

Ph.D. Thesis Research: Where do I Start?

Notes by
[Don Davis](#)
Columbia University

If you are the next Paul Samuelson and will wholly transform the field of economics, pay no heed. If you are the next Ken Arrow and will invent a new branch of economics, these notes are not for you. The aim here is more humble: to provide strategies for identifying exciting thesis research topics for the rest of us.

There is no algorithm that yields an exciting thesis. Too much depends on your energy and imagination. But there are more and less efficient ways of trying to identify exciting topics. And I will try to convey at least my own aesthetics about what interesting research is about. These may vary a bit by sub-field and certainly across economists, so certainly seek out others' perspectives. So, ignore these suggestions if you choose – but have a good reason to do so.

How Do I Find “The Right Topic”?

First, there is no “Right Topic.” What is hot today may be ice cold by the time that you go on the job market. You don't want the nineteenth best paper of the year on a hot topic.

Much more important is to find something that is important and genuinely interests you. There are great papers to be written in almost all fields. You need to settle on an area where you are sufficiently interested that you don't mind making some investments, since these investments are preparing you not only for thesis work but also for your next round of papers as an assistant professor.

Let me underscore that you should focus on an important problem. The lore of economics includes what is sometimes termed “Summers' Law” (yes, after Larry). This holds that it takes just as much time to write an unimportant paper as an important one. Hence . . . you might as well work on *important* topics. (Note: This is not an incitement to work on broad, vague topics!).

This is not as trivially obvious as it might first appear. Lets start with a first fact: Most of economics is boring. No, I don't mean this in the way that the public at large means it; on the contrary, I think that economics done well can be beautiful and fascinating. What I mean is that most writing on economics is boring because: (1) It does not address interesting questions; (2) It has nothing new to add that is itself important; or (3) Even if the researcher does in fact have something new and important to say, the researcher does such a poor job of articulating this that the reader has little chance of figuring this out.

How do I know if I have an interesting topic?

First, be aware that “interesting” inevitably has a subjective, aesthetic component. So we cannot expect to find necessary and sufficient conditions for an interesting topic. Nonetheless, there are useful indicators.

When I undertake a research project, I find it a useful artifice to think of one of the more skeptical members of the profession repeatedly pressing me with the question: “Why should I care?” How am I going to convince this skeptic that she should pay attention to my research?

One part of the answer is that I am asking and answering a question that has some substantive real-world counterpart. Moreover, I would also like to be able to argue that the issue is an important one. Hence, *real-world examples can be influential and magnitudes matter*. In trying to convince yourself (and others), you should be *as concrete as possible* in explaining both the type of problem to which this applies and what the magnitude of the problem is.

Certainly an indicator (not a proof!) that a problem is interesting is that good minds have spent time thinking about it. I note this because most economists will grant the prior that if several leaders in a certain field have struggled with a problem, it is likely to be an important question (i.e. if these people spent most of their time struggling with unimportant problems, they would be unlikely to be leaders in their field!). But you should rely on this only as an indicator. You should be able to tell an independent story about why the area is important. Moreover, working on areas well combed over by the leaders of the field also has a number of pitfalls, discussed more fully below.

Let’s assume now that you have made a convincing case that the problem that you are addressing is one that we do care about – i.e. it is one with a real world counterpart of some significant magnitude. How do you convince your reader that you have something new and important to say about the problem? Let me stop to emphasize both *new* and *important*.

New. In economics, as elsewhere, you are going to be paid by your marginal – not your average – product. Solow’s model of economic growth won him a chair at MIT and the Nobel Prize, but you will be less successful if you write it down again. You may convince us that it addresses an important problem, but there may be nothing in what you have done that is new.

