

Freedom of Research Area

Brian Weatherson

Imagine that a young designer gets hired by Uber. The company has decided that their smartphone logo is terrible, and it needs to be replaced. And rather than using a design agency, they think they can do a better job in house. So they get a designer. And the new employee is, in a sense, really good. They aren't really good at designing logos; in fact they don't make any progress on the logo at all. But they are really good at researching the history of transportation regulation, and writing about this history in a crisp and timely manner. After a while, they spend all their time on this historical research and writing, and the new logo languishes.

I haven't worked for Uber, or any company like much like them, so I can't say for sure what would happen. But my impression is that the new employee would find themselves fired rather promptly. It wouldn't matter how good their work was. In fact, I stipulated that it was good. And it wouldn't matter how important that work was to the overall mission of the company. I don't know Uber's inner workings, but from the outside I suspect that anything that can help them deal with regulatory challenges is considerably more important to their long-term profitability than the look of their smartphone icon. None of this matters. The employee was hired to do a job, they were conspicuously not doing it, and in many businesses, that will mean you get fired.

Academia is, in crucial respects, not like that. Indeed, it is a key aspect of academic freedom that researchers get a rather large degree of freedom in choosing what they want to work on. A department may conduct a search in a very specialised area, and hire someone on the strength of their work in just that area, but if that person conducts research in a somewhat different area when they arrive at the job, there is little the department can do about it. And, some say, this is how things should be. Anything otherwise, they say, would be an unacceptable restriction on academic freedom. Here, for instance, is how Cary Nelson puts the point,

Academic freedom gives both students and faculty the right to study and do research on the topics they choose and to draw what conclusions they find consistent with their research, though it does not prevent others from judging whether their work is valuable and their conclusions sound. (Nelson, 2010)

The idea that academics are not constrained in their research topics has some history. Here is Alexander Bickel, describing an ideal university.

In universities, professionals of many disciplines can follow lines of inquiry determined by themselves, individually and collectively, and dictated by no one else, on grounds either ideological or practical. (Bickel, 1975, 127)

[†] Penultimate draft only. Please cite published version if possible. Final version forthcoming in a volume on academic freedom, edited by Jennifer Lackey

And we see something similar in some university regulations. Here, for example, is what the University of Chicago has to say about research topics.

The basic policies of The University of Chicago include complete freedom of research and the unrestricted dissemination of information.¹

There is a principle that seems to be running through these quotes, and that I mean to focus on here. I'll call it FRA.

Freedom of Research Area (FRA) If part of an academic's job involves doing research, then the academic themselves gets to choose which areas they shall perform that research in. And provided solely that the quality of the work is sufficiently high, this research in their self-chosen area shall count as adequately discharging their duties to their academic employer, at least as they pertain to research.

I think FRA is false, or at least that it should be false. That is, I think academics should not have complete free choice of what they research on. I'm not really sure how many people think FRA is true, though I probably some people think something like it is true. Perhaps more importantly, I think the appropriate qualifications that one needs to add to FRA to make it true are neither obvious, nor reflected in practice. And that's why I think it's worth discussing.

In particular, while FRA does not seem to me to be reflected in the regular practice of academic life, a related principle I'll call FRAD is. That is, I think FRAD is both widely believed, and many people act as if it is true.

Freedom of Research Area within Department (FRAD) If part of an academic's job involves doing research, then the academic themselves gets to choose which areas they shall perform that research in, provided it is within the disciplinary boundaries of their home department. And provided solely that the quality of the work is sufficiently high, and that they work within these disciplinary boundaries, this research in their self-chosen area shall count as adequately discharging their duties to their academic employer, at least as they pertain to research.

FRAD, I'll argue, is also false. Something close to FRAD, however, is true. Academics should have 'elbow room' in their research; they should be able to move from one research project to nearby projects. And if they make such a move, they should count as having fulfilled their research responsibilities provided their research is of a high enough quality. FRAD is similar to the elbow room thesis, but not quite the same as it. And the differences matter in some important cases.

