Uncovering the Hidden Curriculum
Informal Advice from a Few Faculty
(Professors Black, Dean, Gouin-Bonenfant, Kopczuk, and Montiel Olea)
Based on a Meeting of Faculty and Graduate Students
Fall 2020

As a reminder, here is a blurb on what "the hidden curriculum" means:

> *A pervasive concern among graduate students is that there is little transparency or explicit guidance about individual things that would help one succeed in academia. For example: How does one acquire an advisor, or find RA work? When is it best to work as a TA as opposed to work as an RA? How does one get involved in writing referee reports for journals? Which conferences are best to attend? Answers to these questions often rely on students having informal networks among themselves, which can encourage structural disparities in knowledge.*

In this document, we want to make things a little more transparent and equalize the playing field.

<u>What constitutes a dissertation?</u>

- There isn't really a firm answer, and there is a lot of variation in terms of what constitutes a dissertation. In general, it will consist of three papers, but that's not required and it depends on your field and advisor (e.g., an economic historian may consider writing a book instead!)
- The three papers do not have to be particularly related (discuss with your committee to decide what's appropriate)—the typical title of a dissertation is something like "Essays in Macroeconomics."
- These do not need to be seen as three equally-weighted, equal quality papers. For example, your job market paper should be very polished, your second paper is ideally nice and solid, and your third paper can be, e.g., a good idea that is still in progress or has clear limitations. Sometimes you will have three high quality publishable papers, sometimes you will have one high quality paper and two things that were done earlier in the process and are acceptable enough that your committee feels comfortable.
- An example of a reasonable combination is your 2nd year paper, 3rd year paper, and job market paper.
- In practical terms, **the key thing is the job market paper.** The number one thing to make sure you do is to have a really good job market paper.
- However, if you only have your job market paper, people might question your productivity ("how did this person only have one idea"?). There are marginal benefits to having more papers, but the marginal benefit drops off relatively quickly. Once you've gotten to three papers, you've shown you're a productive student and you have more than one idea, so once you've hit three papers it's more about maximizing total quality. That can be done either by adding more papers or developing those three ideas more. In sum, fewer than two papers looks slightly weird, more than three is not necessary.

- If you manage to ultimately publish the three chapters of your dissertation, that can help you get going as an assistant professor. The papers can be of different quality, but if they are all publishable that will help make the beginning of your career less stressful. (Note, however, that many people do not publish all three chapters of their dissertation.)
- It is useful to explore a lot of different topics during graduate school—it might take time to narrow in on a particular topic, and it is a great time to explore new areas.
- In terms of co-authoring dissertation chapters, there are official guidelines saying how much of a dissertation can be co-authored. Learn more here.

What does an advisor do?

- Help you get ideas/someone you can bounce ideas off
- Give you guidance on what ideas can become a good paper vs what ideas are interesting but difficult to turn into papers
- Help you decide when to drop or continue with a research topic and guides you through the literature
- Help you get referee tasks
- Someone you co-author with
- Supports you when you're looking for a job and recommends you to jobs outside the department (also note that your advisor should be the first person to know when you're planning to go on the job market!)
- Provides support of your work inside the department
- In general, professors have a few more years of experience in the profession than you, so they've gone through these processes and want to pass on their insights to you. They won't necessarily always be right, but they have spent more time navigating the profession.

How to choose an advisor?

- There isn't a "best advisor," there's a "best advisor for you." Every advisor has a different style and you need to find one who's a good match for your personality and interests.
- Finding an advisor is extremely important and worth investing your time in. Your relationship with faculty is super important to your enjoyment and success in the program.
- That being said, you are not wedded to your advisor permanently—you can talk to different people each year as your interests change and topics evolve. (And, in fact, you are *strongly* encouraged to talk to different people!)
- You need to be proactive and go talk to people. Talk to as many faculty as you can. Keep in mind that advising is an integral part of a faculty member's job—not just their own students but all students--and they know this, so they expect you to come to office hours and schedule appointments. Force yourself to go talk to people, and you'll find who's a good fit for you.
- You do **NOT** need an idea to initiate a conversation with faculty! If you wait until you think you have a good idea, you may be waiting forever and never do it. For initial meetings, you can talk about broader areas of interest. Professors are happy to bounce ideas with you, even if they're not in your field.

