

A Story of the Job Market

Alison Fernandes

asf2142@columbia.edu

This is a brief account of my experiences on the academic job market 2015–2017. The first year I applied to 52 positions (21 tenure-track), and got no tenure-track interviews. The third year I applied to 19 positions, and got 7 tenure-track interviews, one of which became an offer that I subsequently took. I'll discuss some of the things that seemed to make a difference in my own case. However, as British Rail will tell you when announcing train cancellations, everyone's journey will be different.

Documents mentioned as 'attached' can be found at:

<http://www.alisonfernandes.net/resources.html>

First Year

I started applying for jobs in my final (6th) year of the PhD program at Columbia. I had a full draft of my dissertation, minus one chapter. I had written up the material for this chapter as an article instead, which was accepted for publication at the *Australasian Journal of Philosophy* while I was applying (mid October). I had another paper forthcoming in an edited volume (Harvard University Press). (I'd earlier submitted this to two journals, and after revise and resubmits, the paper was rejected. Perhaps I should have persisted with other journals, or had the initial submissions better prepared, but that's another story.) I used these two papers as my writing samples. I had three letters of recommendation from professors at Columbia, and two from externals. I had designed and taught a summer course at Columbia, and had the usual amount of TA-ing experience. I'd read up the job market, and had my job documents looked over by a few committee members and friends. I had 17 presentations listed on my CV, including conference presentations and a few invited talks. Here's the CV (attached).

I applied to 52 positions: 30 postdocs, 21 tenure-track and 1 visiting assistant professor position. Here's the spreadsheet (attached)—the actual spreadsheet I used had more columns and sheets, but this gives you an idea. I made some progress with 5 research

postdocs, but got no interviews for tenure-track positions. Of the postdocs, I was long-listed for one (a junior research fellowship at Cambridge), had interviews for two (an ERC postdoc, and a German 3–5 year position that included teaching). I also received two postdoc offers that didn't require interviews: a Canadian Killam postdoc at UBC and a postdoc at the Center for Philosophy of Science, University of Pittsburgh. I wasn't successful at the interviews, and ultimately took the Pittsburgh postdoc: even though it was only one year, it suited my research and professional interests more than the Killam postdoc, and moving to Pittsburgh didn't seem as daunting as moving to Canada.

I was very glad to have the postdoc offer from Pitt—I could very much see myself enjoying the year ahead. But I was worried about my eventual prospects for landing a permanent job, given that I'd had publications, and hadn't gotten any tenure-track interviews. I was also disheartened by the job advertisements I saw. Many jobs were open positions at fancy places I couldn't see myself being competitive for. Of my specialties, departments seemed to want philosophers of science who applied their research to practice, or metaphysicians who fitted squarely in the standard analytic tradition. I saw my work as situated somewhere between philosophy of science and metaphysics, and including healthy doses of other areas of philosophy (philosophy of physics, decision theory, philosophy of action, German Idealism, pragmatism). If someone was looking for a philosopher of x , I didn't think I'd fit the bill for any x . In retrospect, I don't think that was a significant problem—there are ways of presenting as a multi-faceted candidate that make you more attractive, not less.

Other possible contributors I considered included the fact that there hadn't been a lot of activity from the placement officer that year. No one looked over my letters, for example. Late in the year one of the new placement officers did—they said the letters were good, but a little ABD (i.e. they couldn't present me as basically done with dissertation, as I lacked a completed draft). Perhaps it would have been better to finish the dissertation draft rather than aim for the publication—I'm not sure. I also worried that I shouldn't have used papers on deliberation and philosophy of physics as my writing samples, since they didn't express my positive views, or reflect my core AOS's in metaphysics and philosophy of science.

