Paul Niehaus

82 Followers About Follow

Doing research

Paul Niehaus Nov 23, 2019 · 9 min read

Some generic advice for PhD students

The loss of structure

In between your second and third year of the PhD program, most students experience a dramatic loss of structure.

By structure I mean guidance about what to do, when to do it, and how well you are doing at it. During your first two years, your teachers give you problem sets and tell you when to turn them in; they tell you when tests will be and what to study for them; they will grade your work. After second year, this structure largely disappears. All of a sudden, *you* become the one responsible not only for doing the work, but also for figuring out what needs to be done.

You should know that most people experience transitions like this to some extent during their careers. An individual contributor who gets promoted to middle management, for example, suddenly becomes responsible not for doing what (s)he is told but instead for figuring out what others need to do. The transition seems to be particularly abrupt and challenging for PhD students, however, especially if you have not experienced periods of being largely self-directed in your earlier work.

It is helpful to prepare for this transition both cognitively and emotionally.

Creating structure: goals vs systems

PhD students tend to be fairly goal-oriented, and to have at least a basic idea of their main goal for the last 3(+) years of the program: write a good job market paper and get a good job. It can be tempting to focus on this goal, and try to backsolve from it to figure out the path by which you can reach it. I think there are two problems with this approach.

First, thinking constantly about the end goal can be daunting. When you see the creativity and polish of other people's job market papers, you may feel hopeless — how could ever get from where you are now to producing something that impressive?

Second, for most people the creative process seems to be far from linear. That is, while you might wish that you could identify a series of steps A, B and C which if you take them in order will reliably yield a good job market paper, this is largely just that: a wish. For most people there seems to be no such linear algorithm.

An alternative way to create structure is to focus on your *system*. By system I mean the habits and routines you develop and practice on a regular basis. For example, attending the seminar each week, writing down three suggestions for ways to improve the paper, and meeting afterwards with a classmate to discuss your ideas is a practice you might incorporate into your system. In designing your system, your aim is to give yourself a high probability of eventually accomplishing your ultimate goal (come up with a great job market paper) even though you cannot predict with any certainty the sequence of events through which this will come about.

You should of course think periodically about ways to improve your system, but you also need to trust in it to some extent: believe that if you are doing the right things day-to-day, this is putting you in a good position to eventually achieve your goal. This can help to reduce the feeling of being overwhelmed by the remaining distance to your ultimate goal.

Specifics vary, but most people find that regular conversation both with fellow students and with faculty advisors is an important part of their system, both to make progress and to stay grounded and happy.

A three-buffer system

"Creativity" often seems to involve noticing connections between things that others had not seen, or had previously seen as unconnected. I find it helpful to think about building systems that give you opportunities to notice these connections. In particular, I

think in terms of giving myself opportunities to notice connections between three categories of things.

The first category is the state of the conversation in your field. What questions are other people asking? What questions are they *not* asking? What things do they believe have been proven? What implicit assumptions are they making? What do they see as the big challenges in making further progress?

The second category is reality. What is in fact *actually happening* in the real world in the area you want to study? What decisions are people making? How is their environment changing, and why? What facts or patterns seem important or surprising to you?

The third category is research opportunities. What datasets exist that might be useful in your area of work? What natural experiments exist that you could potentially use to examine causal relationships? What partners in the public or private sector might potentially be interested in collaborating with you?

At the end of the day, I think most good research ideas involve finding connections between these things — identifying a research opportunity to demonstrate something about reality that advances the conversation in your field. The key idea is that in order to put yourself in a good position to spot these connections, you need to load information about each into memory to give your brain the opportunity to connect the dots.

This exercise is non-linear in the sense that you do not learn about things because you *know* you will need them, but rather because knowing them increases the number of potential connections you *could* spot in the future as you learn other things. For example, you may learn about an available dataset today and see no immediate way of using it, but then months later have a question about a paper you are reading and recollect that there was a variable in that dataset that you could use to address it.

To fill your three memory buffers, you need to build practices for each. For example, you might learn about the state of the research conversation by regularly reading and taking notes on papers, by regularly attending seminars or conferences, etc. You might learn about what is happening in the "real world" by reading more descriptive work (e.g. ethnography), by having lunch once a month with someone working in a relevant area (e.g. the CFO of a company, if you are studying corporate finance), by attending practitioner events, etc. And you might learn about research opportunities by conducting systematic reviews of legal reforms, networking to certain types of

organizations with whom you are interested in partnering, etc. The key thing is to do these on a regular, disciplined basis and not abandon your practices simply because they do not immediately lead to a great insight. Instead, think of this process as a long-term investment, filling your memory buffers and giving your mind more opportunities to make connections.