- Office hours are often the easiest way to start a conversation. Just show up and say something like “I’m in your class and I saw you do research on education and that’s something I’m interested in.” You should also feel comfortable emailing faculty saying something like “I’m interested in labor economics and I’d love to come talk to you about what the interesting areas of research in labor are.”
- The professors teaching your second-year classes are natural people to start talking to (especially since these classes often have paper requirements, so you can discuss your nascent research ideas).
- Even if the professor you’re talking to isn’t doing exactly what you’re interested in, they can point you to related literature as well as help to identify someone in the department or elsewhere who is.
- Personality and close research interests are important for a good match. People who have the best relationships with their advisor are those who become close colleagues with them. Part of that comes from the flexibility of the advisor, and part comes from flexibility on your end. You don’t want to sacrifice your ideas entirely, but your advisor will be able to help you more if your interests align with theirs.
- You should be talking to lots of people. You will have one or two main advisors in the end, but one person will probably not meet all your needs on every dimension. For example, you may have one advisor who is encouraging and positive, one who is really technically good, and one who is good at big picture thinking but maybe doesn’t get into the weeds. You want diversity in who you’re talking to. You should also talk to both junior AND senior faculty, and, ideally, have a senior faculty member as one of your primary advisors.
- Further, you want experts in whatever you’re working on looking at your work (since there will be experts looking at it when you’re on the job market). It’s very unlikely your advisor will be an expert on all the aspects of your paper, so make sure you’re talking to the experts you need for that paper.
- Remember that faculty do not take offense if you decide not to choose them for your advisor. They are still happy to talk to you about your work.

How often to meet with your advisor?

- This depends on you and your advisor, and is something you should discuss with them. How often you meet depends on the personal preferences of you and your advisor.
- Some advisors will want to meet frequently, but other advisors aren’t going to manage you, so you have to be direct to make meetings happen. For example, if you want to meet with your advisor every two weeks, then you need to say that.
- Keep in mind that advisors dislike it if it feels like the student is not listening to them or if the student just disappears. It’s fine if you haven’t done what your advisor asked, but just make sure to discuss it with them. For example, “Actually I don’t think that’s the right direction for the paper…” or “I tried that model and it didn’t work out…” It’s annoying for an advisor to give you a ton of advice and then feel like you didn’t listen to any of it.
- Different faculty have different preferred ways of working, for example whether they like you to submit work in advance of meetings. Finding a faculty member whose working style is compatible with your own is an important part of the matching process. However,

once you have agreed on a process with your advisor, it is a good idea for you to both stick to it as much as possible

How do you become an RA?

- First of all, keep in mind that there's huge heterogeneity across fields in terms of how much RAing is needed. For example, empirical work tends to need more RAs than theory work.
- Go to faculty who do research in the areas you're interested in. In addition to talking about your own work, mention that you're looking for an RA position—"I would love to RA if you have something."
- The more a faculty member gets to know you, the more likely they are to hire you. One way to do this is to go to office hours often.
- In terms of timing, just ask again every so often. It doesn't hurt to send an email or stop by to say, "I just wondered if anything opened up."
- Also keep in mind that different levels of RAing require hugely different levels of commitment from faculty. For example, doing 5 hours a week or working for the summer is low cost for faculty, while being an RF requires buying out all of your time and is thus difficult for faculty to pull off.
- It's worth thinking about the skills that faculty members need. It's easier for faculty to hire you if you already have the skills they need.

How to become a TA for the class I want?

