

## How to Write a Philosophy Paper

Shelly Kagan  
Department of Philosophy

1. Every paper you write for me will be based on the same basic assignment: state a thesis and defend it. That is, you must stake out a position that you take to be correct, and then you must offer arguments for that view, consider objections, and reply to those objections. Put another way: you must give reasons to believe the central thesis of the paper.

Some of you may have never written a paper like this before. So let me contrast it with two other kinds of papers you probably have written. First, I am not looking for "book reports": I don't want summaries of one or more of the readings, and I don't want you to "compare and contrast" what different authors say, or what different moral theories might say, about the given topic. Rather, I want you to "stick your neck out"--tell me what you believe to be the truth about the relevant issue. And then I want you to defend that position. Of course, it might well be relevant, or helpful (or perhaps even part of the specific assignment) to discuss some particular author or view. But even here the emphasis should be on evaluating that author or view. Book reports, no matter how superb, simply don't meet the assignment.

Second, I am not looking for "thoughts on topic X", or "meditations on X", or "remarks on X". It just won't do to simply string together various reflections you may have on the given topic, even if in the course of doing this you embrace various claims, and offer some reasons for your views, before moving on to the next reflection. The paper should instead have a single, central thesis. The point of the paper is to state and defend that thesis. The various contents of the paper should be selected and organized so as best to defend that central claim. (Stream of consciousness, for example, is a poor way to organize material, and is likely to include much that is irrelevant to anything like a main thesis.) No matter how brilliant, mere reflections simply don't meet the assignment.

2. So the paper should state a central thesis, and defend it. More specifically, it should do the following things:

A) It should start with a clear statement of the thesis. This need not be the very first sentence, but it should almost always come in the first paragraph or two. Many papers never give a clear, official, precise statement of the thesis at all. They leave the reader guessing what the author's precise position is. Don't do this. It may seem obvious to you what your position is; it won't be obvious to me.

Other papers do give a clear statement of the thesis--but not until the last paragraph of the paper. ("Oh! Is that what the author was driving at!") This may be a dramatic way to structure a short story; but it is a very poor way to structure a philosophy paper. I shouldn't have to spend my time reading the paper trying to figure out what your view is; I should be able to spend my time deciding whether you do a good job of defending that view.

It is likely that if you cannot state the central thesis of the paper in a clear, crisp, precise

Yale



sentence, then you don't really know yet what you are trying to say. Try stating your thesis out loud. If you find yourself rambling on, then this is a sign you should think harder about what exactly your point is.

(Admittedly, sometimes a reasonably short statement of your view neglects certain qualifications that a fuller statement would include. It is OK to offer the general, slightly imprecise statement of the thesis first, and then say something like "Of course this claim needs to be qualified as follows. . .", or "Of course this needs to be qualified somewhat, and I will note these qualifications, when appropriate, later in the paper." But if you find yourself doing this often, that is a bad sign.)

B) Having stated your thesis, you need to defend that thesis--to give arguments, reasons to believe it. Philosophy is not just a matter of asserting opinions, however profound; it requires backing up those opinions, as best one can, with plausible grounds. It is, of course, difficult to say anything helpful at this level of generality about what a good argument would look like, or what sorts of arguments will be most forceful, most useful, and so forth. To a large extent this will depend on the specific thesis being defended. But I do want to point out that learning how to develop and present plausible and convincing arguments is a skill that takes practice like any other skill. If you work at it, you will get better at it.

I also want to warn you against a common error. Some people try to offer as many arguments as possible for their view. There might be a dozen or more in a 5 page paper, each argument taking up a sentence or two, or at most a paragraph. This is not a good way to defend a thesis, for the inevitable result will be that no single argument is developed with enough care to be convincing. Everything will be superficial. What you should do instead is take what you think is the one best argument for your view, or at most the two best arguments--the most important and convincing ones--and devote the paper to spelling them out, filling in the details, developing them. To be honest, a short paper is often barely enough space to present a single forceful argument with the care and attention it deserves. So you will do a much more persuasive job of arguing for your view if you use the available space to spell out exactly how your single, central argument works.

C) It is almost always possible to come up with something to say on behalf of a view, no matter how stupid that view is. So the real test is to see how well you can defend your position against objections. Thus a good philosophy paper will always consider one or more worrisome objections, and then try to reply to those objections.

Once again, it is difficult to say in a general way what the best objections will look like. Sometimes they will note potential flaws in the arguments you have given (thus giving reason to think that the thesis--whether true or not--has not been well defended). Other times they will directly attack the thesis itself (thus giving reason to think that the thesis must be false). Here, too, learning how to recognize and present the most significant objections is a skill that should get better with practice. But a helpful idea is this: imagine a really intelligent person who is not convinced by what you have said. What, exactly, are they likely to be complaining about?

