Hanging On To The Edges:
The need for discipline

> …far from characterizing [academic disciplines] by theoretical,
> disembodied abstraction, I view them as sites
> for the coordination and embodiment of skill.
> - Timothy Lenoir, *Instituting Science*

On and off over the years, I have sung in choirs of various types. I get by just fine, as long as I can stand in the middle of the basses somewhere, carried with the rumbling tide, but I am not a very good singer, or likely to ever become one. So why do I do it? Well, there are the obvious things: it's a bit social; it keeps me off the streets of an evening; and—not to be overlooked—it makes me get to know and love pieces of music in a deeper way than having the radio on will ever do. But none of these is the thing I enjoy most about choirs. The thing I enjoy most—and this is something I have not publically admitted before—is watching the choir-master (or choir-mistress) do their work.

If you have been to a choral rehearsal, you will know what I am talking about. The choir-master will be sight-reading the piano reduction of the orchestral score, switching what they are playing to bring out one or the other of the choral lines that is most important at the particular point, all the while listening to four different choral sections singing, spotting the kinds of difficulties they are having, and remembering tips that need to be given afterwards. The choir finishes a movement. I am just pleased to have got to the bottom of the page at about the same time as everyone else. The choirmaster, however, pipes up: "I was wondering in bar 73 if that diminuendo should be in the sopranos only, since it is that $c^{\#}$ in the alto line that is heralding the key change we are going to in the next section, and we need to hear the development". And here I am thinking: *how were you wondering anything in bar 73?* You were sight-reading three lines of music and singing a fourth, whilst simultaneously attending to the timing and pitch and volume and pronunciation and breathing of about 40 different people. It takes me about 3 minutes just to work out where bar 73 is. Yet you did all you were doing so effortlessly that you still had spare capacity. You have an immediate and physical dexterity with the components of the music that I can relate to strongly, but never reproduce. "Basses, you seemed uncertain on that *b* natural entry; you are getting swayed into pitching it flat because of the *e* flat in the tenors. Listen the end of the sopranos' tune the bar before, and think, 'happy *birth*day to you'". Will do, boss. Which ones are the sopranos again?

Yes, there's something fascinating—moving even—about seeing people exercising real skill. You might think that observing the highly skilled would be alienating or aversive for those of lesser skill. In fact, the opposite seems to be the case. We flock to hear virtuoso musicians and watch master chefs on television; we seek out dry-stone walling competitions, sheep-dog trials, and demonstrations of glass-blowing. I remember once Melissa and I were having a complex-shaped roof covered in lead on the back of our terraced house. The roofer gadgie duly turned up with his bag of lead hammers and his rolls of lead. I said to our next-door neighbour that I hoped that the builders were not disturbing his peace. On the contrary, the neighbour said—I cannot confirm but would like to imagine a flutter seizing his breast—it's such a privilege to be able to see that man *work*.

It's a privilege to see a skilled person work because, I would contend, it connects us to something deep about being human. We are the species that is good at getting good at doing stuff that is hard to do. We do this rather eccentrically and ecumenically: sometimes in domains with a utilitarian payoff, sometimes not. Some individuals take skill further than others in any particular domain, and this

capacity for individual specialization is itself interesting and consequential. But skill acquisition is not the preserve of a few geniuses: it's just what human beings do. It is because it is so pervasive that we only notice the extreme cases. By the time a chimpanzee is a few years old, it has pretty much reached peak productivity. Yet humans live their lives by expressing extraordinary skills that can take twenty, thirty, forty years to develop and refine; and I have an obscure and possibly sentimental sense they reach their fullest sense of personhood through doing so[1].

It upsets me that our living species gets the name *Homo sapiens*, the human that knows. Watching my choirmaster, or the roofer gadgie, what strikes me is not *that they know* so much as *that they can do.* What they have is not knowledge in some purely informational, propositional sense, something that could be stored on a USB stick. If it is knowledge, it is procedural knowledge, instantiated in and distributed across the whole of the body, and realised in patterns of movement (striking the keys, bending the lead, tensing the diaphragm, modulating the larynx). Indeed, it's often knowledge that cannot be expressed explicitly or imparted verbally, so highly routinized and embodied has it become. Thus, I feel that we lost out in the naming game to our extinct relative, *Homo habilis*. The skillful human. That's what I would like us to be called. I would take *being able to do* over just *knowing* any day of the week.