In focussing on FRA and FRAD, I'm setting to one side most of the questions usually thought central to debates about academic freedom. (Though I trust these questions will get plenty of discussion in the rest of this volume.) The focus here is on which questions academics ask, not on what answers they give. Questions about how free academics should be in answering questions (E.g., Is it ok to defend Pol Pot? Is it ok to use seances to motivate

¹Retrieved from <https://provost.uchicago.edu/handbook/research/research-policies>. Both this quote and the Bickel quote are cited by Richard A. Shweder (2015).

historical interpretations?) are left for others to address. And I'm exclusively focussing on what academics do in the conduct of their work. Questions about whether they should be subject to professional sanction for research activities outside work, and, assuming they are immune from such sanction, whether this immunity ought be related to their status as academics or simply to their status as employees are also being set aside.

So we're focussed on what questions academics ask in the course of their work as academics. There is one more distinction to make to really focus the discussion. Academics have, to greater or lesser extents, both the freedom to tackle different research topics, and the responsibility to tackle certain topics. I'm interested in the responsibility side. What kind of research counts as suitably discharging one's professional responsibility to research? Put more bluntly, the focus here is not on what research questions an academic may ask, but on what questions they must ask.

The American academy has a rather odd structure when it comes to enforcing this responsibility. Junior academics get reviewed after roughly six years, and if their performance is satisfactory, they are awarded tenure. If not, they are fired. Just what counts for tenure varies a lot between institutions, but at research institutions, whether one has adequately discharged one's research responsibilities is a huge part of the equation.

It's not completely true that there is no other point in the American academic's career where there will be an inquiry into how well they are discharging their research responsibilities. But at no other point are the stakes nearly as high. For instance, many departments have a small pot of money to distribute in the form of annual raises each year, and often enough research performance is a factor in that distribution. But the sums involved, especially in cash-strapped times, are tiny. Since it is very rare for one's nominal salary to fall in this process, and inflation is so low, the worst that happens if one completely fails to discharge all research responsibilities is that one's salary falls by a percent or so per year. That can add up over time, but compared to being fired, it's a minor penalty.

So I'm going to focus mostly on that tenure decision here. That isn't because I have any sympathy for the current structure, with the stakes being so high here and so low elsewhere. But it's what we have to work with, so it's what is relevant here and now.²

So imagine the following case. A young scholar gets hired in a US university³ on a tenure-track line. Six years later, they are up for tenure review. And in the interim they have done high quality work, with the quality and quantity of the work being sufficient for promotion to tenure. But the work is in a different area to the work they did before being hired, and this work is not at all what the department had in mind when they were hired. Assuming their promotion file is adequate in other respects (especially concerning teaching and service), should they be promoted to tenure? Or, perhaps more precisely, what further details of the case matter to whether they should be promoted to tenure? Should it matter, for instance, whether the work was inside or outside the disciplinary purview of the department?

²Just to be clear, I'm not denying that the arguments given below could apply to tenured academics just as easily as non-tenured ones. But doing so would be a very radical break from current practice, and motivating such a radical break would need much more careful discussion than I have space to do here.

³I will use the term 'university' here for any post-secondary educational institution that hires professors with an expectation they will produce some research. Many of these institutions have 'college' rather than 'university' in their name, but I'm calling them all universities.

I'm assuming here that principles of academic freedom apply at all to pre-tenured faculty. This doesn't seem too controversial, though it is striking how some universities talk about tenure and academic freedom. Here, for instance, is a passage from the University of Michigan's Tenure Guidelines:

The University safeguards academic freedom through its policy that no person who has been awarded tenure by the Regents or who has been employed by the University for a total of ten years at the rank of a full-time instructor or higher may, thereafter, be dismissed, demoted or recommended for terminal appointment without adequate cause and an opportunity for a review... (The University of Michigan, 2016a)

It isn't hard to read that as saying that it is through tenure that academic freedom is protected, and conclude from that that faculty without tenure don't have academic freedom. But I'm assuming that conclusion is false; academic freedom does extend to untenured faculty. And the question is what it covers.

