Mind over Manuscript: Eight Strategies for Writing Philosophy

Julia Staffel

To appear in a Festschrift for Alan Hájek's 60th birthday

"Philosophy is hard. I feel that I can use all the help that I can get, and I suspect that you can too."

—Alan Hájek

Introduction

Writing philosophy well is an essential skill in our discipline. Philosophical writing must aim for clarity, precision, and rigor, but in doing so, it can often wind up dry, long-winded and boring. It can take many drafts to produce a paper that is suitable for publication in a journal, and many aspiring (and accomplished!) academic philosophers find the process of writing arduous and frustrating. Still, some people make it look easy – if you've read anything by Alan Hájek, you've probably noticed his breezy style that effortlessly communicates complex ideas in simple terms. His concise and witty prose makes even formal epistemology, a notoriously complicated, math-heavy subject, accessible and engaging to readers.

However, while it might *look* effortless, Hájek's spirited style is in fact born out of deep and thoughtful engagement with the craft of writing. His motto is: "Work hard for your readers, so that they don't have to." Hájek's approach is decidedly anti-genius: he believes that having good ideas and communicating them well can be taught, and he has devoted considerable energy to helping his graduate students improve their writing. He has written multiple articles about philosophical creativity, as well as an unpublished lengthy manuscript on the mechanics of writing. I have benefited myself from his advice – at least I believe I have, readers may judge for themselves. My aim in this article is to share a few of his insights that I have found most helpful for myself and for my students. I won't try to summarize all of Hájek's advice, and I also don't claim that all of this is totally new. Rather, I will offer a small collection of greatest hits. I will cover eight aspects of philosophical writing, and for each one, I will explain the basic idea, and then discuss some ways of implementing it for oneself and one's students.

1. Searching for Ideas

1.1 Basic Idea

Philosophy students often think that having interesting, novel ideas that are worth publishing is a matter of luck. You sit around reading other people's work, and suddenly, poof, a publishable idea pops into your head. If you're lucky. If not, you're doomed. But, if not like this, how do philosophers actually come up with ideas? While there are methods for generating ideas, philosophers are not always consciously aware of relying on them, let alone able to describe them and deploy them methodically.

Enter Hájek's articles on philosophical creativity – as far as I know, they present the first explicit collection of strategies for generating philosophical ideas that can be learned, taught and

explicitly deployed.[1] I recommend reading his articles in full, but here are some especially useful heuristics (Hájek 2014, 2016):

Check extreme cases and near extreme cases
To see if a theory really covers all the cases it is supposed to cover, think about cases that are at or near the boundaries of the intended domain. Counterexamples and problem cases can often be found here. Applying a theory to itself to see if it still holds can also be very useful.[2]

Continuity reasoning
Continuity reasoning refers to the strategy of checking whether a phenomenon is graded/continuous or includes natural steps or jumps, and whether the theory correctly accounts for this. Further, we can ask whether a phenomenon is similar across different domains, and whether the theory appropriately generalizes to them.

Transforming arguments from one domain to another
It can often be useful for finding a new idea, or finding problems with existing views, to ask whether a similar phenomenon has been philosophically theorized in another domain. For example, if you are thinking about theories of epistemic value, it will likely be useful to consult theories of value in ethics and think about how their insights and problems transfer over to epistemology.

Trial and error (brute forcing)
Sometimes when the right kind of view or counterexample is not forthcoming, it can be useful to systematically search the whole space of options. One benefit of this strategy is that it forces one to systematically lay out the ground one intends to cover, so it can be searched exhaustively.

In his articles, Hájek offers additional heuristics and illustrates each with many helpful examples, which I highly recommend to readers. I now turn to how I have implemented the use of heuristics in my own writing and with my students.

1.2 Implementations

To make students think about explicit strategies for finding ideas, I have incorporated several lessons and requirements into my classes. One of them is to divide students into groups, and ask

[1] Philosophical methodology has been a topic of discussion for longer, but Hájek's innovation was to shift the focus away from the question of what justifies certain arguments or ways of doing philosophy to the question of how a person can strategically employ heuristics for finding ideas and developing philosophical theories. More recent work roughly in this vein includes Kelley (2024) and Kallestrup (2025). See also Williamson (2024).
[2] For example, in epistemology, it is often emphasized that agents should aim for true beliefs and try to avoid false ones, and this is implemented in various theories of justification. Some of these theories have trouble dealing with tautologies and necessary truths, however, generating the implausible result that beliefs in those are always justified. For more examples of all these heuristics, readers may consult Hájek's paper (Hájek 2014).

them to brainstorm about two questions: 1) What constitutes a novel philosophical contribution? and 2) What are good strategies for coming up with philosophical ideas?

Having students reflect on the first question is important, because they often have an overly narrow conception of what counts as a novel contribution. Most students realize that giving an objection to an existing theory and inventing a new theory would count, but they often aren't aware that they could also, for instance, give a new argument for an existing theory, formulate a new question, generalize an existing theory, or defend a theory from existing counterarguments, among other things. Reflecting on the second idea is important because students often believe that having ideas is a passive process, in which they first learn the material they want to write about, and then wait for inspiration to strike. It's important to change this narrative, and show them that generating ideas is instead an active process that they can learn to get better at. (See the appendix for the actual lists my students generated in response to the two questions from above.)

I also have the students read Hájek's papers about philosophical creativity, but only after this brainstorming exercise. The order is important, because I've found that once students have a set of ideas presented to them, it shapes how they think about the issue and drowns out their own ideas. This happens especially when the material presented to them is perceived as correct and authoritative. Further, it obscures how much the students already know about an issue, by never giving them a chance to articulate their own answers. If, instead, students engage in the brainstorming activity first, they can articulate a broad range of suggestions on their own, and then later use the literature to fill in gaps. Further, they get validation for their ideas when they find them reflected in the literature.

Later in the semester, when students need to come up with ideas for their papers, I ask them to put these strategies into action. They must come up with multiple ideas for a final paper and write a mini-abstract about each. This exercise is intended to combat another unfortunate habit, which is that students often glom onto their first paper idea and stop looking for other ones, even if their first idea turns out to be not very good upon further reflection. I try to instead normalize a work process in which one initially generates multiple ideas, which can then be compared and examined before one chooses the most promising one to develop further.