What name should we give the virtue that we recognize in highly skilled people? I would, for the sake of today, like to call it *discipline*. Discipline, in the everyday sense of self-control, is what is needed to drive oneself through the 40,000 hours of practice and training that high skill requires. But the word *discipline* has broader resonances: it links back to the Roman deity *Disciplina*, with her virtues of skill, self-improvement, economy of action, dedication to the guild, and simplicity of life. What could be more attractive, then, than *discipline*?

§

I began this essay with a paean to discipline in order to wrong-foot you. Because if you are anything like me, you probably have a well-developed sense that in science, *disciplines*, entailing as they seem to do *disciplinary boundaries*, are a bad thing. I have spent my whole career railing against them. Indeed, it's something of an identity marker amongst my people to deplore the balkanization of the study of human behaviour across so many discrete disciplines; to blame disciplinary divisions for our failures to progress; and to claim to be trans-disciplinary, poly-disciplinary, or post-disciplinary, in our orientations. We don't tend to say much about how the landscape of our post-disciplinary utopia ought to be organized for practical purposes; only that, perhaps rather suspiciously, a lot more status and resources would be accorded to people who….well, to people who are like ourselves, really.

Often in academia, we clarify our reservations about some idea by stating that idea in its most stark or simple form: the famous *straw man* strategy. But a very useful complement to the straw man strategy is the *steel man* strategy: try to characterize the idea you oppose in its best, most sophisticated possible form[2]. If you can defeat even the steel man version of an idea, then it really is a bad one. More likely, you will discover unappreciated virtues in an idea that you previously thought of as wholly bad, and adopt a more nuanced position. I feel like this about disciplines. I have read (and written) so many tirades about how segmentation of science into discrete disciplines is bad. Yet academic disciplines got invented, have been perpetuated by a lot of very clever people, and largely continue to exist; indeed, science has been doing conspicuously well since about the time disciplinary structures became established. All of which leads me, in the spirit of the steel man strategy, to ask: what is there that is *good* about disciplinary structures?

[1] Just one reason why the emphasis in universities on giving students *transferable skills*, presumably so that they use them in a series of *bullshit jobs*, is potentially misguided. Virtuosity can be pretty much defined as a high level of non-transferable skill.

[2] I am grateful to Brett Beheim for introducing me to the idea of the steel man strategy.

With this question developing in my mind, it was naturally with interest that I listened, in a restaurant in Helsinki as it happens, to a friend telling me about how a certain university was reimagining the structure of their curriculum. Instead of organizing courses of study around discrete disciplines, students would study a portfolio of modules whose subject matter was defined by phenomena or problems, like migration, climate change, or violence. The problems chosen were exemplary, in that their solutions could not be generated by any discipline acting alone. Within each problem, students would learn how economists thought about it, how sociologists thought about it, how biologists thought about it, and so forth. This would give them the ability to compare, contrast and syncretise different perspectives without being artificially shunted into the confines of any one of them. Great, I thought, inanely: exactly my kind of thing. "How does it work in practice?" "Terrible", she replied. Students could produce generalized and often stereotyped comparisons of different disciplinary approaches, but without enough depth or detail to actually implement (or improve) any of them. They were left with the abiding impression that what you believe about the world is really just a matter of what identity you choose to adopt, rather than the consequence of systematic epistemic work using justified standards. In short, they came out of their studies *not knowing how to do anything*.

The finale of this anecdote reminds us that the origin of the term *discipline* in the academic context is a pedagogical one: the set of training you need in order to be a competent and useful practitioner in a domain. And this training is not reducible to the acquisition of factual statements; not even reducible to the acquisition of factual statements plus frameworks for interpreting and explaining them. Perhaps more of either of these things, the concept describes a set of core *physical skills*. In this sense, the *discipline* of the mathematician or ethologist has more in common with the discipline of the roofer, the stonemason, or the choir-master than one might at first imagine. When we think about practical skills, inter-disciplinarity does not seem like a particular virtue. It might not be bad, but it is a lesser virtue than excellence *within* the relevant domain. Who would you hire: the inter-disciplinary welder— "I can weld a bit, and I can critically compare welding to carpentry!"—or the welder whose welding is simply excellent?