The particular puzzle case I'm interested in is not unique to philosophy, but philosophy is considerably more prone to it than other fields. In many fields, a central part of the tenure file consists of the book that results from the dissertation. In such cases, there is little danger that the tenure file will look radically different from the research profile that was submitted in the candidate's original job application. In many other fields, research is closely connected to getting and spending grant money. And the mechanics of grants make it hard for someone's research to take a sharp change of direction at a very early stage of their career. This is not to say the case I'm interested in cannot arise in such disciplines. But it is much more likely to arise in disciplines that are neither grant-based nor book-based. There are few such disciplines in existence right now, but philosophy, at least in the US, is one. Philosophy is also the discipline I know the most about and, to be honest, care the most about, so it doesn't bother me that I'm writing about a problem that is more common here than elsewhere.

I'm also going to write exclusively about fictional cases. I looked into using some real life cases to make the discussion more vivid. But they ended up being more of a distraction than a helpful illustration. In particular, it was hard to find a case where a candidate for tenure was uncontroversially doing high quality work, but there were concerns about the area it was in. Rather than re-litigating the tenure files of these human beings, I think it most appropriate to focus here on the abstract case.⁴

To be sure, it is hard to precisely imagine a case just like the one I am describing. Doing high level research in any field is hard. If there is no sign of one having worked on something before being hired, the probability that one will be able to acquire sufficient knowledge and skills to do top quality research in that field is not high. And perhaps it will even be hard to get unbiased reports on the quality of the work, if the candidate did not get into the field through the usual channels. But it's not so unrealistic as to be unimaginable.

⁴As I'll return to below, the main way the issues I'm discussing impact everyday academic life is that expectations of how tenure reviews will be conducted affect how pre-tenured academics structure their research profiles. As any game theorist knows, the nature of non-equilibrium outcomes can be profoundly important to the actual world, even if they are never reached.

I hope everyone would agree that doing extra work, well away from what one was hired to do, is not a bad thing. It shouldn't count against the candidate for tenure. Indeed, it is good to show some ability to stretch out. When I say FRA and FRAD are false, I very much do not mean that one is obliged to not do any other kind of research.⁵ The more research the better! But first things first.

So why do I think FRA is false? And why do I think it would be reasonable for universities to insist that their candidates for tenure do some research on what they were hired to do.⁶ Well, let's start by looking at an extreme case that might be thought to motivate something like FRA. The following would be unreasonable behaviour on the part of a hiring department. A candidate is hired on the basis of an excellent dissertation on peer disagreement, which has led to two publications in good journals, and there are two more papers under review from the dissertation. The hiring department expects her to keep doing just this kind of work.

But soon after she arrives at her new job, she surveys the most recent work on peer disagreement and decides the debate is dead. There is, she thinks, nothing more to say about this debate. It is an ex-debate, it has ceased to be, it is no more, it has shuffled off this mortal coil and is now pushing up the roofs of the libraries.⁷ So rather than scream into the void, she decides to take what she has learned in debates about disagreement and apply them to more vibrant debates about testimony, and about judgment aggregation. And between getting hired and coming up for tenure, she writes a series of high quality, widely cited, papers on these topics in respected journals.

But then the hiring department gets upset at time for tenure review. We hired you to work on peer disagreement, they say, and what have we here? Nothing at all on peer disagreement, but all this stuff on these distinct, though admittedly related, fields. That's not enough, we say, for promotion to tenure.

This is poor behaviour on the part of the hiring department, and so unreasonable that I find anything like this happening in a real department almost inconceivable. (Though some departments do have an impressively dogged commitment to unreasonableness, so perhaps I should be careful here about the link between conceivability and possibility.) But we can say more than just why it is unreasonable. There is a good story about what makes it unreasonable.