There are also two further heuristics that I highlight to my students, which help with finding ideas and examining one's own views. I believe I picked these up in graduate school, likely from Jake Ross or Mark Schroeder. The first one is to scrutinize the footnotes and endnotes in the texts one is reading. That's where people tend to hide problems, unanswered questions and other dirty secrets, which are ripe for further investigation. For example, in discussing the appropriate use of footnotes, Hájek mentions that Popper tended to relegate "devastating counterexamples to his own views to footnotes." Hájek argues that good writing should avoid putting important content into footnotes, but since many philosophers don't heed this advice, they are a treasure trove for further exploration (Hájek 2022, p. 199).

The second heuristic is for coming up with counterexamples to definitions and theories, which many students find challenging. A helpful intermediate step is to analyze the relevant view to yield an abstract description of the shape a potential counterexample needs to take. For example, suppose a definition has the form of a biconditional with multiple subclauses, such as "For all A,

an A is X if and only if A has properties G, F, and H." (This is, for instance, the shape of the notorious justified true belief account of knowledge.) Coming up with a good counterexample is easier if we first list all the possible shapes such an example could take, i.e., something that is X, but lacks one or more of the listed properties G, F, and H, or something that has all these properties without being X. Thinking about whether it is possible to devise a concrete case that matches one of these abstract descriptions is much easier than trying to intuit a counterexample without first characterizing it schematically.

Ideally, students will use these methods to come up with multiple ideas for their writing. But how to choose among them? In my experience, philosophy students (and not just students!) are often terrible judges of which ideas are worth pursuing. One reason is, of course, that they don't know the literature well, which is necessary for judging how novel an idea is. For this, they need guidance from more knowledgeable mentors. But independently of this, many philosophers are either overly excited by their own ideas, or don't recognize how interesting they are. After a while, one's own ideas can become overly familiar and thus seem boring. A good remedy is to incorporate opportunities for feedback into the writing process at an early stage, ideally when one is choosing which ideas to pursue. Here, we can ask others not to just comment narrowly on whether an idea seems novel, or has obvious problems, but also to judge how exciting it is. Thinking together with others about what's especially cool about an idea, articulating worries about why it might turn out boring, and thinking about ways of spicing it up and broadening its applications can be immensely beneficial to the writing process.

2. Naming Your Ideas

2.1 Basic Idea

How enjoyable it is to read a paper, and how memorable the ideas in it are, can depend a lot on the naming and labeling choices made by the author. While it's not bad to use the occasional abbreviation, some papers end up with letter salad that is extremely hard to keep track of, along the lines of "A counterexample to the XYZ principle, but not to the XYZ* or ABC principles can be generated by appealing to the PQR theorem." Remember Hájek's advice to work hard, so your readers don't have to. If you need to label different theorems or principles, try to make the names memorable and descriptive of what they stand for. For example, if you have a principle in your paper called the "economy principle", don't call it "EP". You could call it "ECON", which is also only one word, and it is much easier to remember what it stands for.

Hájek also reminds us that "technical terms and abbreviations should earn their keep" (Hájek 2022, p. 84). Why invent abbreviations and labels for things that never get mentioned again? While it makes sense to abbreviate things that get talked about repeatedly, if something gets only mentioned once or twice, it doesn't need a fancy label. The same holds for unnecessary variables. It is rather silly to introduce variables into one's language just for the sake of it. Why say "a person P can purchase a coffee C by paying an amount of money M to the vendor V", when one could simply go with "a person can purchase a coffee by paying money to the vendor"?

2.2 Implementation

I have a few more things to add to Hájek's excellent tips. When choosing names and abbreviations, it's good to try out what they will sound like when someone talks about them or reads them out loud. When you present your work in a talk or when people discuss it, it's important to make sure that your labels are easy enough to pronounce and distinguish. For example, the abbreviation "XYZ**" is easy to read silently, but saying "XYZ double asterisk" or "XYZ double star" repeatedly is clunky and annoying. Or, if two principles are called "VP" and "BP", this is very easy to distinguish on the page, but spoken out loud it can be hard to tell which is which, especially if you have speakers and listeners with different accents.

The same holds for choosing examples. Examples that are as simple as possible, easy to remember, and quick to explain make for both good writing and good discussion. A lot of important philosophical debates are centered around a key case, such as the trolley problem, Mary in the black and white room, or the sleeping beauty problem. Of course, most of our examples probably won't wind up being the centerpiece of an extended debate, but why not choose ones that would be up to the task? For a very famous case that doesn't follow this advice, take Gettier's counterexamples to the justified true belief account of knowledge (Gettier 1963). His work has been discussed for decades, but many people don't use Gettier's own examples, for good reasons. The examples of Smith and Jones with the coins in their pockets and random locations are needlessly complicated and hard to remember. However, avoid the temptation to make an example memorable by making it gratuitously graphic or violent. I like to use Forbes' *Modern Logic* in my classes, but during a semester that included an active shooter lockdown, I really didn't need the predicate "shoot" to be the main example in the chapter on two-place predicates (Forbes 1994).

When naming papers and ideas, aim for being inventive but descriptive. Unless you are very famous and people will read your papers no matter what they are called, a non-descriptive name risks missing out on readers. For example, a very important linguistics paper is Philippe Schlenker's "A plea for monsters" (Schlenker 2002). Cute title, but do you have any idea what it's about? Turns out, in the debate about indexicals, a "monster" is an indexical term whose context of evaluation can be shifted. Schlenker's paper has been very influential, but choosing this title was risky, because only a small in-crowd would know that it was a paper about a specific kind of indexical expression with unusual properties. It's better to try to find fun titles that are also descriptive of the content. Hájek himself is a master of this. Here are some of his greatest paper title hits: "Ramsey + Moore = God" (a paper about conditionals and rational belief), "Declarations of Independence" (about formulating principles of probabilistic independence), "Scotching Dutch Books" (on how to improve Dutch book arguments), "Risky Business" (on the rational permissibility of risk attitudes), "Triviality Pursuit" (on formulating triviality results for the thesis that probabilities of conditionals are equivalent to conditional probabilities), etc. He has many more good ones, but you get the idea.