How do you know if what you are doing is new? One answer is to go back and read the entire history of the literature in your particular area. This is a tempting option as you can surely convince yourself and your advisers that you are working hard. Unfortunately it is also a very inefficient path, more likely to mire you in controversies that are old and forgotten for good reason than to show the path forward. The first step should surely be to talk to someone actually working in the area or at least reasonably familiar with it to find out if someone has already answered the question you are pondering (your adviser is

hopefully a good starting point). Second, one can look at recent surveys of the literature or recent working papers directly on the topic as coming close to providing a “sufficient statistic” for what has been done before in the area. This can be extremely useful, but you should at least be aware that even serious academic work often contains “spin” that may tend to understate the accomplishments of older literatures relative to recent (especially those that the author has contributed to). Third, naturally *Econlit* and the *Social Science Citation Index* can be extremely helpful in identifying related work and should be consulted carefully. If you are not familiar with *both* of these, you should stop reading this very instant and return once you have figured out how they are used!

Let’s assume now that you have convinced us that you are working on an important problem with real-world counterparts and that matter in substantive terms, and moreover that your approach to the problem is new. How do you convince us that the work that you will show us is important? We all know of important papers that have launched vast literatures. But much of the resulting literature ends up in third-tier journals if it is published at all. Occasionally a paper, even in a huge literature, rises to the top nonetheless. Why the different outcomes?

The key, I think, is to convince readers that the novel element in your paper is in fact important. How do you do this? The threshold is that after reading your paper, researchers familiar with the literature in your area should see the world differently. How to do this varies to a certain extent based on whether you are writing in theory or empirics. If it’s a theory paper, one element would be if there is a problem that people have understood is important but have not known how to solve. If you can make an advance of this type, that will be very impressive. A second possibility is that there is an outcome that, under reasonable assumptions, people had not thought possible. If you can show that this outcome is indeed possible, then this can be very impressive as well. Note, though, the clause “under reasonable assumptions”! One important element of the problem may be to establish that in fact the type of assumptions that you make are more reasonable than those that the prior literature makes (or at least no less reasonable).

If you are working on an empirical topic, again it is not sufficient to do something that is new. You have to convince us it is important. Taking someone else’s regression model and adding a new variable that turns out to be statistically significant may be okay for an econometrics exercise, but will it land your paper in a top journal? To start, we have to be sure that you have already met our prior questions that the over all question you address is important and that what you are doing is new. For an empirical paper, we must then ask whether there is good theoretical motivation for the inclusion of the new variable. Are we including it in the regression analysis in an appropriate way? In addition to statistical significance, do we also have “economic significance” – i.e. are the magnitudes economically important. In the end, we are faced with the same question as in theory: After reading your study, will the leading researchers in the field be forced to look at the area in a way differently than they did before and in a way that matters substantively. If yes, then you have a nice paper that you should send to a top journal. If not, then maybe you should think again about the value of the project.

A similar set of questions arise if you make a larger departure in your empirical framework. Is there a strong tie between the theory and the empirical framework chosen? Is the approach sufficiently well motivated, both by the theory and the econometrics underlying the specification, that the results of the study are likely to move people's priors about the economic magnitudes at issue? It is true that the profession tends to grant more latitude to researchers who are trying to address an interesting new problem empirically, partly on the idea that follow-on work may help to elaborate the robustness of the framework. But the basic framework must be sufficiently compelling that the results will have some power in influencing people's priors. That is, the results have to be convincing.

Before moving on to more practical matters, let us summarize what has come before. You should choose a topic that is demonstrably important, that has elements which are themselves new and important, and the resulting study should be both reasonable and convincing. One summary test for this is to ask: If the study proceeds well, can I plausibly hope to have it accepted at a first-tier journal (AER, JPE, QJE, etc.)? If the answer is "no," then perhaps you should spend a bit more time identifying a topic for which the answer is "yes." It is an unfortunate fact that even the things you find very compelling may not ultimately convince the rest of the profession that they should be in a top journal. But you will almost certainly fail to get there if you do not ask yourself this question at the outset.

Where do I start? Strategies for Research:

The foregoing has tried to identify markers of good research projects, questions you should be asking yourself as you proceed in your thesis work. But there are also more pragmatic questions about how to identify good research projects, how to spend your time, etc.