Second Year

By the time I was applying for jobs the second year, I had finished my PhD (defended March) and was applying from a postdoc position at the University of Pittsburgh. I started using a new writing sample—a paper on the core positive view of my dissertation that I'd reworked in response to feedback from one of the late arrival placement officers the year before ('A Deliberative Approach to Causation'). I had an additional external referee. I added additional talks to my CV, and listed talks separately as invited (8) and conference presentations (15). Other changes to my CV included rejigging my AOSs and AOCs:

2015 AOS Philosophy of Science, Metaphysics, Epistemology

AOC: Philosophy of Physics, History of Philosophy (German Idealism), Ethics

2016 AOS: Philosophy of Science, Metaphysics, Philosophy of Physics

AOC: Epistemology, History of Philosophy (German Idealism), Ethics

Earlier I'd been fearful of claiming philosophy physics as an AOS, given I don't do formal work or work very close to physics. I got over that. Previously I'd been claiming AOCs based partly on what I had TA'd and had general interests in (ethics). I later went with AOSs that more plausibly fit my profile, even though I hadn't taught in them (epistemology).

I rewrote all my job documents in full. I'd met someone who'd had good experiences working with 'The Professor is In'—both the advice blog (now a book), and the editing services that the author provides. While skeptical at first, I decided to give these a go. The book cost \$10 or so, the editing around \$200, I think, for three documents to go through 4–5 drafts, where I received comments that I responded to. I chose to work on my cover letter, research statement, and teaching statement (figuring I could do the CV and dissertation abstract on my own). I did this well before the fall to make sure I'd have room on the editor's calendar. All the documents underwent substantial revision, but the biggest changes were to my teaching statement. With due embarrassment, here are the research statements, teaching statements and cover letters from my first and third year for comparison (attached).

The first part of the editing process involved revising the documents in light of the advice the author gives in her book (she also provides samples to help). I thought I'd already done this, but her response indicated I hadn't—when forced to, I really did try to follow the method she prescribed. I had been skeptical that a general method would apply to philosophy (and my own research). But the product was good. All the new documents I had later met with the placement director's approval, and have served me well since. If you were to ask me one piece of advice for the job market that would apply to everyone, it would be to get your documents in gear, and use this book (if not also the editing service).

This year the placement director looked over all my materials, re-worked the dissertation abstract with me, and advised on all parts of the process. I wasn't at Columbia, so I wasn't able to attend job meetings, but the placement officer made themselves very available regardless.

I didn't make it very far into the job market this second year. I'd applied to 9 postdocs, and planned to apply for another 23 positions. But I was offered one of the postdocs in early October and decided to take it. This was a two-year postdoc on an AHRC project at the University of Warwick, 'Time: Between Metaphysics and Psychology'. While I hadn't worked on psychology (or philosophy of mind) before, my interests in agency and temporal asymmetries made me a good fit for the project. After chatting with a job market mentor who was tenure-track in small-town America, I decided a postdoc like this would suit me better than many other things on offer, and I threw myself into preparing for the interview. I read psychology papers and thought about experiment design. I'm told the interview went well, and I was offered the job. The people at Warwick were willing to delay my start date from January to May (in order to let me finish at Pittsburgh). Given they could be accommodating and the project sounded good, I decided to accept the position and take myself off the market. The happy upshot of all this was that I got to relax and enjoy my remaining time in Pittsburgh.

While I don't have much data from this year, things were already looking a little better than the year before. I was long-listed for two JRF postdocs, and I'd only applied to four so far.

But tenure track positions still didn't look too promising: some were repeat advertisements from last year, and none looked a particularly good fit.

What surprised me about the results this year was that not much had changed, by my lights, from the previous year. I had no new publications, for example (although I had new things under review), and no new teaching experience. What had changed, however, was the fact that I had a PhD, and some better job documents (although I hadn't got a chance to use most of them yet). I was also lucky to have a postdoc advertised that I was a good fit for—nothing like that had come up the previous year.

Third Year

I was now applying from my second postdoc position. A big change was that I'd had an article accepted in *Philosophy and Phenomenological Research*. This paper presented the main positive view from my dissertation, and was the paper I'd started using as my writing sample. A second big change was that I'd had 5 months work on an interdisciplinary project. I'd given conference presentations, and could talk confidently about what interdisciplinary work involved. I'd also thrown myself into lots of new research and presenting: three new papers under review, three new papers presented at conferences, including fancier events (APA) and by invitation. (I gave more than 16 talks in 2017—that strikes me as far too many.) I had another letter of recommendation, which I started using on all my applications. I had just started teaching a Masters level research methods course at Warwick. Even though the workload was minimal, as it was co-taught by four, it still counted. I'd begun to use my reworked job documents, with only minimal changes.