- In general, there's a matching process for TA jobs. This process takes into account both the survey students fill out in the spring and direct requests from faculty and students.
- You can approach faculty to request working with them, and they can make requests to have you as a TA (though even if they put in such a request you won't always get it).
- If you want to TA a PhD class or math camp, Amy now sends out a call for interest and you can respond to that.
-

How to find out about summer schools and reading groups?

- Summer school announcements start to come in February and March, and Amy posts them on the "Resources" page on the website.
- Faculty also get emails about summer schools that they can circulate to their students, so be on the lookout for that.
- Reading groups are very ad hoc; as of now, they require hearing from the faculty organizing them. More advanced students in the field might also have information about these.
- For students interested in summer reading groups or colloquia, it may be best to reach out to student colloquium organizers at the end of the spring semester, because they will typically know about these meetings or if anything is being planned.

What happens are the job market placement meeting?

- Students are discussed, and this discussion is less about specific rankings of “who’s the best/worst student” and more about coordinating who to pitch for which jobs according to where each student fits best. There’s not a list with rankings and quotas for schools—it’s more about where the student fits and getting everybody on the same page. If there are 30 students that fit top 5 schools in a particular year, then the faculty will try to place 30 students in the top 5 schools.
- When you go on the market, you should have a conversation with your advisor about reasonable expectations about your placement and what would be a good fit for you. Try to be as broad and open as possible.
- Then faculty meet and discuss so that everyone is on the same page about which types of jobs to pitch which students to.

How important is the internal job market talk?

- Think of this like a signaling game with productive effort.
- On the one hand, if you’ve spoken a lot with faculty and your advisor, even if your talk doesn’t go well, they may know that you’re better than your performance in that one seminar, and your advisor can advocate for you. Your advisor should not be learning anything new about you at your internal job talk. They will also be your advocate in meetings.
- On the other hand, everyone knows your advisor is advocating for you, so the other faculty are looking for some verifiable information that confirms that what your advisor says about you is true, and the job talk serves as that verifiable information. At your job market talk there is a broader group of faculty who will see your work, and this is occurring at a time when you can’t make dramatic changes to your paper, so the talk is considered a strong signal about where you stand. It is harder for your advisor to push for you if things went badly in your job talk.
- On the whole, the job talk is important: you want to enable your advisor to say the best possible things about you in a way that people will believe.
- The practice job talk is also very useful practice for the real thing!

Talking to your advisor about non-academic jobs

- Do **NOT** think that there’s a strict ranking saying that you have succeeded if you got an academic job and you failed if you didn’t. Your advisors don’t think like that and neither should you. It is absolutely fine to not want to go on the academic market.
- Additionally, don’t assume that you can just decide last minute that you want a non-academic job and suddenly all these firms will line up to get you. You need to prepare properly for the non-academic market as well, so you should discuss this with your advisor early on so that they can help you. They will not drop you as a student just because you don’t want to go into academia!
- Sometimes you don’t know whether or not you want to pursue academia. Be open to different types of jobs since you don’t know everything that’s out there.
- Think about the PhD as 5 or 6 years where you can learn more about yourself and learn what type of environment you want to be part of. What type of career path might be

better for you, your preferences, and your lifestyle? The goal isn't just an outcome—think about your own long-term happiness and well-being.

At what point should you present a working paper?

- Internally, you should talk to people all the time and as much as you can. The point of being a grad student is to get feedback along the way.
- Externally, talk to your advisor about when/where to present/submit papers.
- Faculty know the different journals and where papers should fit.
- You get one chance to make a first impression when you take a paper out in public, so when you're new, you want to balance getting practice and getting good exposure. You don't want to go too early, but you do need to get some exposure. Decide this with your advisors. You can say something like, "I'm thinking of submitting a draft of this paper to this conference, does that sound good to you?"
- There are also an increasing number of conferences targeted to graduate students, which can be a valuable lower-stakes way to get practice and feedback.

When do you stop working on a paper and submit it?