How does one come up with good titles? I am strongly in favor of asking others for help here. Some people just have a gift for naming things, and most of them would be delighted to name a paper or a view. My friend Justin Snedegar has this talent, so when I am looking to name a view, I often ask him. He helped me come up with "transitional attitudes" (a type of preliminary attitudes we have during ongoing deliberations) and "pro tem rationality" (a type of weak rationality

standard for the transitional attitudes we form during ongoing deliberations). Crowdsourcing is a good method too. My student Amir Ajalloeian is writing a dissertation about Lindley's notion of probability through a transform, which is now aptly titled "Transforming Bayesian Epistemology" thanks to the joint efforts of my PhD student workgroup. It can also be helpful to consult AI. When I was looking for a title for this paper, I asked ChatGPT for suggestions, and while most of them didn't impress me, they helped me brainstorm in different directions until I landed on "Mind over Manuscript." I liked the alliteration, and the allusion to the principle "mind over matter", which suggests that we can solve a (physical) problem by using our mind. Adding the subtitle then made the paper title sufficiently descriptive to (hopefully) find its intended readership.

In sum, the matter of naming one's views, principles and papers should not be treated as an afterthought but given proper consideration, since it influences how easy it is to read and discuss the work, and also whether it reaches its intended audience.

3. Taking Notes

3.1 Basic Idea

One thing you'll notice about Hájek when you meet with him and talk philosophy is that he is always taking notes. He carries a tiny notebook in his pocket, and any time something comes up in the conversation that seems possibly important, he writes it down. You might think, everyone takes notes, big deal. But in my experience, people often don't take notes when doing so would be most beneficial, and when it could save them a lot of time and effort down the line. Hence, here are a few notetaking strategies that I use myself and that I try to instill in my students.

3.2 Implementations

Taking notes in meetings: I am often baffled that students come to me for advice about a paper, we have an extended, complicated discussion about it, and they don't write down a single thing (unless I remind them to). I find it hard to believe that they are really taking advantage of my advice if they don't have a way of recording our conversation. Hence, I always try to encourage one of the following: Take notes by hand, old school. I often pause to make sure the student has time to write things down before we move on to a new point. Another good strategy is to take pictures: If we're writing on the board during a meeting to solve a formal problem, come up with a paper structure, or illustrate something with a diagram, I encourage students to take pictures on their phone as we go. I often do this too, so I have a record later of what the student and I discussed. Lastly, given that we often meet on Zoom and similar platforms now, I allow students to record or auto-transcribe meetings, eliminating the need to write notes while we're talking.

We often present our work at conferences to get feedback before we publish it, but it can be difficult to take notes during the Q&A or right after when we want to be discussing with people. A great trick for not missing anything is to establish a buddy system, where someone you know in the audience volunteers to write down the questions that people ask, ideally even with their names if possible. A few times, I have given talks where the chair volunteered to do this for me, which was fabulously kind and helpful. This makes it easy to come back to all the questions later and think

more carefully about how to answer them. If you have people's names, it also makes it easy to later credit someone for a good point or follow up with them.

Keeping track of all the questions people have asked about one's research comes in especially handy in preparing for job interviews. I encourage my PhD students to keep a running list of the questions people have asked about their research over time, because there is a high chance that they will be asked very similar questions during job interviews. Being able to think about how to answer in advance is a huge advantage. It's very hard to perform well in interviews without doing this.

Lastly, it's important to have an organization system that works. There's no benefit in taking lots of notes if one can't find them down the road when they are needed. I'd recommend against a system that is overly elaborate, because that makes it less likely that one sticks with it. The best system, in my experience, is one where I have a central gathering point for information that I can easily access from my computer and my phone, so I can quickly dump my notes where I can find them later. I use a cloud storage service (dropbox) that I can access from my phone and my computer, and each project has its own folder. I can add notes by writing them down on my computer and putting them in the folder, add written notes from my phone, and I can also quickly take pictures of handwritten notes or of a whiteboard and save the picture in the folder. Thus, all the notes end up in one place, and the cloud storage system means that my notes won't be lost if I lose access to my computer or phone. There are of course many different ways of organizing notes, but good ones meet the following criteria: They are easy to access from different devices, backed up so losing a device is not a problem, and there is a clear way of associating information with the project it belongs to.

4. Getting Yourself to Write

4.1 Basic Idea

Another major hurdle in writing philosophy is to actually do the writing. It is comforting to see that even a prolific author like Hájek sometimes struggles with this. In his manuscript on writing heuristics (Hájek 2022), he shares his favorite strategies for getting it done. He endorses some popular methods, such as having accountability buddies and using timed blocks for writing, but I think his most important method is to break writing tasks down into manageable chunks that don't seem daunting. One way of implementing it is his "night-watchman strategy." The idea is to deliberately start a new project at the end of the workday with only a little time left. Hájek says

> This makes the exercise less daunting. I think: 'I only have to do this for a short time – how bad can that be?!' In the worst-case scenario, I'm uninspired, nothing comes of it, but I have wasted little time. (…) In the best-case scenario, I'm inspired, something good comes of it, and I have some momentum going for the next day. (p. 237)

Another way to break up the writing process into smaller units is by creating detailed outlines of papers that can then be filled in. By thinking of a project in terms of modules that have sub-modules, a large task becomes a smaller task of filling in the parts. On a good day, one might

tackle a hard module, such as the defense of a central premise, or a compelling way of framing a problem. On a bad day, one can still make progress while working on an easier module, such as a summary of existing literature or hunting down works to cite.