Because I already had a two-year postdoc, I decided to apply for permanent positions only. I applied to 19 before stopping, and had another 8 applications planned—somewhat more than previous years. Of these 19, I received 7 interviews, and was short-listed for another 2 positions. Of the 7 interviews, 3 were final-round interviews that included campus visits. I was offered the first of the positions I interviewed for, at Trinity College Dublin, and subsequently accepted it. (I didn't get the second, and didn't wait to hear back on the third.) I stopped applying to more positions when I got the Trinity offer, which was early in the cycle (mid-November). I kept most of the other applications in play until I got the written

offer from Trinity, which took four weeks—longer than one would expect. I didn't hear back from 6 of the 19 positions before withdrawing my application—so, strictly speaking, I don't know what would have happened with those.

What made a difference this year? I hadn't expected to apply to so many positions, but it was a surprisingly good year for metaphysics and philosophy of science. There were more mid-ranked places looking for specialists in these areas. But it wasn't only this—there were also well-ranked places advertising open positions that also offered me interviews, or short-listed me—places I hadn't taken myself to be competitive for. What I think made a bigger difference was the publication. It was only one more publication than in the first year, but it was in a very good place, and helped tie together who I was as a candidate—I was now the person who did agent-based accounts of scientific relations. I could now present as someone having strong research potential: I had interesting projects on the go that had been well-received within the profession. Fitting into categories like philosopher of science or metaphysician no longer worried me as much. Thanks to my interdisciplinary work and some new projects I was developing (on time travel, and chance), I also ticked a few more interdisciplinary and public-outreach/impact boxes than in previous years. My letters had been revised to reflect some of these changes—one of my letters was rewritten from scratch, and I had one I hadn't used before.

Here's my summary of what I think made a difference between years one and three.

What did matter (somewhat ordered by degree of making a difference):

- Strong publication on positive view
- Strong job documents
- Letters from the people you really need them from, revised to reflect updates
- Interdisciplinary and other eye-catching additions (e.g. time travel)
- More active and coherent research profile
- PhD done and dusted
- Great placement support
- Luck (there were a few more good-fit positions)

What didn't matter (very much, if at all):

- Additional teaching or teaching qualifications
- Fitting in as a 'standard' metaphysician or philosopher of science
- A long list of publications (I had 3 by the end)
- Fancy photos on my website (I never had the time, though I did have a website)

Overall, here are three things I took away from my experience of the job market. Firstly, it sometimes seemed that things weren't going well, particularly when I compared myself to others, or what others (sometimes) expected of me. I wish I'd seen the job market more as a process, and something you can learn to be better at, rather than a direct evaluation of my abilities or potential as a philosopher. Secondly, and relatedly, I think the job market warrants some patience. I was tempted to think that how I presented (and did) in my first year on the job market would make all the difference. It didn't. In my own case, and many others I've heard about, having finished one's PhD (and continuing work) enhances one's chances of gaining a good position, rather than lessens them. For many people, it takes a few years working in temporary positions before securing permanent employment. While not everyone will want to go down that road, don't be tempted to write yourself off because the first years don't go easily. Thirdly, I learnt to think differently about what I had. A postdoc is a great stepping-stone, for example, but I didn't always appreciate that. You'll always find people with more of this or that—be it publications, teaching experience, institutional prestige, whatever. But you'll also meet people with less of each of those things. Work with what you've got.

Bonus Topics

Interviewing and job presentations: Things can go wrong, and do. Try to be so well versed in your presentation, and so confident-sounding in your interview answers, that you act entirely unfazed when these things happen. With you presentations, make sure you're well under the time limit (by at least a couple of minutes), to have time to deal with the unexpected.

Things that happened to me:

- You're presenting over Skype, and there's massive audio feedback, requiring you to talk over the top of yourself every second sentence. Take lots of pauses.
- You're presenting over Skype, and the entire interview team appears as out-of-focus blobs. You can mention it, but there might be nothing done to fix it.
- You walk into a presentation room someone else has set up, and there's no screen in view while presenting—so you need to know the presentation off by heart.