Here as well, it is important to avoid the common mistake of simply trying to state as many objections as possible. Better to take the one or two objections that seem the most devastating, or most interesting, or most worrisome. Concentrate on those.

Having given the objection, or objections, it is of course crucial to try to offer some sort of reply to those objections. After all, if the objection is truly worth offering in the first place, it

should be raising a serious worry. Very well, then, where exactly does this objection go wrong? Why doesn't the objection convince you to abandon your view? Is the objection based on some sort of mistake? Or does it force you to modify your view in some way? You should not pretend that the objections are completely idiotic. (If they are, you shouldn't bother giving them.) You should just give your best attempt at answering them.

D) What I have said so far may give the impression that the right way to organize the paper is like this: Thesis, Argument, Argument, Objection, Objection, Reply, Reply. In fact, however, this is a poor way to organize a paper, for with an approach like this the various parts of the paper are not logically integrated. Better to raise the objections at the relevant stage of the argument. (Or perhaps you should state the argument briefly, and then go on to raise objections to it.) Better to reply to the given objection immediately after having raised it, and then move on to a new objection, or perhaps to a new argument. If you do this, you might end up with something more like the following: Thesis, Argument, Objection, Reply, New Objection, Reply, New Argument, Objection, Reply. Obviously, there is nothing magic about this outline either; a great deal will depend on the details. But you should do what you can to integrate the various elements of the paper.

E) There are very few "knock down" arguments in philosophy. It is almost always the case that there is more than one plausible view on a given subject, and there are arguments for and against the particular position you are defending. Thus, if you have done a fair job of presenting your thesis, offering arguments, offering objections, and replying to those objections, then the fact of the matter is that when the dust settles down, you will have to admit--if you are honest--that there is something to be said on both sides of the issue.

That doesn't mean that you must think it is a tie. Even if both sides are still "standing", you might still think that, on balance, one side of the argument has the better of it: it seems more convincing, more plausible, better able to defend itself. Hopefully, this is the side you set out to defend. (If not, switch sides, and write the paper again!) So a good place to end the paper is by briefly reviewing the major advantages and disadvantages of your view (as compared to the leading alternative, perhaps)--and explaining why, on balance, you think that your view remains the most plausible position.

Again, your goal should not be to try to convince everyone who reads the paper: that's a goal that it's almost always impossible to meet. Your goal, rather, should be to give the best possible defense for your view, recognizing that, even when all is said and done, others may still disagree.

3. In writing and working on the paper, what should you do? The first thing you should do is this: think. You should do a lot of thinking before you try to write anything. You need to think about the topic. You need to think about the particular issue you will be writing on. You need to think about the specific assignment (if there is one). You need to think about your views. You need to think about what position seems plausible to you. You need to think about your reasons for believing that view. You need to think about the weak points in your position. You need to think about possible replies. And you need to think about whether--in the light of all of this--you need to think again. So don't expect to sit down the night before the paper is due and throw something off. That just doesn't leave you enough time for thinking (and changing your mind).

The first thing you should do, then, is give yourself several days just to think about things, and to try to figure out what you believe.

Once you've got a reasonably good fix on what you want to say, write a draft of the paper. Don't think of it as the final product. That may only cramp your writing and make you afraid to discover new ideas (new problems, new questions, new arguments, new objections) in the course of writing. Don't worry about getting it perfect. Worry about getting it down. Then, having written a draft of the whole paper--put it away! Sleep on it. Look at it again in the cold light of morning, or after a day spent thinking about something else. With fresh eyes you will be better able to see what makes sense, and what doesn't, what's relevant, and what should go, what's convincing, and what needs further development, what's clear, and what isn't, what's good, and what's garbage. Then, with an eye to all of this (and after having thought about the new difficulties you've discovered) write the paper again. Maybe you will want to discard the original, and write it all over, from scratch as it were. Or maybe you will want to use the original as a basis, and revise, expand, correct, and cut, as needed. Either way, the way to make it better is to rewrite it.

I don't think I can emphasize this last point enough. There is absolutely nothing you can do to improve the quality of your writing as much as getting into the habit of rewriting your papers. Nothing.

