle for perception." *Philosophy and Phenomenological Research* 92 (2): early view.
- Moss, S. (2013). "Epistemology formalized." *Philosophical Review* 122 (1): 1-43.
- Neta, R. (2002). "S knows that p." *Noûs* 36 (4): 663-681.
- 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. (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. (forthcoming). "Social construction as grounding; or: Fundamentality for feminists, a reply to Barnes and Mikkola." *Philosophical Studies*.
- 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.