Things that happened to someone else:

- The audio quality so bad you can only pick up every second sentence. Ask for some repetitions, but then do your best with guess work.

Finally, there were a couple of times I knew a relatively junior person on the interview committee. For whatever reason, they presented as some of the most apparently-unfriendly people I ever encountered while interviewing. It's disconcerting when someone you know acts cold, so be prepared for it.

Regarding general preparation for the interview, firstly, make sure you've researched the department really well. This might not include reading the work of people on the panel, but you should be broadly aware of what they do, and the activities of the department. (At one interview someone was really pleased I knew about the undergraduate philosophy club, for example.) Sometimes I read paper abstracts or skimmed papers if the work looked really relevant to what I do (so that I could talk about potential research collaborations), but I never read full papers, or read in areas that I wasn't genuinely interested in.

Secondly, I was advised to have a series of 'platform planks' or significant points that I wanted to get across in my responses. For me, they often included things like interdisciplinary work with psychology, engagement with German Idealism (unusual for an analytic metaphysician), and use of new technologies in teaching—things that make one stand out from others as a candidate. Having these in mind can help you think of appropriate interview answers. (But note, you really do have to listen to the answer and respond to what they ask—rather than say generic things on the topic.) A number of interviews I took ended with questions like 'do you have anything to ask us?', or even, 'is

there anything about yourself you'd like to mention that wasn't covered in the interview?' Final questions like these are excellent opportunities to cover plank points that haven't yet been brought up in the interview. For example, you can ask if they have research support for interdisciplinary work, or if they have interest in the kind of teaching approach you'd like to take. It's good to have specific questions prepared that you run by other people in advance. For jobs that have 'key selection criteria', these are also a guide to what you should cover.

Some of the trickier questions I was asked included 'discuss three philosophers whose careers you have influenced', 'what was the most significant innovation/discovery in your field in the last 15 years?', and 'I see you haven't worked on x . Can you discuss how your work might be relevant to x ?'. Often tricky questions came from non-philosophers on the panel, and required you to think quite broadly about the discipline, and your place in the profession.

Letters of recommendation: I used generic interfolio letters of recommendation for almost all my applications. The exceptions were for Junior Research Fellowship (JRF) at Oxford and Cambridge, where the instructions often specified that recommenders had to upload the letters themselves (typically because there were additional questions the recommender had to answer). Other JRF applications specified that interfolio letters were acceptable. If in doubt, I didn't use interfolio for JRFs. For all tenure-track positions, I used generic interfolio letters. Given the number of jobs I was applying to, I didn't want to be chasing down my letter writers for lots of additional letters. I might have done otherwise if I had a very unusual position I was applying to—I would suggest speaking to your advisor about whether they think it would be important in a particular case.

One comment on UK, some other European, and Australian jobs: often these didn't require letters at the application stage. Letters were only requested for short listed candidates. In these situations, I always used the actual details of my letter writers (rather than an interfolio email address). Sometimes these also specified that one of your referees had to be your current or immediately prior supervisor/employer. I'd feel free to ignore this requirement. It's a feature of generic HR forms, not of academic jobs, and no one seemed to notice or care whether I did or didn't abide by it.

Networking: I enjoy getting to know philosophers and conferencing, and I travel a bit, so I'd say I'm reasonably well networked. I think this makes a significant difference to my professional life in general, but I'd say it didn't have a large direct impact on my application results. For the 3 postdocs I was offered, I didn't personally know anyone involved in the hiring process at all, or really know anyone in the department. I knew one person in the department of the tenure-track job I was offered, but I sincerely doubt it made a difference. I think being networked might have made a difference in being offered some interviews (particularly if I wasn't otherwise a good fit for the job). I never reached out to people in the department to hear more about what they were looking for (although on occasion, people did reach out to me). Provided the people you reach out to aren't on the hiring committee itself, this strikes me as a reasonable thing to do. But I usually felt I had a sufficient idea of what they were after based on public sources of information. Moreover, there's only so much you can do to tailor yourself—particularly at the final stages of applying, it seems the best you can do is put your best self forward.

Best of luck!

Comments on this document are welcome—please email me at asf2142@columbia.edu.

Thanks to Thimo Heisenberg for useful comments and questions.