

Approaching the Dissertation

The way you approach your dissertation should be determined by a clear understanding of your incentive structure. The dissertation's role in a lot of graduate students' lives is: something that *looms*. But it shouldn't be. If you have clear idea of what it needs to do for you, you'll have a much clearer idea of what it should look like, and so of what steps you need to take, and in what order.

On the one hand, don't think of it as something so large and daunting you can't see how to get started on it. It's your final homework assignment: *homework*—not the next *Critique of Pure Reason*. (If you think of it as your magnum opus, something dauntingly and impossibly big, it's easy to become paralyzed, and spend a lot of time doing things you don't need to do, like procrastinating and twiddling your thumbs.) Your dissertation consists of five or so chapters on suitably related topics. So don't get too bent out of shape: after all, you've already written lots of term papers already, right?

On the other hand, if you're going to do it right, you should expect to work harder, and with more focus, than you've ever worked before. Let me explain why... and *how*.

First, when you write your dissertation, you're building up intellectual capital, capital that you'll be spending down over the course of your half-life doing one-years and as an assistant professor. If you're TAing right now, you may *think* you're overworked; when you start up as a full-time faculty member coming up to speed on a difficult new set of skills, *that's* when you find out what it's like to be overworked. If you're doing a string of one-years before you get your first tenure-track job, that's even more true: being on the job market is itself quite time consuming. This means that you won't have *time*, as a junior faculty, to do a lot of new research. But, to keep your job, you'll have to publish a great deal. That means you'll normally be reworking material you by and large already have into articles and maybe a book. What material? Almost always, that means your dissertation. So you want your dissertation to have content that you can mine, refine into articles (or perhaps a book), and so get tenure.

Second, when a department makes a hire, they mean to hire someone who's an expert in some area of philosophy. It's normally the process of writing a dissertation that makes you into that expert. This means that your

dissertation should be executed in a way that makes you as knowledgeable as anyone about its topic (and fairly knowledgeable about closely surrounding areas). Plan it so that a side effect is expertise in your AOS.

Third, remember that philosophy puts a premium on originality. Plan to come out of the dissertation with a set of ideas (and arguments) that are all your own. They *do* have to be connected with some familiar body of extant philosophical work, or people who might hire and publish you won't necessarily see the point of your project. But they also have to be original, new, and yours. That means that you have to be inventing as you write.

So here's what you need to do.

First, select a topic that will sustain your interest, not just for the two or three years it takes to write a thesis, but for the at-least-several years afterwards that you'll be mining the dissertation. In philosophy, you can change areas of research interest, but realistically, you won't be able to do that until you're several years out. While you're at it, select a topic that will be interesting to other people, too—the ones who will hire and publish you.

That means being able to motivate your dissertation: to explain, convincingly, to yourself and to others (others who don't know a whole lot about it) *why* it's interesting. The dissertation proposal is the stage at which to make sure that you've got that down. Take a look at "Dissertation Proposal Guidelines," on my web site, before working up your proposal.

Second, you need to schedule a regular pace of reading, producing arguments, writing, and polishing.

Reading: My estimate is that if you read one paper or chapter a week, you'll be in the position of understanding and having critically assessed most of the relevant positions in a reasonably circumscribed literature. (Some areas, especially in the history of philosophy, may require a heavier reading load.) Of course, that won't work if you're reading randomly: read selectively, working your way through the most discussed pieces in the circumscribed area of your dissertation topic. Read actively: extract arguments, outline the articles, and criticize the positions and arguments in what you read.

While you're reading, bear in mind that there are model dissertations out there: working through one or two of them will give you a sense of what you're aiming for. A few years back, I taught John MacFarlane's dissertation (online at his Berkeley web page), which is exemplary in the effort it puts into motivating his project. There's also a series, called Garland Dissertations in Philosophy; these are successful dissertations that have been published, as is. There may well be one of them in your own area, and if there is, take a

look at it.

Producing arguments: You need to be making up your own arguments; after all, they're the raw material of the dissertation. You'll need more arguments than you end up using. A lot of them will just be thrown away, if you're doing it right, and you want to have enough to be able to select the good ones. Ideas and command of the material will come as you make up arguments, and normally not otherwise. So this should be an ongoing background process. (Every weekend, ask yourself, "Have I made up a new argument this week?") A good aid here is to keep a journal, and make sure you write regularly in it.

Writing: As ideas come together, write them up. Don't leave this till the end. Write regularly. Revise and polish your writing, several times, before you turn it in to your advisor.

Polishing: Last, least, but important. You may have come to graduate school thinking that you know how to write, but chances are your writing needs work. (Are you thinking, "That doesn't mean me?" Think again.) The gloss level of your writing has to end up very high indeed. Bear in mind that in the real world—i.e., once you have a job—as far as grammar, style, etc. goes, *everything* has to be perfect, *all* the time. If it's not, your journal submission will just be thrown out, your job dossier will be thrown out, and so on. From here on in, perfection along this dimension is the *minimum*.

Time to practice that. It's fine if the ideas and arguments you turn in are rough: that's what an advisor is for. But your advisor shouldn't have to line edit your work for grammar or style. Develop a pool of colleagues you can turn to, on a regular basis, for editing and proofreading (and of course for criticizing your content, but that's not what I'm stressing right now). *Before* you show a piece of material to your advisor, run it by them, and make sure the prose is clean. When they point out problems with your writing, figure out what the problem was, and internalize the fixes.

While you're doing all of this, you're going to have a dissertation advisor. Let me say something about my own expectations as a dissertation supervisor.

If I'm your advisor (and, probably, even if I'm not), it's best if you don't approach the dissertation with a 'tell me what to do and I'll do it' attitude. There's an intrinsic reason, and an extrinsic one. The intrinsic reason is, of course, that you need to engage in the intellectually rewarding process of finding your own direction, making up your own ideas, and getting there

yourself. If you're not doing that, there's no point in being in the business: you'd have to be crazy to be in philosophy for the *money*. And if the direction is coming solely from the outside, over the long run, you'll be bored stiff.

The extrinsic reason is that after you complete your dissertation, you will need a letter of recommendation from your advisor. If all the letter can describe is handholding, it won't be good enough.

My own preferred role is that of a filter. Bring me ideas, arguments and draft material. I'll tell you which I think are more promising, and which look less promising to me. I'll make suggestions as to how you might pursue or extend the ideas you have. I'll object to your arguments, and complain about your prose.

If things are going well, you'll be meeting regularly with your advisor. In my case, that means having something written for me to read before we meet. (In other words: If I'm your advisor, I want you to be bringing me the new ideas and new arguments, cleanly written up, on a regular basis.) If something is going wrong—if you're stuck, if you're not making headway on your topic, if you're procrastinating, whatever—keeping a low profile just about always makes it worse. So don't hide!

Expect your advisor, whoever he or she is, to convey regular and frank progress reports to the DGS and the Graduate Committee.

A dissertation can and should be an enormously rewarding experience, and a stepping stone to an academic career. I may not have been making it sound this way, but a dissertation can be *fun*—hard, but *fun*. And once you see how it's broken down into bite-size tasks, a dissertation is, with proper organization, quite manageable. But if a dissertation is going to do its job, it's also an enormously demanding phase of your life. To make the dissertation work, you have to make it your top priority. You have to read (selectively, regularly, actively). You have to produce raw material—arguments and written draft—on a regular basis. You have to polish your writing to a high gloss. Doing all of this is the path to a strong and successful dissertation.