

SOME THOUGHTS ON DOING RESEARCH

PHILIP H. DYBVIG*
WASHINGTON UNIVERSITY IN SAINT LOUIS

July 27, 2013

*I am grateful to Southwest University of Finance and Economics (SWUFE) for support.

1 Introduction

Recently, I have been spending a lot of time giving advice to young scholars who are starting their careers, mostly doctoral students but also junior faculty, and students thinking about pursuing doctoral work. Since I was saying the same thing many times, I thought it would be efficient to write down my thoughts. Some of my main suggestions are:

1. Work on problems you care about.
2. Have faith but have doubts periodically.
3. Your paper, as well your slides and interview pitch, should derive from two sentences summarizing your main result.
4. Know a lot of things nearby your actual analysis.
5. If a paper isn't working, set it aside.

While these are ideas that have worked for me, your mileage may vary, so you may find that a completely different list of ideas works for you. In any case, you shouldn't apply these or any other ideas too rigidly.

2 Problems You Care About

From time to time I talk with someone who reports working on a problem because it is popular or because "I know they can publish a paper on this topic." Maybe they need one more "A" publication¹ to get tenure and this

¹Most schools have formal or informal lists of how they rate academic journals, often on a similar scale to grades. In some schools, especially when people in different areas

kind of paper is “sure to be published.” However, my casual observation is that the paper won’t get published. If you are not interested in your paper, nobody else will be either. Maybe that topic was interesting to people a couple of years ago but if you are not interested in it now, probably nobody else is interested either. Perhaps more importantly, if you are not interested in the paper, it will be a chore to work on it for three hours a week and you will not do a good job. However, if you are very curious about the topic and you have a burning desire to know what is really happening, it will be a pleasure to work 50 hours a week, and you will do a very good job.

How do you know what you are interested in? When you read a paper, it is your job to read it actively. This includes thinking about alternative explanations, whether you believe the main arguments, your assessment of the quality of the workmanship, and in short, your opinion about what you like and don’t like about the paper. It is not going too far to put pencil to paper and check the model under alternative assumptions, or scour the tables (or even look at some data on your own) to look for patterns not discussed by the authors. If you do this for a lot of papers, you will have a good idea about what topics you like and do not like. Also, this can be a direct source of problems to work on. A lot of good papers have started because the reader likes the topic of a paper but does not like the paper and writes a new paper that fixes what the reader views as problems or shortcomings in the paper.

The same comments about reading a paper actively also apply to going to seminars. You should go to a lot of seminars across disciplines. One view of creativity is finding new and interesting combinations of things you know about, and going to seminars will give you lots of raw material for your mind to work with. Just as when you are reading papers, you should be active in going to seminars. Read the paper beforehand and ask questions at the seminar. Most people are nervous about asking questions at seminars and are afraid they will look stupid. However, it is good practice learning to talk about what’s on your mind, and the nerves will go away with practice. Also, the people who look bad are the ones who just sit there. If you ask questions, like everyone else who asks questions, a lot of them will be simple, but occasionally you will ask a great question (if only by accident) and you

don’t trust each other or don’t trust themselves to evaluate cases, the research evaluation is based entirely on the number and journal ratings of publications.

will get a lot of credit for that. When you start going to seminars, maybe you will understand the first five minutes, and on average how long you can keep up will increase over time, and asking questions will accelerate the process because it will keep you in the game longer and because you will be more engaged.

Doctoral coursework and seminars teach different things. The coursework gives you a good foundation in traditional wisdom, ideas, and techniques in your area of study, while seminars give you a peek at the leading edge and how research looks before it is fully polished. Seminars also give you a lot of information on how to give a seminar and how not to give a seminar (some are so bad that nobody anything after five minutes, and this kind of disaster can be an unforgettable lesson on what not to do). Most PhD students spend too little time in going to seminars and, at least in the thesis stage, much too much time reading articles. Reading articles seems safe, but after a point it is a waste of time. It is also important *not* to read articles about the problem you are starting to work in. Making assumptions for your theory or empirical model is hard work, and if you study the literature first you will fall into the standard assumptions in the literature, and the result will be a footnote on the literature. However, if you set up your own model first and then look at the literature, you will find you like your assumption better in this part and the literature's assumption better on another part you can usually build a really nice model if you combine the assumptions you like best (your preferences again) from both your paper and the literature.