§

It might be useful to separate analytically between two components of disciplinarity: the *declarative* and the *practical*. On the declarative side, disciplines sustain particular structures of explicit belief and understanding, and they do so in part through political and ideological operations. (I am not saying that science is mere ideology or mere politics—on the contrary, its content is in the end constrained by nature. It is nonetheless a social process that proceeds year to year through ordinary human manoeuvring.) Disciplines exert power and sustain ideologies through control of what gets published (and in what form); through control of funding panels; and through control of academic hires and curricula. By such means, disciplines can define what types of question can be asked, in what way, and what constitutes an acceptable answer. They provide handy non-reasoned authority and legitimation for particular decisions and inferences. I remember one conversation with a colleague about why she was interpreting a particular behavioural phenomenon as evidence for a particular cognitive mechanism in the children she studied. Her reply was simply that, in developmental psychology (her discipline, not mine), that's how researchers interpret it. Whilst her answer was pragmatically realistic, it was epistemologically unsatisfactory. I was tempted to recall George Berkeley's jibe against scholastic philosophers: "when a Schoolman tells me 'Aristotle hath said it', all I conceive he means by it is to dispose me to embrace his opinion with the deference and submission which custom has annexed to that name"[3]. Judging this a bit heavy for a Tuesday lunchtime, I held my tongue.

It's the declarative aspects of disciplines, particularly the way they trammel and police researchers' explicit cognition, that are the easiest to use in an indictment against them. They normalize assumptions that should be exposed, provoke cognitive conformity, and delimit possible moves and

[3] Berkeley's jibe is from the introduction to his *Principles of Human Knowledge* (1710).

juxtapositions. Thus, they blind us to aspects of the phenomena, or theoretical resources, that might hold the key to progress. I'll come back to this argument, which is the one I have habitually relied on in advocating post-disciplinarity. For now, let us concentrate on the fact that this argument gives no recognition to the practical components of disciplines.

The practical components of disciplines are the physical skills they serve to inculcate, transmit and refine. When you hire a mathematician, you want them to be able to do matrix algebra, and when you hire a molecular biologist, you want them to be able to pipette. You can expect that they will be able to do so skilfully in virtue of the disciplinary training that they have received. Disciplines in this sense are guilds of artisans. It seems obvious that to drive levels of skill higher and higher, there will need to be specialization, deep apprenticeship, and assortment of artisans with the others of their guild. The wood turners may get on well with the potters, but they will naturally want to spend a lot of time with other wood-turners to mutually improve skill, and to pass it on to the next generation. Viewed from the practical rather than the declarative perspective, then, the existence of separate disciplines seems much less pernicious and much more natural. And indeed, historically, it is the development of particular laboratory or field practices, as much as declarative theoretical commitments, that gives rise to new disciplines[4].

You might concede the need, on practical grounds, for disciplinary specialities such as electrophysiology or molecular biology. You might however still cling to the view that the broad domain covered by psychology, sociology, anthropology and economics should simply be one open field, since these disciplines all look at the same thing, namely human behaviour. That's only right to an extent. After all, electrophysiology, fMRI and EEG all look at the same thing—neural activity—but nonetheless require discrete skills. Capturing and analyzing large amounts of data on monetised transactions; getting to know and understand the lives of a particular small social group; designing and interpreting large-scale attitudinal surveys; making experiments to isolate the causal structure of particular cognitive processes: these are different operations that each require deep reservoirs of skill to do well. The lack of expensive equipment and white coats should not blind us to the complex and distinctive skills required in each case. Talking to people in an informative way is a skill. Designing surveys is a skill. Making experimental materials is a skill. Gathering and analysing data on prices and sales is skills. And I know from experience how easy it is, as a disciplinary novice, to do these things so badly that the results are basically useless.

§

Another musical anecdote: Melissa and I play baroque recorder duets. She's the more skilful player. I am known for proclaiming, as we embark on a new piece: I think this movement should have a *misterioso* quality. Her translation: what you mean is that it's too difficult for you to play properly. Your *misterioso* is basically made up of slowing down on the hard bits, getting some of the notes wrong, and a distracting smoke-screen of emphatic head movements. If I practice a bit, guess what, I learn that it sounds even better if, rather than *misterioso*, it is just played well.

In like fashion, I have been involved in several research projects that were self-consciously trans-disciplinary, usually mixing some kind of evolutionary biology theory with some kind of social science data. As we accumulated rejection slips from the best journals in the field, we consoled each other with the insight that journal editors and reviewers are cognitively trapped in disciplinary silos, and prejudiced against any kind of attempt to transcend them. Our indignation affirmed our sense that we were right and what we were doing was important. Indeed, at times, it seemed like the more the disciplinary specialists rubbished it, the more vindicated we felt in our world-view. I shouldn't need to tell you that this is an extremely disturbing epistemic direction to be headed in. Looking back over this history now with the benefit of a bit more experience, I can see a plainer truth: the studies we were

[4] For this view of the history of science, see Lenoir, T. (1997). *Instituting Science: The Cultural Production of Scientific Disciplines.* Stanford: Stanford University Press.

doing often contained interesting germs of ideas, but were *just not very good*. Our trans-disciplinarity was, not always but all too often, the *misterioso* of science: a pretext for sketchy methods, careless design, hasty data analysis, inferential over-extension, and lack of theoretical precision.