Academic debates die. Everything that needs to be said is said, and it's time to start talking about something new. When that happens, it is wrong to keep plowing these barren fields. And the people best positioned to spot the death of a debate are experts, with dedicated knowledge. Outsiders may suspect that if nothing new is happening, the participants are just tired following a prolonged squawk. Or, perhaps more likely, outsiders will confuse mere squawking for actual progress. The people best positioned to determine whether it is worth investing more resources in a debate are participants to it. If our imagined candidate has decided that the debate is dead, then it will usually be reasonable to defer to her

⁵I wasn't hired to write articles like this one, but there isn't, I hope, anything wrong with my writing it.

⁶There is an important caveat here. It would be completely unreasonable for a university to decide, just as a candidate comes up for tenure, that they really wanted the candidate to have been doing different research for the past six years. The question is what behaviour on the part of universities would be reasonable if it were clearly communicated well in advance.

⁷Full disclosure: I have a book manuscript under review with a long discussion of peer disagreement.

expert judgment. I'm assuming here that when it comes to particular areas of debate, an assistant professor will be an expert, even compared to her senior colleagues. She will know, and they won't know, the details of what has been happening in the very recent literature, and how much those details matter. Sometimes that won't be true; her colleagues will be experts. But it will be usually true, and so it is safe to assume it is true when considering hypothetical cases for the purposes of policy development.

One other relevant fact about academic debates is that they are not isolated. An expert on one debate won't automatically become an expert on all related questions, but she won't be a novice either. In the peer disagreement example, it is very natural to think that our expert will know a lot about testimony and about judgment aggregation. Those debates are both highly relevant to disagreement. So it is reasonable to expect that if our candidate slid into those debates, she would produce excellent work. And, recall, that is exactly what happened in the example.

Putting these two thoughts together, we get the following conclusions. Allowing people to drift between nearby areas of research will not, on the whole, reduce the quality of their research. And allowing the people who are experts in a particular debate to choose when to move between nearby fields, we can leverage their knowledge of those fields to ensure that their work remains relevant to lively debates.

These considerations support a freedom to drift, to move from one area of research to adjacent areas without needing approval from a central authority. And that's already a kind of academic freedom. The motivation here has a family resemblance to Hayek's argument that a virtue of markets is that they provide a way for the system to leverage the expertise that market participants typically have, at least about areas immediately relevant to them (Hayek, 1945). And this Hayekian flavour to the argument shouldn't be surprising. The alternative to a model where academics have some freedom to choose the direction of their research is one where a central planner chooses everyone's research topic for them. And arguing against the success of such central planning models was a central concern of Hayek's throughout his career.

So academics, even junior ones, should have elbow room (to borrow a metaphor from Daniel Dennett). But it's a long way from this to endorsing FRA. Indeed, the considerations that supported a freedom to drift could not possibly support FRA. For one thing, the fact that one was good enough to be hired in one particular sub-field does not indicate that one will have particularly expert judgment on whether it is a good use of university resources to have (more) research conducted on a particular field distant from one's own. And for another thing, the fact that one was hired in one field is little to no evidence that one would be the right person to conduct that research, even if it were in the university's interests. So if there is a wide ranging freedom to research on whatever takes one's fancy, it will need radically different justification to this.

It is hard to see what that justification could possibly be. There are people who are given awards that are meant to support any kind of research that they find interesting. The MacArthur Fellows Program, the so-called 'genius grant', is like this. And it seems suitable for people who have made spectacular contributions, and will likely continue to do so. It seems particularly suitable for people who have already shown an ability to create great

works that require leaping between seemingly distant fields. If, for example, you can use hip-hop to turn the story of the most elitist of the founding fathers into a popular phenomenon, then someone should probably give you untied funding to just see what happens next.

But a junior professorship is not a MacArthur Fellowship. Indeed, it is dangerous to think that it is, or that it should be. It encourages the idea that universities should be looking to hire geniuses, rather than hiring people who have put in the hard work to get to where they are in their field, and are likely to keep getting further results by a continued application of just that kind of hard work. At least some of the time, perhaps most of the time, the question of just how smart the candidate is should be considerably less relevant to a job search than the question of what they have achieved, and what those achievements signal about their likely future research contributions.⁸ Yet if everyone who was hired was been given a free rein to work on anything whatsoever, if every hire was the equivalent to bestowing a MacArthur Fellowship, then whether the candidate was some kind of genius would be a central, perhaps sole, criteria.