Ideally, even this revised version of the paper would not be your final version. You might revise it again. And maybe again. But even when you finally have a version that you are truly happy with, there is still something else you could do that would make it better still: show it to a friend. Find someone else in the class, or just some friend who seems reasonably intelligent, and ask them to take a few minutes to read your paper. Then ask them to tell you--in their own words--what your basic thesis is, and how you go about arguing for it. If they can't do this, then the odds are pretty good that the paper isn't as clear as it should be, or isn't as well organized as it should be--and this will give you a reasonable idea of what still needs to be rewritten. If you are really lucky, in the course of the conversation your friend may point out some important objection you overlooked. Or you may end up seeing how to make your point more clearly, more simply, more convincingly. So then you can go back and improve the paper some more. (Be sure to include a footnote acknowledging your friend's ideas.)

Now I am under no illusion that many of you have the time (or the inclination) to do all of this: write, revise, revise again, share, and revise yet again. I realize that you have other classes, and other responsibilities--and you also have a life. But you can think of this elaborate process as some kind of ideal, and do what you can. For rewriting really is the secret of good writing. The more of this you do, the better your writing will be.

4. Let me mention some of the virtues that are especially important in a good philosophy paper. Obviously enough, it is important to have good arguments: most of your grade will be a function of how good a job you do of defending your thesis. But there are other, more general, aspects of good writing that need to be emphasized as well. (I've already mentioned most of these, but they bear repeating.) First, the paper should be well-organized. That is, the underlying outline should reflect a logical progression of ideas. Closely connected, second, the paper should only contain material relevant to the official purpose of the paper. The idea of the paper is to state and defend a central thesis. Many ideas may be reasonably relevant to the general topic you are

discussing, but nonetheless altogether irrelevant to your particular position or to your particular arguments. Cut out anything that isn't actually doing some work for you.

Third, write in simple, jargon-free English. Long, convoluted sentences can be difficult or impossible to understand. And jargon can get in the way of communicating, or--even worse--it can mask ambiguities and confusion; it can get in the way of your seeing just what it is that you really don't understand. Sometimes people think that profound topics need to be discussed in a language that looks profound. But that's just a mistake. There is no reason why you can't express your thoughts in simple, straightforward English. (Of course sometimes a certain amount of philosophical jargon may be unavoidable, or even helpful. But it should be used sparingly; and the meaning of any such terms used should always be carefully explained in the paper.)

Fourth, and closely connected to the last point, it is absolutely crucial that you strive to be as utterly clear as you possibly can be. Spell out your ideas as carefully as you can. Make it explicit. Make it all explicit. Don't "hint" at things, or "point" in the direction of your thoughts. Don't assume that it will be "obvious" what you have in mind. (It won't be at all obvious, and the chances are good that I will be confused, or at least uncertain. I haven't been living in your head all these years.) Spell out all the steps of the argument, even the ones that seem to you to be too obvious to need stating. Explain exactly what the train of thought is--how and why we are moving from A to B, but skipping over C, to get to K! Put in "road markers" for the reader, reviewing what the paper has shown so far, and announcing where the paper still has to go to in order to arrive at its destination. Say things like: "In order to defend my position, I am going to offer two arguments. The first argument has three premises. . . Now let's consider each of them. The first premise was. . . Here is why I think this premise is plausible. . . But here is an objection to that premise. . . and here is my reply. . . That covers the first of the three premises. The second, you'll recall was. . ." --and so forth, and so on. Spell it all out, as clearly as you possibly can.

You may reach a point where you think you are beating a dead horse on this clarity thing. That's fine; I want you to beat the horse anyway. In all my years of teaching I have never seen a paper that goes too far in this direction, but I have seen hundreds that don't go anywhere near far enough. So please: do everything you possibly can to make your paper as clear as it possibly can be. (If your paper is the one that finally goes overboard in this regard, I will happily tell you to ease up slightly on the next one.)

One further aspect of spelling things out clearly may deserve special mention. If you do quote from someone in your paper--even if it is from one of the assigned readings--this should always be followed up by your own explanation of what the quotation means. (Typically, after all, quotes don't wear their meanings on their sleeves.) You may need to explain any technical terms that the author uses, or summarize the author's main reasons for holding that particular view. In any event, quotations (like jargon) should be used sparingly; keep it to a minimum.

5. Here are two things that you don't need to do in writing a paper for me. First, you do not need to do "research"--at least, not if what you mean by research is doing extra reading, looking at other works by one of our authors, or at secondary literature on some topic. Obviously, there is an important place for such scholarly research, and for many classes doing this well is an important preparation for writing a paper. But for my classes there is (usually) no need to do anything like this at all. Indeed, I would strongly prefer if you not do this kind of research. For what I am really

interested in is seeing what you can come up with through thinking about the relevant philosophical issues. I want you to wrestle with the issues first hand, rather than getting your ideas or inspiration from the thoughts of another. In short, I want you to devote your time to philosophical reflection, rather than scholarly research. (I also want you to come to appreciate the way in which philosophical reflection is a way of doing research in philosophy--indeed, the primary way.)