4.2 Implementations

I have implemented these methods both in my own writing and with my students. For my own writing, the modular approach has been especially helpful in writing books. For my first book, *Unsettled Thoughts*, I was fortunate to get a yearlong fellowship without teaching. I knew I needed to have a full draft of the book written by the end of that year in order to be on track for the book to be published before my tenure review. I thus wrote a modular outline of the book (or at least, a preliminary one that I could follow), and calculated how much time I had for each chapter. I determined that I had six weeks per chapter, so I divided each six week period into a research period of about two weeks where I would read additional literature and develop my arguments, and a four week period where I would write one or more preliminary drafts of this chapter. My goal was to have a draft of each chapter at the end of each period, and I would move on to the next chapter no matter how rough my drafted chapter was at that point. This method turned out to be immensely efficient. I had a full draft of the book by the end of the fellowship period, with a pretty clear idea what revisions I wanted to make. Many of the revisions I needed to make in earlier chapters only became apparent once I had moved on to later stages of the book. Moving on rather quickly proved to be very efficient, because it would have been pointless to spend more time earlier on polishing chapters that I was going to change later anyways. But the biggest benefit of this strategy, in my view, was psychological. Rather than being faced with the daunting task of writing a book in a year, I could focus my attention on much more manageable goals that I could accomplish in a few days or weeks, knowing that if I kept it up, the larger project would come together automatically.

A simple but incredibly useful phrase I kept repeating to myself when things got difficult along the way was given to me by my former colleague Casey O'Callaghan: "Trust the process." It can feel very scary to discover that an initial idea has a fatal problem, or that a train of thought is a dead end. But of course, when we explore uncharted philosophical territory and develop new ideas, this is completely normal. We just have to keep going and we will eventually reach firmer ground again. We just have to trust the process and move forward.

However, sometimes even a modular strategy that breaks the work up into chunks seems too daunting. I was once invited to revise and resubmit a paper to a journal that I very much wanted to publish in, but the reviewer requests were so condescending, and their demands so hard to tackle, that even opening the document made me feel nauseous. Here, I implemented a method similar to Hájek's night watchman strategy. I decided that I could bear 15 minutes of dealing with this project in a day without despairing, and that I would allow myself to always stop after 15 minutes. It was still unpleasant, and it took a while, but I ultimately figured out how to modify the paper, and the revised version was accepted. Remembering that almost anything is tolerable for just a few minutes is sometimes the key to success.

I also encourage my students to use a modular approach in their writing. This has benefits for them and for me. For them, it has the usual benefit of making a large project easier to tackle. But it also has benefits for me as their supervisor. First, it helps me gauge how to best advise them. I usually check in with students and give them comments after they've done a "chunk" of their project, but for different students, the size of the chunks they can tackle between check-ins is very different. The modular approach helps me determine the right chunk size for each student. For example, a student who is prone to procrastination and going down rabbit holes might benefit from more frequent check-ins where I continuously ask for small portions of work, such as reading notes from their research, a list of bullet points of ideas, or a couple of 200-word abstracts. For other students, this might feel way too much like micromanaging, and they can handle completing an extended paper outline or drafting multiple sections in between meetings. Modularizing writing like this means that students don't spend a long time working on a text or idea that isn't panning out, and as a result, I don't have to spend a long time working through a piece of writing that is fundamentally wrong-headed. Problems can usually be identified and tackled at a much earlier stage with a modular approach. This also helps maintain motivation on both sides, as the modular approach keeps the workflow going and leads to noticeable progress.

The last thing I want to mention in this section is that we should not underestimate the amount of psychological resilience and pain management being an academic writer requires. Writing academic philosophy is akin to an Olympic sport. It requires putting in long hours and working on minute details over long stretches of time, there is intense competition at the top (think of the low acceptance rates in top journals!), and even putting in a lot of effort does not guarantee success. This means writers must endure boredom, mental exhaustion, the anxiety of waiting for decisions, the pain of rejection, and the stress of not knowing whether a project will pan out. Experiencing these emotions is normal, and it's helpful to recognize this as part of one's job. It's important to develop an inventory of coping strategies to help deal with these emotions, such as taking enough time off, exercising, cultivating other hobbies, commiserating with friends or talking to a therapist. And we need to remind ourselves that it's a choice whether we want to put up with these difficulties – are they sufficiently outweighed by the joys of writing to make them worth enduring? No one gets a special cookie for being a tortured writer, just like no one gets a special medal for being a miserable marathon runner. We should remember all of this when we give feedback to our students. It's good to acknowledge that hearing negative feedback is hard and to share some of our own strategies for learning from criticism and for managing our emotions. We don't need to pretend to our students that we're somehow immune to feeling bad about setbacks and rejections. We can be much better teachers if we show them how to cope.

5. Knowing Your Audience

5.1 Basic Idea

A key ingredient in good writing is to pitch it at the right level for one's audience. Recall Hájek's principle: "Work hard for your readers, so that they don't have to" (p. 7). Making your papers easy to read is good advice not only for altruistic reasons. Papers that are poorly written, or technical and difficult to understand tend to get less uptake. Hence, it is in the writer's own interest to be

kind to their readers. Hájek himself applies this strategy aggressively in his own writing. Although he frequently writes about very technical issues in formal epistemology and decision theory, he eliminates unnecessary formalisms whenever possible, and never uses a formula without immediately explaining in plain prose what it means.

In order to make our writing more approachable, he advises us to think about writing with a particular person in mind, rather than an abstract, general audience of philosophers. For example, one might imagine writing something in an email to a friend (who is perhaps not the world-leading expert in the relevant subject…), or in an overview for one's students.[3] This also helps with the tendency to overestimate one's audience. I find that, especially in formal epistemology, graduate students in particular think their professors are far more knowledgeable than they really are (or at least than I am…), and thus write incomprehensible papers with an audience of impossibly smart professors in mind. Yet, doing this will almost certainly scare away all but the most dedicated readers, either because people think they won't be able to understand the paper at all, or because they don't want to invest the time to work through it.

But how does one learn to better address one's actual audience? Hájek has several great strategies, which I will now describe.

5.2 Implementations

There are a few strategies that help making formal philosophy in particular more approachable for readers. Omitting unnecessary formalisms is very important. Writers who are immersed in a formal system often don't reflect on whether it would be better to avoid a formal expression, because they perceive the formalism as equally easy to understand as an English expression, and they might even prefer the formalism for being more concise. Things look very different from the reader's perspective, however. A reader who must familiarize themselves with a particular formalism for the first time will have to slow down a lot to understand formal expressions, and they might not even engage with a text if the perceived difficulty level is too high. Hence, employing a formal expression should be a conscious choice, and deliver a payoff that can't be achieved by using prose instead.