- Again, ask your advisors before circulating externally or submitting for publication. Just say "is it ready to send off?" or "I'm thinking we should send this off, what do you think?" And "where do you think we should send it?" You are balancing visibility vs. making a good impression.
- Internally, you should always circulate your work.
- When you submit a paper, it either gets desk rejected or sent to referees who are experts in the area. Then referees submit their reports, and the editor decides. You will always get either a rejection or revise and resubmit—papers are rarely accepted outright.

Does the first rejection have any bearing on later submissions?

- You might get the same referees depending on what your area is. This is field-specific and journal-specific.

How do I/should I become a referee?

- To become a referee, it's just who you know (or, really, who knows you). Talk to your advisor and see who they know. It is not common for graduate students to have referee experience.
- The payoff to being a referee is probably close to zero, aside from getting experience and making a good (hopefully) impression on the editor. It doesn't matter for anything in practice.

What are the benefits of professional affiliations?

- While these can be both nice and useful, in general they aren't at all necessary for getting a good job.

- One possible benefit is that you can publish on their working paper series (though much of this is now available via people's websites, or on services such as Arxiv).
- There might also be meetings that you can only attend if you are a member, or networking opportunities.
- One thing that can be valuable is to sign up to mailing lists, for example from the big journals, and various field specific associations (for example in behavioral/experimental/theory there is the Economic Sciences Association and a Decision Theory mailing list). These can let you know what people are working on, inform you about summer schools etc., and occasionally post doc and job opportunities. As usual, ask your advisor which ones would be useful for you.

When to make a website? When to put your work on your website?

- It's never really too early to have a website—you can have a website just saying you're a grad student and you're interested in X fields. Then, once you have projects you're working on, you can post abstracts.
- Talk to faculty about when to post full papers. You generally don't want to post early versions of papers, as you want a certain level of polish before circulating. Most people will only read your paper once and you don't want them to read a bad version.

How do you make the transition to Research?

- Some people are great at coming up with ideas and some aren't, but it is best not to wait around for 'the perfect idea'.
- Do not let the perfect be the enemy of the good: get started on something, and it will often lead you to other ideas.
- If you are struggling to come up with a research idea, there are many ways you can find research inspiration:
    1. Read recent working papers and ask "Can I do this better?" For example, among applied micro researchers, the IZA working paper series is a great start.
    2. Read old papers and think about ways to improve or update it, either through theory or improved/more recent data.
    3. Get started on *an* idea (not *the* idea), something bite-sized. By researching a topic for a small idea, you may stumble into a larger and better idea. Even while writing a 'bad paper' you will acquire skills and ideas to write better papers in the future.
    4. It is possible that you will have an idea but no data, so you can try to think of creative ways to get the data you need.
    5. **Meeting with faculty is crucial, and should be done often**.
    6. There is a lot of learning by doing and heterogeneity in individuals' processes. Guidance from an advisor is huge (if they're at Columbia, they clearly know something about doing good research), so try and leverage their experience and expertise (ask what they think about your ideas). Also, it is not uncommon for JMPs to be ideas inspired by or given to students by advisors.
    7. Don't ever underestimate the power of having fresh eyes; your experience is different from those who have already written on a subject, use that to your advantage.

8. Don't just consume papers and ideas (you shouldn't spend all your time reading papers), you have to produce ideas too.
9. Learn how to structure a paper from more senior people (co-authoring or working as an RA helps, but even looking at well-written papers is useful).

- Advisors are important because they help you distinguish between good and bad ideas; some people get their ideas straight from their advisors, but advisors can be coauthors too.
- Ideas don't grow on trees, they arrive in a stochastic process (you can't decide "Today, I'm going to generate *the* idea").
- Depression and poor mental health are common in grad school, particularly when switching to research because it is easy to get lost during the transition from structured (classes) to unstructured (research) work. To combat this (in part) you need to break the unstructured stuff down into structured pieces. Again, meet with faculty to help you stay motivated and on track.