If we were going to say any academic should work on whatever they like, or even whatever they like in their department's research purview, we need to do one of two things:

1. Show how this freedom is consistent with the idea that departments can, in hiring, take area of research into account; or
2. Argue that the very widespread practice of taking area of research into account in hiring is indefensible.

I actually have some sympathy for option 2 here, but it would be an incredibly radical step.⁹ So let's investigate the prospects for option 1. I think they are rather dim.

The motivation for hiring by research field seems straightforward. Departments have (allegedly) an interest in having researchers working on diverse fields. And hiring people who have worked in diverse fields is one way to meet that interest. But given FRA, or even FRAD, there is a big gap in this motivation. All that we can know by looking at a job application file is what a person has worked on. The department, presumably, has an interest in there being diversity in what its members will work on. And we need a bridge between past work and future work here.

One way of bridging this gap would be to insist that the newly hired academic work continue to do (some) research on (roughly) the areas they were hired to work on. I think that's the right way to bridge the gap, but it is inconsistent with FRA and FRAD.

⁸Sarah-Jane Leslie's work (with various colleagues) has shown that there is a strong correlation between how strongly people think that brilliance is required for producing good work in a field, and the gender distribution of faculty in the field (Leslie et al., 2015; Meyer et al., 2015; Storage et al., 2016). Thinking that whoever is hired can work on anything, and it will probably be good if the work they originally did was good, seems similar to me to taking raw talent to be the primary requisite for successful work.

The primary argument of this paper has been the requiring academics to do at least a certain amount of work in a particular area is not a violation of academic freedom. It is no violation of academic freedom to set up something like NYU's Marron Institute (discussed further below). The considerations of this paragraph suggest something stronger, that it is positively bad to not require (most) academics to do work in a particular area, because to not do this encourages an invidious cult of genius. I'm not endorsing this stronger claim, but these considerations do look like the germ of an argument for it.

⁹As of November 16, 2016, there were 209 jobs advertised on PhilJobs.org, and by my count only 36 did not put some restrictions or desiderata on the research area of the hired candidate.

Another way would be to take past research interests as noisy indicators of future research interests. So if you want to hire in philosophy of biology, you might hire someone who has worked in philosophy of biology to increase the probability that that's what they'll work in. The problem with this reasoning is that hiring the person who is most likely to do the best work in the area you want to hire in will lead to some bad choices in realistic scenarios. Imagine you want to hire in philosophy of biology, and you have three candidates.

- A is the best philosopher of the bunch, but has at best a passing interest in philosophy of biology.
- B is the best philosopher of biology, but also has a very strong interest (including a book manuscript in progress) on a completely different field.
- C is nearly as good as B at philosophy of biology, and has no other philosophical interests.

If you want to maximise the expected value of research your department does on philosophy of biology, and FRA or FRAD are in place, the best thing to do hire C. After all, there is a non-trivial chance that B will just work on their book manuscript and related papers, and indeed use it to get tenure. If you want to maximise the expected value of research your department does in philosophy, the best thing to do is to hire A. They are the best philosopher. What's hard to see is the motivation for making what intuitively is the right choice here, hiring B. The solution, I think, is to ditch FRAD, and hire B with the explicit requirement that they do a certain amount of work in philosophy of biology.