There is no unique path to identifying a good research project. Some might find inspiration in Adam Smith. Some might find inspiration in Fred Flintstone. So the following suggestions point to areas where I think the probability mass is concentrated.

If you want to write applied theory, read empirics.

My aesthetics are that the most interesting work in economics must have some real substantial contact with both theory and empirics. The number of internally consistent theoretical economic models that can be written down is unbounded. But which are interesting? Which are papers that you might want to send to a top journal? If you are Gerard Debreu, you may end up writing very abstract models, but the profession as a whole does not have any problem understanding the importance of a consistent statement of conditions for the existence of a competitive equilibrium. For those who are going to do more applied theory, the threshold for it being interesting rises substantially in terms of finding an empirical counterpart. Interesting applied theory is not just looking down the matrix of combinations of possible assumptions to find cells that have not been filled in. Again, the number of these is unbounded. Instead, the key is to find why, having filled in one of those cells, the reader should think that this is an interesting cell to have filled.

Being able to point to empirical facts that would be hard to understand given existing theories is one very important way to convince your reader that your paper is essential, not clutter, and the more important those facts, the more important the contribution of the theory (holding fixed the “wow” factor of the technical contributions).

If you want to write empirics, read theory.

For those who plan to write in empirics, there are several good reasons to steep yourself in theory. The first is simply because you would like to have your empirical work place some *intellectual capital* on the line. What views of the world will we affirm or abandon (strengthen or weaken) on the basis of your empirical work? If you do not have an answer to this, then the empirical work will not be very exciting. Yes, sometimes we just want to estimate an elasticity and we can tell a story about why we care about it. If the approach to estimation has some novel and important element, that can be its own justification. Failing this, the excitement in empirical work is to cast doubt on/rule out some views of the world that people might otherwise have maintained. A second reason for reading theory is simply that the more closely your empirical work is tied to the underlying theory, the more convincing will be the resulting estimates.

There is a “Research Frontier”; Your job is to find it.

Some questions in the field have been answered, or approaches so exhaustively explored that it is nearly impossible to identify topics or questions able to move people’s priors. On the other hand, there is often a set of questions that the leaders in the field are currently struggling with and may be very far from having definitive answers. Being able to weigh in on these problems with a new insight (and avoid dead topics) is an important step. So much of your work is “finding the frontier.”

Go to weekly departmental seminars in your field.

This may be a direct source of ideas for research. After all, the speakers are selected for being leaders in the field and they are presenting their research that is usually at the working paper stage. In addition, it is important to watch how those who have been successful in the field structure their inquiry. Do they convince you that they are dealing with a question that is important, that they have something new and important to contribute to this, and that the contribution they make is reasonable and compelling? Often the answer will be no. It is important to see why this is the case, where they fall short. These will be important lessons as you develop your own research.

Go to seminars of potential new assistant professors at your school.

They are in the position you want to be in within a couple or a few years. Why not go to see which ones fly and which ones dive and to figure out why. In addition, if they happen to be in an area that interests you, they are likely to be very much at the frontier.

Read the working papers of the intellectual leaders in your narrowly focused research area.

This combines two ideas. The first is that within any reasonably-narrowly defined area of economics, there is usually only a small set of people who consistently push forward the frontiers of research. One of your early exercises is to identify this research community and find out what the problems are which they are struggling with currently. (Of course, do be aware that sometimes these leaders may be at seemingly unlikely places!). Of course, the rise of the web makes this vastly simpler than it was only a few years back. Check out their web pages; check the NBER; check the CEPR. Again, the premise is that current work is close to (but not exactly!) a sufficient statistic for what has come before. Take advantage of this.

Read the best journals selectively.