(Needless to say, if you do disregard my advice and do outside research anyway, and then find yourself incorporating ideas from other works into your own paper, you must properly and explicitly acknowledge your debts in the paper. And the same is true, of course, even if your debt is to a "mere" friend, rather than to some famous philosopher.)

But if you don't do scholarly research, to find out what views have already been discussed, and what sorts of arguments and objections are out there, isn't there a serious chance that you will just end up saying something that someone else has already said? Yes, it's true, there is a significant chance that your paper won't be "original" in that way. But that's OK, too--for the second thing that you don't have to do in writing a paper for me is come up with something that no one has said before. I want you to get some experience in thinking through philosophical problems for yourself. You can do this just as well--and learn just as much from the experience--even if you end up reinventing the wheel. After all, if you do reinvent something as fantastic as the wheel (or the philosophical equivalent), that would really be quite an achievement indeed. Of course, it will only be an achievement if you do genuinely reinvent it: if all that happened was you read about the wheel from some unassigned reading, and then wrote about it, that would hardly be evidence of any kind of creative talent on your part. (So that's another reason I usually don't want you to do "research".)

6. Before finishing, it may be helpful to say something about picking a topic and a thesis. For many of you, this won't be much of an issue. Usually, for large, lower-level courses, I tend to assign a choice of specific paper topics. In that case, of course, once you decide which question appeals to you--and which one you have something to say about--choosing a thesis will be fairly straightforward: it's mostly a matter of deciding which side you come down on.

But for some classes (usually upper-level courses) I prefer to leave things fairly wide open. If that's the case, then you can generally write on anything at all connected (reasonably closely) to the subject matter of the course. You might pick a topic that we've talked about in class, even though none of the readings discuss it. Or it might be something the readings discuss, even though we haven't talked about it in class. Or it might be an issue which goes beyond anything we've talked about, or read about, but which is logically connected to issues that have been discussed. I won't quite say that anything goes--the topic really does need to be connected to the subject matter of the class--but I do tend to be fairly liberal about this. (If you're not sure whether a given topic is legitimate, just ask me.)

In any event, assuming you can tell what's relevant to the course and what isn't, the really crucial matter is finding a topic where you have something to say! My hope is this: at one or more points during the term, you have found yourself thinking that the class discussion just didn't do justice to some point of view that seems attractive to you. Or perhaps there was some view that got mentioned in the readings, and even though we didn't discuss it, you think it's worth defending. Or maybe you think it would be illuminating to point out where the view goes wrong. Whatever.

Think back on the moments when you felt you had something more to say, something that goes beyond just repeating points already made in the readings, or already made by others in class discussion. Pick one of those issues as your topic. Then you can concentrate on trying to figure out exactly what you want to claim. That's a matter of clarifying and refining your thesis.

It can certainly take some work to find just the right topic, and just the right thesis. Part of the challenge, of course, is to pick an interesting topic and an interesting thesis (there is no point in defending trivial or uncontroversial claims). But it obviously won't help to pick something fascinating and controversial if you don't have anything plausible to say in defense of your chosen point of view. Sometimes you just have to admit that despite your interest in a given issue, or despite your belief in a given position, you just don't have enough to say. Then you have to move on and try something else.

At other times you will have to admit that you have too much to say, that you can't possibly defend your position adequately in the limited space available. It is important to find something "the right size". If the paper is too ambitious, you simply won't be able to spell out and defend your view adequately, and you will end up being superficial and hasty or unclear. Fortunately, with a little work, you can often find some way to scale back on overly ambitious projects--find a piece of it that makes for a more manageable paper. Other times, however, you'll just have to go back to the drawing board, and try to find something else.

7. No doubt, some of you have a fair amount of experience writing the kind of paper I have been describing. But others, as I noted at the beginning, may have little or no experience doing this sort of thing. Either way, it's pretty likely that for a lot of you, your initial attempts won't represent the best you can do. But that's OK too (though I certainly hope you will give each paper your best shot). I am not especially interested in whether or not you come into my class knowing how to write the kind of paper I am looking for. I mostly care about whether you come out of my class knowing how to do this. So if your work shows a clear pattern of improvement over the course of the term, I'll count the later, better papers much more heavily than the earlier, initial attempts. That's a promise.

Copyright © 2007 Yale University. Some rights reserved. Unless otherwise indicated on this document or on the Open Yale Courses web site, all content is licensed under a Creative Commons License (Attribution-NonCommercial-ShareAlike 3.0).