When using formalisms, the following strategies help the reader understand the material better and more quickly: First, authors should explain their notation and say what all of the variables and constants are supposed to stand for. They should avoid using non-standard notation if more widely used notation will work just as well. Sometimes, just a couple of unfamiliar symbols can be a major hurdle for readers. Second, authors should say what every formula means in English prose either right before or right after the formula appears in the text. Even if the English gloss is not as precise as the formal expression, it will be tremendously helpful to the reader. Third, authors should illustrate every general theorem with a simple example. Writers are often tempted to introduce a general result or theorem first, and then introduce the example afterwards. However,

[3] However, writing for a friend should not be taken as an invitation to presuppose too much shared background, whether it's philosophical or cultural. The philosophy community is highly international, and we need to keep in mind that cultural references that are easy to get for one's friends and colleagues might be completely obscure in ten years or to readers in other countries.

for the reader, this is not necessarily the most useful order. Understanding the general theorem might be very difficult, and so readers might get bogged down before even getting to the example. Explaining the relevant phenomenon with a simple example first, and then showing how it generalizes (perhaps even relegating general theorems to the appendix) is far more reader friendly. It gives readers a chance to understand the basic idea before encountering it in full generality. This is especially true for giving presentations – the audience will get far more out of understanding a simple case fully than out of half-following a general proof, and they can always read the whole paper later if they want to know the details.

I further propose that we forever eliminate the following phrase from our writing: "The proof of this result is simple and left as an exercise to the reader." It sounds condescending and is almost guaranteed to make readers feel bad. Given that the reader is usually much less familiar with the subject matter than the author, chances are the reader will not find the proof especially simple. This communicates to the reader that they are unfortunately much less capable than the author expected. Moreover, the idea that the author gets to give the reader an "exercise" reinforces the idea that the reader is not the author's equal interlocutor, but more like their student (though of course an exceptionally smart student who can easily complete this exercise). Readers who are taking the time to engage with our work deserve better than being talked down to and being made to feel stupid. There are far more respectful ways of saying the same thing, such as "The proof of this is not included here, but if the reader wishes to reproduce it, the proof strategy is roughly to do X, Y, and Z." This way of phrasing the thought does not make assumptions about the readers' level of expertise, and does not sound condescending.

These tips presuppose that one's papers will be read by a wider audience. But for this to happen, one must be able to write for an audience of journal reviewers, otherwise one's papers will have little chance of ever reaching a wider readership in philosophy. A fabulous teaching strategy I picked up from Hájek when I was his postdoc at ANU is to coauthor referee reports for journals with graduate students. While one can of course explain to students how to write for publication in journals, I have found nothing as effective as turning the tables. To do so, I wait until I am asked to referee a paper by a journal on a topic that one of my students has some knowledge about. I usually don't have to wait all that long, given how desperate journals are for referees. The student and I then both read the paper, and we meet to discuss its merits and problems. I also explain to the student the intention and format of the referee report. After writing a brief summary that highlights some of the strengths of the paper, they should first list major comments and then minor comments, concluding with a verdict. I tell my students that the main task is to explain to the *editor* what the merits and problems of the paper are, and the secondary task is to give feedback to the *author*. Of course, if the verdict is "revise and resubmit", the author needs to be given enough information to know how to make suitable revisions. But if the judgment is "reject", it is not necessary to list every single reason why the paper is rejected, and it is also not the reviewer's job to fix the problems for the author. I sometimes show students old referee reports to illustrate what I mean, and then I ask them to write a first draft of the report. Once they have done that, they send it to me and I make final edits before I submit the report to the journal. I also show the edits to the student and explain why I made them. I sometimes edit to improve the structure and clarity

of the report, but I also often edit to make the tone less harsh. Graduate students are still learning how to clearly convey a critical point without being mean, and this is an excellent way to teach them how to do that. When I submit the report to the journal, I indicate to the editor that I collaborated with a particular student on the report, but that I have final responsibility for the content. Editors usually welcome this practice (if they comment on it at all).

Doing this once or twice tends to have a profound effect on how students see their own writing. They usually slip into the role of being a critical reviewer very easily, and once they look back at their own work with this perspective in mind, they see that their own papers have similar problems as the papers they've reviewed. It's always difficult to look at one's own work in progress with fresh and critical eyes, and this is an incredibly effective way of gaining a new perspective. Further, it helps me communicate better with them about their work. Once they've been in the reviewer's shoes, they understand much better what I mean when I tell them "a reviewer might get tripped up by this" or "a reviewer would probably not be satisfied with that."

The same advice holds not only for research articles and graduate students, but across the board. I became a lot better at writing grant applications for interdisciplinary review panels and recommendation letters addressed at people outside of philosophy after serving on a few review panels myself. In particular, I learned that people can't judge how good something is unless you tell them explicitly. On one occasion, I was reviewing applications for an interdisciplinary research prize at my university, and one of the applicants from the music department had performed at venues I had actually heard of, such as Carnegie Hall, and also recorded a couple of albums. But I found myself wondering – is that exceptional? Roughly what every decent music faculty member should be doing? Less than stellar? I had no idea, and the application materials provided none of the context I needed. This experience has really helped me in writing better recommendation letters and application materials, since I now reflect more explicitly on how to explain the value of a person's accomplishments to someone who is unfamiliar with the expectations of a particular discipline.

In sum, I recommend to my readers to sign up for reviewing tasks not just for altruistic reasons, but because being on the other side is an excellent opportunity to learn how to improve one's own writing.