Let's say that FRA is false then, and conclude with a more focussed look at FRAD as it applies to tenure cases. This is a somewhat more practical matter, since FRAD is more like the rule that is applied in actual tenure cases. Indeed, here is what the handbook at the University of Michigan (my employer) says about tenure review,

After the appropriate probationary period (see section 6.C "Tenure Probationary Period"), tenure may be granted to those instructional faculty members whose professional accomplishments indicate that they will continue to serve with distinction in their appointed roles. Tenure is awarded to those who demonstrate excellent teaching, outstanding research and scholarship, and substantial additional service, **each of which must be relevant to the goals and needs of the University, college and department.** The award of tenure is based on the achievement of distinction in an area of learning and the prediction of continued eminence throughout the individual's professional career (The University of Michigan, 2016b, emphasis added)

As far as I can tell, in practice the bolded clause is interpreted in line with something like FRAD. The wording is ambiguous; it could just as easily be read as supporting the elbow room standard that I prefer. But I don't believe that's how things work in practice.

We have two questions to answer then. First, how similar are FRAD and the elbow room standard? And second, in cases where they differ, which provides a better model for building a university. I think they are not particularly similar, and the elbow room standard is much better.

There is, of course, a certain similarity between the two standards. FRAD says work on whatever you like, provided it is in the same discipline as the work that got you hired. The elbow room standard says work on whatever you like, provided it is sufficiently similar (along some salient dimension) to the work that got you hired. And being in the same department is a dimension of similarity. But it isn't, ultimately, a particularly important one. Making it of central importance, as FRAD does, leads to numerous avoidable errors.

For one thing, FRAD gives some academics more freedom to switch fields than the elbow room standard could possibly justify. Philosophy is a very broad field. Just because one is doing really excellent work in one field is very little evidence that one will be able to do excellent work in another field. Thinking that it is evidence is to just relapse into a restricted version of the myth, or cult, of genius. So at least in some cases, the elbow room standard will be more restrictive than FRAD. But in other cases it will be less restrictive, and those are perhaps more important in practice.

Disciplines have boundaries. Those boundaries are vague, but they are there. Some people work on topics that are very near to a boundary, and some work on topics that are far from a boundary. FRAD impacts these two groups in very different ways, and the difference is unfair. Someone whose work is near a boundary can't just drift into any nearby field, since the nearby fields may be outside the disciplinary bounds. To take one clear example, a researcher hired for work on the semantics of modals can easily drift into other areas of semantics, or onto modal fallacies in argumentation, but not into work on the syntax of modals. For whatever reason, we've decided the boundary between philosophy and not-philosophy around here is very close to the the syntax-semantics boundary. And this is a violation of the elbow room principle, since this kind of move from the semantics to the syntax of a particular class of expressions is just the kind of move to a closely related field that the elbow room principle is designed to protect.

It won't help here to say that there should be overlapping areas of research concern between departments. The crucial question is the boundary between X and not-X, not the boundary between X and Y. If someone is near the X/not-X boundary, they could have a tenure home in X, and then drift into not-X. They deserve protection in their research, and FRAD wouldn't provide it. It would perhaps help if literally every possible area of research was such that there was some department it was not just in, but in without being near the boundary. But a world with academic departments organised that way feels very different to the one we are in.

It won't help to say we should just make the boundaries larger. Unless we abolish the boundaries altogether, the problem will persist. And abolishing all the boundaries would create more problems than it solves. The boundaries play useful roles right now. It is good that hiring and tenure decisions for, say, a position in metaphysics are made by people with a broad range of philosophical backgrounds, and not (in the first instance) by an arbitrary collection of people from across the university. The ideal here is not no boundaries, but porous boundaries.¹⁰ Boundaries exist so that local experts, and people with special local interests, get extra say on questions of local concern, but they are porous so they get in the way of freedoms. Replacing FRAD with the elbow room principle gets the balance right.

¹⁰Academia is hardly the only place where this is the ideal.

It also doesn't help that the boundaries are vague. In general, if something is true no matter how a vague term is made precise, it is a good bet that it is true.¹¹ But we can say a bit more about why vagueness doesn't matter in this particular case. The following principle looks both true, and the best bet for why we might think vague disciplinary boundaries make FRAD more palatable.

- If someone's work is clearly within discipline X, then any related area they could reasonably drift into under the elbow room principle will not be clearly outside discipline X.