What if you have trouble finding anything you are interested in? This might be a symptom that doing research is not a good fit for you. This is a point that Steve Ross made in a slightly different way in his speech at the first FARFE conference. He said research is hard work, and if you don't like doing it you should probably find a different line of work. Research is a good fit for me, and I have some idea why. When I was a kid, I loved doing puzzles, and now my job is writing down and solving interesting puzzles. Note that this is about preferences, not ability (I like solving chess puzzles even though I am not a great solver). Everyone admitted into PhD programs is smart enough to succeed; if a student fails a prelim² more than once, that is a strong

²A prelim (or preliminary exam or qualifying exam) is a test taken by PhD students covering coursework taken early in the program. Institutions vary greatly, but this is

symptom of not being so interested in the research papers and a poor fit for doing research as a career. People entering PhD programs typically have little or no actual experience with doing research, and it is pretty common for people to discover that they like the idea of doing research much better than they like actually doing research. I hope this is not the case for you, but if it is, I hope you discover this sooner rather than later and make a switch.

When you take the preliminary examination, you should really nail the exam. Take the view that learning all this material is fun. You want to have all the tools and ideas you are studying at your fingertips for years to come. Learning things to the point where you forget it after the exam is worth nothing, and learning it to the point where you have a vague idea but have to look it up every time you use it will waste a lot of time later compared to learning everything backwards and forwards so you never forget it.

3 Faith and Doubts

Although part of your mind knows there is a small probability that a PhD program may not be the right place for you, you need to act and think as if you are very sure it is a good fit. The alternative of continuously thinking you might not belong in the program wastes too much time and can ultimately lead to failure. In other words, instead of being 96% certain you should be a researcher 100% of the time, you should be 100% certain 96% of the time and forget about your doubts and have complete faith while you are doing your coursework and research.

The same principle applies when you are working on a problem. If you are doing theory and working on trying to prove a result (which should be testable and/or having interesting policy implications), you should assume and believe the result is true (this is the faith part, 96% of the time). If you work hard on a proof with all your heart, and all possible proofs seem blocked, then eventually you should think about whether maybe your faith was misplaced (this is the doubt, 4% of the time), and possibly then you

typical.

should decide and believe the result is false and start looking for a counter-example (this is the faith part again, 96% of the time).

Similarly, if you are doing empirical work, you should have a maintained hypothesis (which should be truly based in theory). You should work hard to develop a test that tests the hypothesis and you take a hard-nosed look at whether the data really support your hypothesis, and you should try to find a consistent story for why not only the narrow hypothesis test (which should support your story) but also the other stylized facts and institutions are consistent with your hypothesis. During this phase you should have faith that it will work out. However, if it doesn't work out in a reasonable amount of time, then you should have doubts and think about changing or at least modifying your hypothesis.

These examples are simplified, and need to be modified some do apply to experiments or other types of research, but I think you get the idea.

4 Two Sentences

The job of academics is to conduct and communicate interesting and useful research. For many scholars, especially those who are more quantitative than verbal, the communications part is daunting. The good news is that most of us enjoy writing more and more over time. Writing always requires a significant time commitment, but there is a simple device for making it easier to write a focused paper and making it less likely you will have writer's block. The advice is likely to seem naive or simplistic if you have a lot of experience and training in writing, but then you don't need my advice anyway. The simple advice is to start by writing down two sentences that summarize the most important things you want the reader to take away from the paper. After that, having the two sentences will simplify your work, and it will be much easier to write the Abstract, the Introduction, the Conclusion, slides for a talk, or a summary of the paper for job interviews.

Why two sentences? People are busy, and they do not want to know (and will not have the patience to read or listen to) everything you know about your

topic. It is your job to tell them the most important things. If your paper is cited, maybe someone is writing an extension of your paper and you will be cited again and again throughout, but more often, the citation will be more modest. In fact, most commonly it will be one of a list of papers in an area. A typical nice citation will give a one- or two-sentence summary of what your paper is about, and then discusses how their paper is the same and different. It should be your goal to write these two sentences and explain them clearly. One reason is that if the reader writes the sentences you will not like them, since they will probably be inaccurate and may miss the most important part of your paper. Maybe even more importantly, if you do not tell the reader your two sentences, the reader may not understand what the paper is about at all. Suppose the reader concludes "This paper gives an empirical analysis of capital structure, which seems complicated but the econometrics are well-done." Then, the reader will not remember this paper and will never cite it.