6. Choosing a Narrative

6.1 Basic Idea

Any idea in philosophy has a particular origin and genealogy. For example, I encountered the problem of how to formalize degrees of rationality in a Bayesian framework (the topic of my first major research project) when I was writing a paper about the semantics of gradable adjectives. I was thinking about different linguistic theories that spelled out scale structures underlying various types of gradable adjectives, and I realized that I had not seen anyone explain how to spell out scales for "rational" or "coherent", even though those are clearly gradable. I spent several years working on this problem and ended up publishing multiple papers and a book about degrees of rationality in formal models. However, in presenting my ideas and motivating them, I found that it was much more compelling to choose a narrative that was distinct from this origin story. Instead

of talking about the semantics of gradable adjectives, I motivated the need for a theory of degrees of rationality by pointing out that theories of ideal rationality need a non-metaphorical account of how human thinkers can approximate rational ideals.

It is often difficult to see the narrative possibilities for presenting one's ideas that don't follow one's own trajectory of discovery. Thankfully, Hájek has us covered. He points out that there is a certain inventory of narrative arcs that tend to work well in philosophy papers, and that one can try out for one's idea to see which one is most compelling. Here are a few that he lists (Hájek 2022, p. 145):

- Hegelian dialectic: A thesis is countered by an antithesis, a resolution is found through a synthesis.
- Solving a problem: Set up a problem, explain why existing solutions fail, give a better one, finish with objections and replies.
- Desiderata (possible): Set up some desiderata that a theory needs to meet, compare different theories regarding how well they meet the desiderata and show that yours comes out on top.
- Desiderata (impossible): Set up some desiderata that a theory should meet, but then show that it's impossible to meet all of them. Explain which one(s) to drop, and which theory meets the surviving desiderata.
- Setting up a paradox or puzzle: Set out premises that are individually plausible yet jointly inconsistent. Squeeze out the least plausible premise.

And here are a few more from me:

- Corpse Reviver: Show that a theory that is thought to have fatal problems can be made to work.
- Counterexample + improved new theory: Give a counterexample to existing theories, use it to diagnose the nature of the problem and give a new theory that avoids it.
- Possibility Proof: Prove that a desirable result can be attained.
- Better argument: Give a better, simpler, more elegant argument or proof for an existing result.
- Generalization: Show that an argument or theory has a wider application than previously thought, or can be generalized to a new set of cases.
- Bug = Feature: Argue that a property of a view that is generally thought of as a problem is not really a problem but a desirable feature of the view.
- New concept: Show that a phenomenon can't be explained well with existing conceptual resources, then introduce a new concept that can do the job.

6.2 Implementation

The above list is of course not complete. A helpful task (which students should start on early) is to start a list of narrative arcs in papers. This can easily be incorporated into seminars. It's useful to set aside a few minutes to record which narrative arc a paper one is reading for the class is following, and to think about whether the same idea could have been presented compellingly with a different

narrative strategy. Once students have their own inventory of narrative structures to draw upon, they can try out different options for their own ideas before settling on a particular one. It is a good exercise, once students have settled on an idea to develop for a paper, to write up a few different mini-abstracts that try out different narrative arcs to find the one that is most compelling.

Regardless of the narrative structure one chooses for one's paper, Hájek reminds us that any paper benefits from a diamond shape:

> Begin with a 'satellite view' – very big picture, broad brush strokes. (The elevator pitch.) What is the topic, and why does it matter? Then, an 'aerial view': in more detail: what are the battle lines, the main positions? Who are the main players in the relevant debate? What side will you be on? Then, get into the 'trenches': address the relevant arguments in detail, and present your own. (Presumably, this will be the bulk of the paper.) Then, in the conclusion, zoom back to the satellite view, this time making clear what you have contributed. (Hájek 2022, p. 139)

The diamond shape helps with a common problem new writers have. They pay so much attention to the minute details of a debate and the nuances of their arguments that they forget to remind their readers why they were supposed to care about all of this to begin with. As a result, very interesting papers often undersell the importance of an argument, because the author never gets back to explaining how their view fits into and changes existing debates at a broader level.

However, I just as often see papers having the opposite problem: even if they contain some very interesting and compelling points, they oversell their overall contribution, which is a surefire way to irk readers, including reviewers. Hence, it is useful to add to one's inventory of narrative arcs different ways of explaining the strength of a contribution in order to prompt readers to form accurate expectations of what is to come. At the strongest end of things, one might have a general proof, a knock-down argument, a decisive counterexample, or a triviality or impossibility result. At the weaker end, one might have: a result that holds for typical cases, a compelling alternative to an existing view, some troubling edge cases, undesirable consequences, a tension between claims, a speculative interpretation, or an argument sketch. And so on.

Using this inventory of descriptions deftly can make a world of a difference in how one's arguments are received. Mark Schroeder once gave the following example to illustrate this point: Suppose you have two arguments for a point. One is very good, the other one still compelling but weaker. Telling readers in the introduction to the paper to expect two good arguments will leave them disappointed, as they will notice the weakness of the second argument. However, telling readers instead that the paper contains one good argument and one speculative idea regarding how to further support the point will adequately manage their expectations, leaving them satisfied with the result. In short, promise, and then deliver exactly what you promised.

7. Practicing

7.1 Basic Idea

The assumption that being good at philosophy requires some kind of innate genius is unfortunately still widespread (Leslie et al. 2015), and it suggests that practicing is either futile (because it won't make you a genius) or unnecessary (because you are a genius already). If, with Hájek, we instead accept that being good at philosophy is something that can be taught and learned, then we must emphasize the importance of practicing. Regarding practicing our writing, he reminds us that

> (…) it may be salutary to remind yourself that while achieving high standards in your writing is laudable, one way to *guarantee* that you *won't* achieve them is not to write anything! It may also help not to take your current piece of work too seriously. You're not writing the Gettysburg Address, after all! You should get into the habit of writing regularly, much as you should get into the habit of exercising regularly, even if you do it badly at first. Practice will make it better. Don't let 'writer's block' be a way of romanticizing laziness. (Hájek 2022, 229)

The best nugget of advice here is to not take any particular piece of writing, or writing session, too seriously. Many people who struggle with writing tell me that the conditions need to be just right for them to be able to write – having just enough time, feeling fresh, being in just the right workspace, and so on. These people hardly ever apply such demanding standards to the conditions under which they do other things. Somehow, they get their lessons planned, their exams graded, their dinners cooked, and their laundry done even when conditions aren't perfect. Writing should be no different. This also means that it should be normalized to "write for the trash." The job of a first draft is just to exist. That's it. Once it exists, I can decide what, if anything, I like about it and write a second draft that's a little better. I always save all my drafts individually in case I want to go back to an earlier version, but I am not precious about anything I've written. If it's not sparking joy, it must go.