If everyone who is hired is clearly working within X, then FRAD might be no more restrictive than the elbow room principle. After all, any permissible drift will not take one clearly outside one's home discipline.

The problem is that only hiring people whose work is clearly within the hiring discipline is a terrible idea. Indeed, it is a worse idea than FRAD. It ensures that one will only hire in safe, traditional areas of research. And that's a plan for stagnation, not for doing the best research. Sometimes we have to hire people on the frontiers, and sometimes their work will drift clearly outside one's discipline. That's just a cost of having a dynamic research program, and attempts to avoid paying this cost will make things even worse.

There is one problem that vague disciplinary boundaries does help with. Vague boundaries are easier to shift than precise boundaries. That's because there was never any agreement on where they were in the first place, so no agreement has to be overturned. And if enough people work who were hired in X start working on an area that used to be outside X, we should just start treating that area as inside X. But we don't have to make this conceptual shift every time a good philosopher does good work on a nearby topic. Some good work is in other fields, and that's ok.

There are other odd features of FRAD. If applied consistently, it would lead to treating some like cases in very unlike ways. There are several fields that are set up as departments in some universities, and as programs in other universities. For example, in America right now there is a divide among universities about whether to set up things like Women's Studies, Cognitive Science, and PPE as independent departments, or as programs run collaboratively by a number of different departments. There are administrative considerations on either side of this choice, and these considerations vary somewhat between different universities. If they are set up as programs, then anyone hired in to them will have a tenure home in one of the constituent departments. And given FRAD, that will put certain limitations on their research. Those limitations will be very different to what they would face if the unit were its own department. But it seems very wrong to think the administrative decision to set up a unit as a department or a program impact the freedoms of people hired to work in those units. Where FRAD applies, however, this difference is dramatic. Someone hired in a philosophy department to support a Cognitive Science program could move to work on philosophy of religion, but not experimental developmental psychology. Someone hired by a Cognitive Science department would have the opposite set of freedoms. I'm not sure what is optimal here, but it is very odd that an administrative decision should impact researchers in this way.

¹¹This principle traces back at least to Keynes (1936, Ch. 6).

But the biggest problem with FRAD is that it makes the disciplinary boundaries too important. Young researchers shouldn't have to second guess whether a particular development of their research is inside or outside a vague, shifting, boundary. The solution isn't to abolish these boundaries, any more than rights of free movement across an area is a reason to abolish all political boundaries within that area.¹² Rather, the solution is to downplay them. If exercising their elbow room rights takes an academic outside the purview of their home department, that's just a cost of having a dynamic research program.

And it is really the effect on these younger scholars, trying to pre-judge what their reviewers at tenure time will think, that I'm most interested in here. As I noted above, I don't know of any clear cases where someone was turned down for tenure because their research was in the wrong area. But I know many cases of academics who have put off more speculative research projects until after their tenure review. And the reason, typically, has been that they are nervous that the outputs of the new project would be discounted, merely in virtue of their subject matter, at the time of tenure review. This feels like an undesirable feature of the status quo, and one that could be remedied by rethinking why we care about what an individual academic works on.

So while I disagree with the strongest statements of academic freedom, I think the position I'm endorsing allows greater freedom in practice than existing practices like FRAD. There are, I would guess, many more people who are worried that their research is drifting away from what their colleagues will regard as *really* part of the discipline than there are people who would like the freedom to jump to an area they have no training, expertise or background in. Defending the elbow room principle, or freedom to drift, will take care of those concerns. And the principle is much easier to defend in theory than FRA or FRAD. So it, I think, is the core important principle concerning academics' rights to direct their own research.

Conclusion

I've focussed here exclusively on research, and not at all on teaching. But in many ways what I'm saying here could be summarised as the view that the norms concerning topic choice are fairly similar in research and in teaching. If I'm given a course on history of political philosophy to teach, then I better teach history of political philosophy, and not, say, formal logic, or Australian geography, or baseball statistics.¹³ It isn't in any way a violation of academic freedom if I'm required to teach the subject I signed up to teach.