Your two sentences should contain specific results from the paper and should include the most important economic ideas but almost never all the important assumptions. "We formulate and solve a two-period banking model in which rational agents have liquidity shocks," is less useful than "We show that bank runs can arise as a rational phenomenon in a model in which liquidity creation by banks makes them susceptible to runs." It is okay to describe results here without all the qualifications; the reader has been around enough to know it necessary to read the paper to understand exactly what the two sentences mean and under what maintained assumptions or in what dataset the results hold. The reader will, of course, feel cheated if it turns out that your results are not anything like what you described in the two sentences.

How do you use the two sentences? The two sentences are the main idea that you want to convey in a seminar or in a job interview. So, when you are preparing your slides or your pitch for a job interview, use the two sentences to decide what to include. If a slide doesn't clarify the two sentences, cut it (and if you are like most people, you have too many slides and it will be useful to have some guidance about which ones to cut). In a literature review, don't discuss the whole area (like the literature review required in PhD theses in some schools): you should only discuss the papers that help to clarify your two sentences, usually by explaining what you do and don't do in the paper. If a section in the paper doesn't contribute to the two sentences, cut it! Or,

if you think the section is more interesting than your two sentences, revise the two sentences. The abstract will be the two sentences plus a few more details.³

5 Things Nearby your Analysis

It is very important to know a lot about things nearby your results and not just the narrow path through the argument in the paper. You should think about (and ideally derive or test empirically) what happens if you change the assumptions in the model. What happens if you add transaction costs? Could your empirical results be due to managers overestimating their own ability? You assume risk aversion is greater than one; what happens if risk aversion is less than one? The deep understanding you get from this approach is interesting and is likely to suggest an even better way to set up the analysis or exposition of the paper. It will also help you to develop your intuition and to answer questions better. If someone asks what happens to your analysis if there are taxes, maybe you haven't thought about taxes but if you have thought about transaction costs, you might say "I haven't thought about taxes, but I have thought about transaction costs and they are similar because both put a wedge between prices in the market and the effective prices faced by the agent. Transaction costs increase volatility of prices because agents are less eager to trade to exploit mispricing. For taxes, the effect is asymmetric since taxes discourage a large scale of transaction and encourage making the transaction smaller. Therefore, it might increase volatility on the upside and reduce volatility on the downside and therefore increase the average price." Okay, I made this all up and I don't have a particular research setting in mind, but the idea is that if you know a lot of things nearby your analysis you can probably make an educated guess about

³I have another little piece of writing advice, which is don't use "she" to refer to agents in your model. Sharon Oster chewed me out for doing this when I was an assistant professor at Yale. She said it is bad grammar and distracting. Now, I have to agree: what is this little political message doing in the middle of this paper that is supposed to be communicating your research? Also, the political message is ambiguous: do you find it scintillating to fantasize about women being CEOs of large corporations? These days I write my papers in a gender-neutral way that shouldn't upset anyone who likes either traditional grammar or the substitution of female pronouns.

the answer to any question that comes up. This answer might not be correct, but it is reasonable and the listener understands that you are speculating and can learn something.

6 Sunk Costs

Graduate students understand completely the idea of sunk costs, which says that the current circumstances, not how much was spent getting here, is what matters in decision-making. However, in my experience even the most brilliant students have trouble applying this idea in real life. My advice: whether you have spent five minutes or five years on a paper, if it is not working, then you should set it aside. This is similar to the abandonment option we study in finance, only you don't actually have to abandon it and you can come back to it later if you have an idea on how to make it work.

How do I see this? Typically, the student and I agree the paper is not working. I suggest setting it aside and working on something else, and the student says something stupid like "I already spent eight months on that, and I'm not going to throw it away." So, the student spends an additional two years making a paper that's not working into a pretty good paper, instead of spending six months developing a brand new excellent paper. What a waste of time and talent!