Making connections

The actual process of making connections is the most mysterious part of the creative process. Most people find they cannot force it, and only become frustrated and disappointed with themselves if they try. That said, I think there are a few practices you can build into your system that likely help.

One is synthesis. As you examine some new thing, ask — what is the bottom line? What is truly unusual about this? Etc. Keep brief notes on these things (partly for future reference, but largely because the discipline of writing them will increase retention in memory).

Another is conversation. Talk with fellow students about what you are seeing, thinking, etc. The act of verbalizing our thoughts often helps us to better understand them ourselves, and the back-and-forth of conversation can help us draw more interesting and indirect connections — speaker 1 connects A to B, then speaker 2 connects B to C, etc.

A third is sleep. There seems to be pretty good lab evidence that the brain does some of its most creative connection-finding while you are asleep, and particularly during REM sleep. This is why dreams often involve strange, unusual or unrealistic combinations of ideas. Give yourself the opportunity to get a good night of sleep every night.

Failing fast

Once you have identified a potential research idea, a key parameter is the speed with which you discover whether it is in fact worth working on it to completion (i.e. a finished paper) or discarding it and moving on.

To see this, suppose that one in five of your reasonably serious research ideas ends up being worth writing a paper about; this seems like a pretty typical “hit rate” for many students. Imagine that you take every “bad” idea 25% of the way to completion before you discard it; then to write one finished paper you have to do the work of writing *two*

finished papers — 1x100% on the paper you actually write, and 4x25% on the four papers you do not. Now imagine that you take each "bad" idea 50% of the way to completion; you would have to do the work of *three* papers to produce one actual paper. You get the idea. The upshot is that if a project is going to fail you want it to "fail fast" so you can move on to something better.

One practice that can help you fail fast is to meet regularly with an advisor or a fellow student to run through new research ideas and get rapid feedback on them. Which have limited upside and should be discarded immediately? Which are worth further exploration? Which are good enough to become your primary focus?

Another helpful discipline is to identify and test a new project's *fatal weaknesses* first. Suppose that for a project to come together, three things would need to be true; the first two are true with probability 80%, while the third is true only with probability 20%. All else equal, you should first try to figure out whether the third thing is true or not. More than likely it is not, and once you discover this you can discard the project and move on without wasting time on the first two things. (The temptation can sometimes be to do the opposite — to spend time on the parts of the project that are likely to work and thus feel more satisfying, while holding out hope that you will find a magic solution to the seemingly insoluble problems.)

Finding your team

In the short run, you will be focused on writing a job market paper that is both blindingly insightful and, ideally, single-authored. This is how you prove to the world beyond a shadow of a doubt that you are capable of doing great research.

You should also be aware, however, that in the longer run you will (judging by the statistics) write very few if any single-authored papers. Almost all papers in economics these days are co-authored; research is a team sport. And that means that part of your task early in your career is to find your team.

Finding teammates with whom you work well is important for a few reasons. First, there are real *technical complementarities*, e.g. between someone good at negotiating data access and someone good at using it to estimate models. Second, there are less obvious forms of *complementarity in personality* — for example, one person may be more optimistic and carry the project through hard times, while the other is more skeptical and makes sure problems have been caught before the work gets presented. Third, some people are simply better teammates than others. Governance of research

projects is essentially non-existent, and your co-authors performance (and your own success) will often depend on their personal reliability and integrity. You want to find good people to work with.

You should consider practices you can build into your system to help you search for potential longer-term collaborators. Try working with different classmates on relatively low-risk, early-stage projects to find out how you like working together. Observe what kinds of people you tend to work well with, and pattern match.

Taking care of yourself

Take care of yourself physically and emotionally. These things are both intrinsically important and also contribute to your success as a researcher.

Consistently get enough sleep. As I mentioned above, sleep has well-documented impacts on cognitive performance including functions that are critical to research such as memory and pattern-matching. Sleep also powerfully affects mood, ability to handle stressful situations, etc. Practice good sleep hygiene and find habits that work well for you.

Seek emotional balance in your life. If your research career is the only thing that feels important or meaningful to you, your successes and failures as a researcher will drive extreme emotional highs and lows. Value and invest in other activities and relationships that help to define a more complete and balanced "you," giving you emotional flywheels.

Try therapy. Seeing a therapist might sound uncomfortable or potentially embarrassing, but in fact many people regularly see a therapist and benefit from it, and in some communities it is the norm rather than the exception to regularly do so. I see a therapist at the moment to work on a problem I have waking up in the middle of the night ruminating. You can also learn about and benefit from some forms of therapy even without seeing a therapist. For example, I have picked up a number of useful practices from *Cognitive Behavioral Therapy for Dummies*.

Education Economics Research Metacognition