There are a couple of issues here. The first is that material in the journals is inevitably dated. An empirical project may involve conceptualizing the problem, waiting for a grant approval, gathering and cleaning data, getting the software programs up and running, doing first runs, writing a paper, issuing it as a working paper, sending it to journals, getting rejections, doing revisions, submitting a final draft, and waiting for it to finally appear. Thus the paper in the issue that arrived today may reflect the state of thinking five years ago! On the other hand, you should expose yourself to material broader than your own research project, for two key reasons. The first is that there may be unexpected synergies between work in other fields and your own inquiries. Many economists have made a career out of exploring just one or a couple of those synergies. Second, by reading some of the best research and by looking at it with the appropriate questions in mind, you can come to understand concretely what the profession recognizes as outstanding research.

Talk, Talk, Talk! Write, Write, Write!

Interaction with your professors and your fellow students is where a lot of your ideas should come from. Moreover, this is not a passive process. Often it is in the course of trying to articulate something that you think that you understand that you find the weak point in the logic of prior work, which then points you in the direction of something exciting. Trying to articulate things, both orally and in writing, is an important part of the process.

Question Authority!

Economics, or academics more generally, is not a place for reverence! Read what is being written in your field, recognize the contributions that have come in the prior literature, but do not be awed by it. Question everything. Try to state the arguments in your own words. Do you find the arguments convincing? Are there some lapses in the broader claims that are made? Often these will be the paths open for new and interesting papers. While one should respect prior work for having brought the field as far as it has come,

every step forward begins by recognizing the limitations of what has come before. If you look at the prior work too reverently, it will be hard to see these steps forward.

Don't Take Courses!

By the third year of a PhD program, your job is research, not more courses! You can take more courses (of course), but you should have a very good reason for doing so. Acceptable reasons include (a) It is a course that takes you to the frontier of research in an area in which you plan to do research or (b) It develops mathematical or econometric techniques that you plan to use in short order. The reason that I advise not taking courses is that it is a convenient, comforting, and seemingly rationalizable way of avoiding the harder, more frustrating, but necessary conversion from being a consumer of research to being a producer of research. Focus on your primary task – developing your own research program.

Don't teach!

. . . more than you have to. For many, teaching is attached to a stipend or is otherwise economically unavoidable. In this case, do what you must! Moreover, there are some real intellectual and practical advantages from doing a couple of terms of TA work. Explaining the concepts to others is very useful in consolidating them in yourself. But beyond this, the returns become strongly negative. Your job is research – and anything that distracts you from this is a heavy cost. The first cost, which may seem remote at the time that you are deciding on the teaching, is that it could delay completion of the thesis by a year or more. An even larger cost is if it crowds out time to write a really great thesis. As a PhD student, your time is very valuable; treat it that way.

Dealing with advisors.

Advisors want you to succeed. We would love to have Harvard or MIT pursuing all of our students. Engage your advisers with ideas. Do not be afraid to speak up – the risks of saying nothing far outweigh the costs of occasionally saying something stupid (so long as you also occasionally say something interesting!). These contacts can be very important in allowing the adviser to eventually speak of you with confidence at the time you go on the market. Also, don't wait to write a whole paper before running ideas by your professors. They may be able to save you lots of time by asking pointed questions early. You don't have to accept what they say, but have a good reason for ignoring their advice.

Your advisor is too nice!

Believe it or not, your advisors like you! They like you both as a younger colleague and as a human being. And therein lies a big potential problem for you: Your advisor may be too nice! The job market, by contrast, can be cruel. Potential employers, such as professors at other schools, just don't share the same warm, fuzzy feelings for you as your advisors. They are going to pay good money for a product (you) that, for better or worse, in sickness and in health . . . they will have to live with for years to come. One consequence of this asymmetry is that, in spite of their best efforts, advisors may fail to

ask some tough, probing questions about your thesis work that you will not be able to avoid once you are on the market. How do you deal with this? The first is simply to ask your advisors to be as frank and critical as they are able when reviewing your work. Better to have this done by someone who likes you and wants you to succeed than for it to be done by someone who just relishes the opportunity to dissect a job market candidate. Second, diversify. If you can't find more than one or a couple of advisers who think that what you are doing is interesting and important, then perhaps you should think over the topic again.