I was lucky not to have fallen into this trap when I was a student and assistant professor, mostly because I was young and not very serious about my job. I was just having fun writing down models and solving them. If a model wasn't working or I didn't know how to make a result into a completed paper, it wasn't fun anymore and I set it aside. That was an extremely productive time for me, perhaps largely due to falling into this strategy. I estimate that 60% of the projects I started during that time and got to the point of some sort of draft stayed in the drawer. Not wasting my time on projects that weren't working left me lots of time for the fun projects that were working. And, because I worked mostly on fun projects, it was not a chore to spend lots of hours per week on research.⁴

⁴I also had a tremendous advantage of the best possible advisor (Steve Ross) and

7 Preparing for the Job Market

My subjective impression is that the job market in academia is much more competitive than it used to be. Perhaps that is partly because I had the best advisor ever, Steve Ross, and at that time most of his students went to top places. In any case, the market today is very competitive and good preparation is important. Being better prepared means you can go to a better school, which will make your subsequent career easier because you will have more time for research, better colleagues to work with and give you comments, better support for going to conferences, and a better research environment with better seminars.

The first step to preparation for the market is to do lots of good research. (Okay, this is easier said than done, but this is the goal.) One thing this means is that you shouldn't be satisfied with the first draft of your paper. You need to sit down, probably after a couple of weeks of not looking at it, and look at the paper with critical eyes as if it is someone else's paper. "What is author trying to say here – why not just tell what is going on?" Try to avoid falling in love with your paper as it is now and look for any possible improvements, in the writing as well as the analysis. Present your papers as often as you can (these are in effect practice job talks – see below – and can start early), and get comments from faculty and friends to make the paper as good as you can. Don't settle for a paper that only meets a minimum standard – remember that it is a competitive market. Having several excellent papers (four is a good target), ideally with at least one acceptance or revise-and-resubmit at a top journal, will maximize your chances. And, don't take years and years to put together your portfolio; there is a significant stigma to taking more than five years to finish and you are very unlikely to get a good job if you take seven or eight years to finish. It is easy to get the wrong signals about things like this, since students often look to other students for cues and there is a self-perpetuating "culture of mediocrity" that develops in a lot of PhD programs.

For the actual job search, the first step is preparing the packet to send to schools. Prepare your CV and research summary very carefully and have

getting a young start that is usually unavailable for students coming from abroad.

a writing specialist check the writing (especially if the job search is not in your native tongue). Having very well-written papers (e.g. following the two-sentence idea above) is especially important at this point, since faculty deciding on who to invite for interviews have limited time and if they don't understand quickly what you are doing, they will just reject your application. Also, have a writing specialist look at your papers, with a focus on the title, abstract, intro, and conclusion.

Assuming your packet captures the imagination and attention of a number of schools, the next hurdle is usually the job interview at professional meetings. In the job interview, you want to signal to the interviewers that you would make a good colleague, and this means not only that you are smart and will do good research and will do well as a teacher, but also that you would be fun to have around, that you respect them and the research they do, that you might write papers with them, and that it would be a pleasure going to dinner with you. Importantly, the interviewers are also concerned about whether your papers are mostly your own work or whether your advisor directed you every step of the way. How are you to handle this complicated signalling problem, especially given that this kind of interview will tend to make you nervous?

The best way to prepare for the real interviews is to do practice interviews. Fifteen is not too many, maybe scheduled in clusters of five spaced weekly so you have time to update your pitch in between; if you are a finance student in a small department, maybe you will have to have practice interviews with some people in economics or some people in accounting or management to get your numbers up. Pretend you are not nervous in the practice interviews, and as you get more experience you will actually be less and less nervous (a little nervous is actually good because it will keep you alert). The interviews will alert you to likely misunderstandings from your pitch and will tell you what intuitions people find meaningful. Also, you will hear a lot of the same questions as you will hear in the actual interviews as well as get practice answering various unusual questions. As mentioned earlier, knowing a lot of things nearby your analysis will help you to answer the unusual questions. Unusual questions are often intended to see how you think and whether you have your own ideas beyond your advisor; being well-prepared will make you look good.

Assuming the interview goes well, most schools (at least in the US) will invite you for a recruiting visit with a job talk. Preparations for the interviews at the meetings will prepare you pretty well for the office visits on campus, but it wouldn't hurt to have a couple of practice interviews from the faculty at your school. More importantly, you have to practice your job talk many times, especially if you will not be presenting in your native language. Each time, get comments on your slides and other specific comments. A lot of the time, you will know which parts are working and which aren't, but the audience members will see more than you do. As you practice the job talk, you will get a clearer and clearer idea about your paper's contribution and what are the main results people will be interested in. This is a good point to go back and reread the introduction: you will probably be surprised at how far the introduction is from your new understanding of the paper and how easy it is to improve based on your experience in the seminars.

8 Conclusion

I hope that this core dump of some of my ideas about research strategy are useful for you.