7.2 Implementation

There are a few ways in which I encourage my students to practice writing without taking any particular draft too seriously. First, I encourage them to share things early in the process, including outlines, partial drafts, and abstracts. Of course, not every stage of development is fit for every audience – while I am happy to slog through a first draft written by one of my PhD students, they should probably wait to submit it to a conference or ask a senior expert at another school to read it until they've improved it more. But getting feedback early in the process means that problems can be spotted quickly without spending too much time heading down a dead end.

Also, I discourage students from paying too much attention to proofreading their work early in the process, especially if English is not their native language. I am not a native speaker myself, and when I first started writing in English, I spent an inordinate amount of time just trying to perfect my word choices, grammar and spelling, which had little to do with the philosophical substance or overall organization of my writing. There's no need to waste time on doing this during the early writing stages, as long as the content is sufficiently clear. In my view, professors who mark minor errors of this kind in preliminary drafts are just wasting everyone's time. As long as these

errors are fixed in the final version, all is well. (Also, since AI tools can now proofread almost as effectively as humans, this is a task I will happily outsource.)

I've also found that it has a profound impact on students when I share details about my own writing process with them. While many of them understand that practice leads to improvements, they think that once someone is a professor, this means that they have reached the stage where compelling arguments flow from their pens in perfect prose on their first attempt. They tend to be quite surprised when I tell them that this is not the case. Rather, I explain that there will always be a stage of the "purgatory of ideas" during which we're wrestling with what to say and how to say it. With more experience, the projects we take on become more ambitious too, so there's always a phase during which our ideas are still premature and in flux (Staffel 2025). Realizing that this flux stage is not a deficient stage, but a necessary and important part of the process of doing philosophy well has a very beneficial effect on philosophy students. It makes the uncertainty of not knowing how a project will turn out easier to bear, because it is reevaluated as being normal rather than problematic.

Writing is of course not the only thing philosophers are well-advised to practice. Giving a research presentation in a seminar, at a conference, or during a job interview, will go far better if the presenter has practiced it ahead of time. By practicing, we can make sure the argument flows naturally, we can speak without stumbling over sentences or terms, and the length of the presentation fits into the allotted time. Our audience invests a considerable amount of time into hearing us speak, and we should do them the favor of coming well-prepared. Making sure we stick to our allotted time is doing the bare minimum, and it is hard to do so unless one has practiced and timed one's presentation.

Practicing a talk can of course be a bit dull, and one's partner or pet might not be the most receptive audience. It's good to make practicing more fun. For example, in graduate school at USC, the PhD students held impromptu meetings of the *Highly Speculative Society*, often at someone's house, during which one of us would practice a presentation while the audience was enjoying snacks and drinks. Now when I need to practice a talk, I often do so while going for a walk. I have a nice trail through a nature reserve near my house that takes 45 minutes to complete, which is just the right length for many presentations. I print out my notes or handout, grab a cold beverage, and mumble to myself while enjoying my walk. If you're feeling self-conscious, put in headphones and pretend to be talking to someone on the phone.

Another skill where practice makes perfect is giving job interviews. I won't go into too many details here, since much of the advice from before applies equally well: it's good to think about who ones audience is and how to explain one's ideas to them in the simplest and most compelling way. Further, it helps a lot to practice interviewing from both sides, just like it helps with writing to see a paper through a reviewer's eyes. I recall Al Hájek telling me that he practiced answering every question he could possibly think of for several weeks leading up to the interview for his dream job at the Australian National University (which he landed, of course).

8. Polishing

8.1 Basic Idea

Now that you've become comfortable with the idea of writing crappy first drafts and practicing a lot, you might still wonder: But how do I get from there to a piece of writing that is really excellent? Getting feedback from others is of course incredibly helpful in improving one's work, but other people's time is a finite resource. It is thus good to have some strategies we can apply to our own work to find problems and fix them. We might want to improve the ideas in our papers, the organization or the writing style. I will propose a few tactics for each.

8.2 Implementation

First, how can we improve our ideas and find problems with our own arguments? A good first move is to turn Hájek's heuristics against ourselves. While they are of course helpful to come up with ways of arguing against other people's ideas, that doesn't mean we can't also use them to stress-test our own theories. For example, checking whether our own view generalizes to related cases, whether it is appropriately sensitive to whether a phenomenon is continuous or contains jumps, and whether it properly deals with edge cases is something we can systematically investigate. Further, by thinking about the structure of our own theory, we can think about what counterexamples to it would have to look like and see whether we can find any. It is also helpful, if one hasn't done so already, to put one's argument into premise-conclusion format and think about whether it meets the basic conditions for being a good argument. Are the premises true? Is it either valid, or a well-formed inductive or abductive argument? Explicitly using argument analysis methods from "baby logic" class can be shockingly helpful in finding weaknesses in one's own reasoning.

Another helpful tip, which I learned from my dissertation advisor Jake Ross, is to try to identify the "skeleton" of the argument. What are the absolute minimal assumptions that are needed to make the argument work? Where are the important choice points that the argument turns on? By stripping out anything that is not absolutely necessary, we can identify what to focus on and how to clarify the main point of the paper. This also includes identifying when we can help ourselves to other people's labor. We usually have to make choices about what to argue for and what to take for granted in a paper-length piece of philosophy. By identifying the skeleton of the argument, we can figure out the places on which our main argumentative focus must lie, and also identify where we can make starting assumptions that can be supported by existing arguments. For example, in my book on degrees of rationality, I realized that for most of my theorizing, I was relying on a value-based conception of epistemic rationality. I didn't want to devote a significant part of the book to defending this conception, so I directed my readers to some compelling arguments for it in the literature, and explained briefly how one might modify my view to accommodate alternative conceptions of epistemic rationality.