But in practice, and in theory, there is a lot of freedom within the boundaries of a course. If I want my history of political philosophy course to include thinkers who are not commonly central to the story Anglophone political philosophy tells about itself, I should (and typically would) be free to include them. If I think the most relevant secondary literature is by people in departments other than philosophy (e.g., history, political science, women's studies, etc) then I should be able to base my syllabus around such writers. Now I personally haven't taught history of political philosophy since I was a post-doctoral fellow

¹²By analogy, it's a good thing that Ann Arbor has a city council, and it's a good thing that there are no barriers to moving in or out of Ann Arbor.

¹³Compare the discussion of academic freedom and the requirement to teach the topic of the class in Cole et al. (2015). The example of teaching Australian geography in a class that is not on that is taken from one of their survey respondents.

who was too nervous to consider any such plan. But it's exactly the kind of thing academic freedom should protect - and I suspect in most cases it is what academic freedom would protect.

The same I think should go for research. If one is hired to research history of political philosophy, then it is reasonable for the university to require that one do just that. It isn't reasonable to require one do only that; people should be allowed to explore what they want. But it is reasonable to require some work on what one was hired to do. Yet if doing that takes one outside the bounds of what is (hereabouts) considered philosophy, that should be fine too. Do what you're hired to do is a good principle; FRAD is not.

The picture of academia I'm trying to promote is one where more units are free to operate the way that Paul Romer describes the Marron Institute of Urban Management at NYU as operating.

[I]nstead of giving its faculty members the usual freedom to study anything that that seems interesting, the institute lets the problems that cities face set its research agenda. Because these choices are not the usual ones on campus, many people complained. (Romer, 2016)

If universities want to give people complete freedom to set their own research agenda, I'm not going to complain about that (much). What I want to deny is that setting up things like the Marron Institute is a violation of academic freedom. There is a lot to be gained by hiring people for relatively specific research tasks, and it isn't a violation of their freedom to hire them in this way. And that's especially true if the constraints on their research agenda are set just by the questions that their research team is focussed on, and not by the disciplinary homes that house the thinkers they engage with.

References

- Bickel, Alexander. 1975. *The Morality of Consent*. New Haven, CT: Yale University Press.
- Biglrami, Akeel and Cole, Jonathan R. (eds.). 2015. *Who's Afraid of Academic Freedom?* New York: Columbia University Press.
- Cole, Jonathan R., Cole, Stephen, and Weiss, Christopher C. 2015. "Academic Freedom: A Pilot Study of Faculty Views." In Biglrami and Cole (2015), 343-394.
- Hayek, F. A. 1945. "The Use of Knowledge in Society." *American Economic Review* 35: 519-530.
- Keynes, John Maynard. 1936. *The General Theory of Employment, Interest and Money*. London: Macmillan.
- Leslie, Sarah-Jane, Cimpian, Andrei, Meyer, Meredith, and Freeland, Edward. 2015. "Expectations of brilliance underlie gender distributions across academic disciplines." *Science* 347: 262-265.
- Meyer, Meredith, Cimpian, Andrei, and Leslie, Sarah-Jane. 2015. "Women are under-represented in fields where success is thought to require brilliance." *Frontiers in Psychology* 6: 1-12.

Nelson, Cary. 2010. "Defining Academic Freedom."

Romer, Paul. 2016. "Everybody wants progress; nobody wants change."

Shweder, Richard A. 2015. "To Follow the Argument Where It Leads: An Antiquarian View of the Aim of Academic Freedom at the University of Chicago." In Biglami and Cole (2015), 190-238.

Storage, Daniel, Horne, Zachary, Cimpian, Andrei, and Leslie, Sarah-Jane. 2016. "The frequency of words like "brilliant" and "genius" in teaching evaluations predicts the representation of women and African Americans across academia." *PLoS One* 11: 1-17.

The University of Michigan. 2016a. "The University of Michigan Faculty Handbook, 6.A General Principles."

—. 2016b. "The University of Michigan Faculty Handbook, 6.B Criteria for Tenure."