Is it valuable to think about and push ideas that do not necessarily fit within the standard economics literature (and may not be published in Econ journals)?

- If you want to go on the econ job market, then you want to figure out which econ departments will hire you if you decide to stray from studying economics.
- There is value in thinking about what is happening in science in general, but you need to make contributions to economics to be in economics.

What is the value of working with professors outside of the Econ department at Columbia?

- If they are helping you develop techniques to solve econ problems, that's great. If they are helping you develop techniques to solve problems outside of economics, that's less useful for the econ job market.

Do you find good data then come up with an idea? Or have an idea then find the data?

- It can go both ways depending on your style. You can investigate what's out there and what's possible with existing data, or you can find out about a cool data set (you can apply for funding to buy data) and then think about what can be done with it. It is often an iterative process: you may have an idea → look for relevant data → find the data, explore, and come up with a new question…

How do you balance theory vs policy relevance?

- There is no universal answer, this is where talking to advisor is the way to go.

How can you tell if an idea is good or great (or bad)?

- It depends, and it is hard to tell early on. Talk to your advisor OR someone with more expertise. It's helpful to talk to 'exactly the right person' even if it means going outside Columbia, and that right person can save a lot of time.

- Do not make the best the enemy of the good (don't die wondering, at some point you have to write a paper). Just start and see where it goes.
- How you answer the following questions informs a paper's potential: 1. Is the research question interesting? 2. Is the answer correct? 3. Is the answer new/novel?
- When you pitch your idea (even in an informal setting, i.e. elevator pitch), how do others react? If people react well and are excited/interested, then it's probably a good idea; but if not then maybe reconsider how you're approaching it.

<u>How should we find the right expert to speak to about specific papers?</u>

- If someone has written a on the subject at Columbia, then go talk to them.
- If you don't know, do the hill-climbing approach: ask the wrong person (perhaps your advisor or any faculty member) and they will point you to someone who knows more. Maybe they'll be the right person or maybe they will point you to someone who knows even more than they do.

<u>Should you have multiple research projects?</u>

- It depends on the projects—if you have down time while working on one project, or are waiting for data, then it makes sense to work on something else at the same time.
- However, it is important not to start too many and never finish any. Ask your advisor!!!

<u>Mental Health</u>

- Mental health is super important. You need to develop strategies on how to get through graduate school.
- To-do lists are good.
- Make sure you're talking to people (e.g. your advisor). And make sure not to neglect your mental wellbeing—grad school can be isolating and tough.
- Avoid becoming a 'ghost', someone who is always working but never making progress or getting anything done. It will make things feel insurmountable and awful. A good way to do this is to check in regularly with faculty, who can help you stay motivated and on track.

<u>Are there any tangible strategies for how to add structure each day in the third/fourth years?</u>

- Make sure you have regular interactions with other people so you are not too isolated.
- Set goals that are doable and then do them (break big tasks down into doable chunks or else you will seek out easy to-dos that don't make progress on research).
- Maintain a 9-to-5 schedule or find a regular routine (sleep/wake up at same time each day) and have a fixed amount of time when you focus on work (and really focus, so mute your phone). This does not work for everyone—some people just do other stuff when they don't feel productive (go for a walk, work out, cook, relax).

- When feeling overwhelmed, sit down and think about what needs to be done; break tasks into blocks of importance/urgency (get the important and urgent stuff done first, and put other stuff off or delegate it).
- Exercising is a helpful mental break.
- You need to find your own tricks. Try to have a separation between short-term (the class you have to TA tomorrow or the problem sets you need to grade this week) and long term tasks (your defense, 3rd year paper, etc.) Separate the tasks out and put them in perspective, having long term to-dos gives you a sense of progress and accomplishment.

Quick Wrap-up Comments:

- Talk with your advisor more than you think you should.
- The econ department has resources for purchasing data, so talk to faculty/your advisor about it.