Suppose you are happy enough with your arguments and your main ideas, but the writing and organization of the paper need improvement. What can you do? One helpful strategy for improving the organization is to reverse-outline one's paper. This method is especially useful if one didn't use the modularization strategy from earlier, but wrote more freely without a tight structure in mind. The idea is to take the existing paper and write an outline that traces the main moves and arguments as they currently appear. This can reveal structural problems that can then be fixed in

the outline before going back and rewriting the draft in a better order. A similar idea is to take the paper and try to make it into a handout for a talk or a slide presentation (and to perhaps even try giving it as a talk, with or without an audience). Gaps in the argument, unnecessary detours, clunky repetitions, an excess of examples and confusing ways of ordering ideas can usually be identified with these methods.

If there are specific sections of the paper that seem not quite right, it can be helpful to step away from one's current text and to try to write the section from scratch, perhaps even in multiple versions. This can reveal new possibilities that remain hidden if one merely tries to improve one's current text. Changing the narrative at a micro-level can also help: If a concrete explanation doesn't quite work, try an abstract one, or vice versa. Try switching out the examples to see if new ones work better. If arguing directly for something isn't working, what about a reductio? Also, sometimes it's easier to explain something orally. Record yourself and transcribe what you said.

I recommend keeping a running list of one's writing problems that other people have identified in the past. Those problems tend to escape one's notice unless one explicitly hunts them down. For example, I have a tendency to often write "I think" or "it seems to me", and similar phrases. I now devote one round of proofreading just to looking for these expressions and deciding whether to delete them. I also specifically hunt down overly long sentences, weird uses of transition words like "however", and unnecessary repetitions. For other writers, the flaws to hunt for are likely different, but I bet everyone has a few writing habits or expressions they're keen to eliminate. It's a good idea to read one's paper out loud at the final editing stage, or to have a computer read it. It's an excellent method for checking whether one's prose flows nicely and there are no stumbling blocks, awkward sentences or weird transitions that need fixing.

Final Thoughts

This is just a very small slice of the wealth of insights and good advice that Alan Hájek has shared over the years with his students and readers. I've attempted to curate a manageable selection of tips and added to them my own variations and implementations. I tried to choose suggestions that I've found especially useful for myself and my students, and that are perhaps not used as widely as they should be. I hope to leave my readers feeling emboldened and equipped with some new tools to conquer uncharted philosophical territory and to tackle the manuscripts that await them on their journeys.

Acknowledgements

I am grateful to my PhD students Amir Ajalloeian, Bret Donnelly, Tyler Porter, Jonathyn Zapf and Yuri Zhu for reading an earlier draft of this paper and giving me very helpful feedback and advice.

Appendix:

My epistemology graduate seminar in the fall of 2019 came up with the following two lists, which really impressed me:

Q1: Ways of making a philosophical contribution

1. Show that an argument begs the question
2. Show that an argument or popular position has an undesirable conclusion
3. Show that a theory is both over- and under inclusive
4. Show that a new concept is less specific/useful than a previous concept
5. Reject a common-sense intuition
6. Respond to a critique
7. Offer a new distinction
8. Offer a new argument for a position
9. Explain the importance of a view or question
10. Explain or explain away a disagreement
11. Ask new question(s)
12. Reconcile existing views
13. Show an implication of/expand a current view
14. Amend a current view
15. Clarify an existing position
16. Give an overview of the literature
17. Set up a dilemma
18. Formulate a puzzle
19. Give a formal model
20. Offer a new interpretation of a historically important text

Q2: Ways of coming up with philosophical ideas

1. Start thinking early in the semester
2. Talk to people that know what's up
3. Compile a lit review, look for gaps
4. Get emotionally invested: write against the view you like
5. Scan footnotes/conclusions for questions authors leave unanswered
6. Cross-pollination: Combine views from disparate areas for exciting results
7. What would theory X say about cases of class C?
8. Go to talks outside your main area of philosophy for new ideas
9. What are your favorite authors saying? What about authors you dislike?
10. Take a break
11. Just start writing - it will help you see which ideas are promising
12. Explain your ideas to non-philosophers
13. Crowdsource ideas from class discussion
14. Don't grant authors' assumptions
15. When someone says something is a big/small cost of a view, show that it isn't
16. Show that friendly/rival views are incompatible/compatible
17. Compile a document of choice quotes
18. Peruse the news for real-life cases

References

Forbes, Graeme. 1994. *Modern Logic*. Oxford University Press.

Gettier, Edmund L. 1963. Is Justified True Belief Knowledge? *Analysis* 23 (6), 121-123.

Hájek, Alan. 2014. Philosophical heuristics and philosophical creativity. In: E. S. Paul, and S. B. Kaufman (eds), *The Philosophy of Creativity: New Essays*. Oxford University Press, 288–318.

Hájek, Alan. 2016. Philosophical heuristics and philosophical methodology. In H. Cappelen, T. S. Gendler and J. Hawthorne (eds), *The Oxford Handbook of Philosophical Methodology*. Oxford University Press, Chap. 19.

Hájek, Alan. 2022. *Heuristics for Philosophical Writing*. Unpublished manuscript, draft of 22 February 2022.

Kallestrup, Jesper. 2025. *Methods and Skills for Philosophy. An Advanced Guide*. Routledge.

Kelley, David. 2024. Philosophical Moves. *Australasian Journal of Philosophy* 102(3), 537-550.

Leslie, Sarah-Jane; Cimpian, Andrei; Meyer, Meredith and Freeland, Edward. 2015. Expectations of Brilliance Underlie Gender Distributions Across Academic Disciplines. *Science* 347 (6219), 262-265.

Schlenker, Philippe. 2002. A Plea for Monsters. *Linguistics and Philosophy* 26(1), 29-120.

Staffel, Julia. 2025. In the Purgatory of Ideas. On the Transitional Nature of Rational Philosophical Attitudes. In: Sanford C. Goldberg & Mark Walker (eds.). *Attitude in Philosophy*. Oxford University Press. Forthcoming.

Williamson, Timothy. 2024. *Overfitting and Heuristics in Philosophy*. Oxford University Press.