Present your work whenever possible.

Sign up to present in student seminars. Deadlines help to focus the mind and you learn a lot both about what works and what doesn't by practice. Ask the students on the job market currently: Are their seminars at the end of the market much better than at the outset. Almost inevitably the answer is "yes" and by a large margin. Experience matters.

Consider writing your *first* paper jointly.

One of the biggest obstacles in writing a thesis is getting the first paper written. One way to make this first step easier is to write a joint paper. There are several advantages to this. The first is that it is much harder to become thoroughly stalled on a project that you are working on with someone else. Neither wants to be seen as the sluggard. Second, you are likely to write a better paper together than either separately, simply because you bring different skills. Third, this may give you a good start to having a publication even as you go on the job market. Finally, it is fun. So who do you write with? Writing with another PhD student is one good option. You start out on equal terms, can share all aspects of the project, and can usually devote large chunks of time to it. An alternative is to write one paper with one of your professors. This has some big pluses, but potentially also some minuses. How you figure the balance depends on the particular opportunities you have. One big plus is simply that they have more experience in judging whether a particular line of research is likely to be fruitful, what methods are appropriate, and how to write the paper up in a manner that is appealing to the journals. After all, these are the skills that got them their position in the first place! The biggest plus may simply be the opportunity to see at first hand the choices and decisions that are made at various stages of a research project by someone with a track record for successful research. But there are some potential minuses as well. It is a fact of life that the profession tends to assume that the intellectual heavy lifting in a paper was done by the professor even if the professor stands ready to swear that it was a fully equal project (and even if the reality is that the student may have done a more than equal share!). This is a good reason why you do not want your main job market paper to be joint with a professor (and why it is also best not to have the job market paper be *any* joint paper). But getting the first paper written and possibly accepted at a journal even as you are writing your main job market paper on your own can be a big plus.

Writing matters.

Your job as a researcher is not only to create new knowledge, but also to communicate it effectively. You cannot persuade your reader that you have done something important if they cannot figure out what you did or why even you think it is important. Bad writing often accompanies muddled thinking. State theses clearly and precisely and you may be able to see where the gaps are that need to be filled in. If your topic is boring, even transparent writing cannot rescue it. But leaden prose may lead many readers to give up on a paper that, written more clearly and precisely, they might find pretty interesting. Moreover, especially early in your career, the reader is unlikely to have a strong commitment to slogging through your writings. If you make the task loathsome, the reader will simply stop. Make life easy for your reader. Help her to identify simply and precisely the contributions of your paper.

Presentation matters.

The same lessons hold for seminar presentations – only more so. You should be able to summarize what question you are asking, why it is important, what is new, and what you will do to convince the seminar attendee in no more than a few sentences. If you cannot do this in perfectly intelligible English, then you do not understand your own topic well enough. All other versions of your presentation should be looked on as simple elaborations of this core set of ideas. Why? The profession needs a simple take-away idea from your paper that is memorable. The successive elaborations reflect the fact that in different fora (face-to-face meeting, formal job interview, job market seminar), you will need to take the same message and make it successively richer, more nuanced. This is never more important than when you are on the job market, when you have to speak to a broad range of economists rather than specialists in your own field.

Inspiration is where you find it.

Maybe this is a disclaimer. In the end, there can be no rules for finding a thesis topic, since it can't be mechanical. Much depends on your creativity and inspiration, your insightfulness and energy. A bit of magic is required. If your adviser tells you to stop working on such and such problem and to return to problem X where you were working, they probably know what they are talking about, but then again they may be wrong. You should listen to what they have to say, but be willing to make the substantive judgment that they are wrong. How do you create your own magic? Some people say their best ideas come when they are in the shower, or playing raquetball, or I'm not sure the answer is that I should direct you to take lots of showers! You have to find your own muse. Success and failure, in the end, are in your hands only.

February 2001