Pat Bateson, nearly fifty years ago now, distinguished between roundheads and cavaliers in the study of behaviour. The roundheads are methodologically impeccable in every respect, but, in his words, unwilling to flirt, let alone dance, with ideas. The cavaliers exhibit great dash and intellectual exuberance, but this does not come without cost: they "also are notoriously unsound and constantly confuse inference with evidence"[5]. It would be too much of a stretch to equate roundheads with disciplinary specialists and cavaliers with trans-discipliners. There can after all be people who are cavalier within the bounds of one discipline, and trans-discipliners who develop beautiful and rigorous methods. There is no doubt in my mind, though, that across the study of behaviour, greater trans-disciplinary comes with a more cavalier attitude to methods on average. Which of course links us back to the practical components of disciplines.

I can think of many examples from my career where the growth of a trans-disciplinary research area produces high-impact publications that are simply unsound, and whose unsoundness could easily have been prevented if there had been a little more discipline. There are the famous economic experiments 'demonstrating' that people have an altruistic concern for others' financial outcomes as well as their own. These experiments consist of artificial financial dilemmas, in which many participants can pay into a group fund that will benefit everyone at personal cost to them, when there is no requirement to do so and no future or current penalty for not doing so. I firmly believe, by the way, that humans have other motives than just maximizing personal gain, including prosocial motives. It's just that these particular experiments don't provide good evidence that this is the case, because they do not include the right control conditions. In particular, they do not include control conditions to separate the motivation that others should benefit from, for example, failure to truly understand the rules of the game, or a dislike of using the extremes of a scale. Once you include control conditions to rule these alternative accounts out, it becomes clear that the reasons people often fail to maximize their income in these artificial dilemmas may have nothing to do with concern for others' outcomes[6]. The original experiments—often published in the highest-profile inter-disciplinary journals—were done by brilliant economists who lacked deep discipline in experimental psychology. Experimental psychology has many faults, but one of its virtues is real skill to designing experiments and, in particular, the almost ubiquitous need for multiple control conditions in order to make inferences about the meaning of an experimental effect.

To take another example, in primate neuroscience and primate cognition, it's common to find people taking 1000 or even 5000 trials to train their animals on a simple discrimination (say, between two colours) where one option is rewarded and the other not. Even after thousands of trials, performance is not always very good. There is in fact a view out there that smart animals like chimpanzees do not readily acquire arbitrary discriminations in the way that rats or pigeons do[7]. The fact that these monkeys and apes take so many hundreds of trials, though, is an artefact of the way they are trained. Intuitively, it seems like the way you would make an animal learn an association between a stimulus and a reward is by pairing them more and more often; so in these training paradigms there may be hundreds of pairings of stimulus and reward, with very short intervals between them, in a single session. It turns out that intuition is spectacularly wrong in this particular case: the way you make an animal learn an association between a stimulus and a reward is to make the stimulus, and hence the

[5] Bateson, P. (1970). What is learning? *New Scientist,* 25 June.

[6] Burton-Chellew, M. N. and S. A. West (2013). Prosocial preferences do not explain human cooperation in public-goods games. *Proceedings of the National Academy of Sciences* 110: 216-21.

[7] See Hanus, D., & Call, J. (2011). Chimpanzee problem-solving: Contrasting the use of causal and arbitrary cues. *Animal Cognition* 14: 871–878; and Bateson, M., & Nettle, D. (2015). Development of a cognitive bias methodology for measuring low mood in chimpanzees. *PeerJ* 3: e998 for some discussion.

reward, very rare. Once you do this, the number of pairings required comes down by orders of magnitude[8].

Old-school rat and pigeon animal learning theorists knew this very well, and they could train an arbitrary association perfectly in a dozen or two of trials at most. They also knew how best to navigate the way to learning a full discrimination: first training the association between the positive stimulus and reward, then introducing the unrewarded alternative, then building up to choices. Just sticking two colours in front of an animal again and again in quick succession is only going to teach them something by a near-endless war of attrition; the unrewarded stimulus ends up temporally proximal to the rewarded one, and the required informative contingencies in experience are lacking. So the interesting question is why the skill held by animal learning theorists has not found its way into the communities studying primate neuroscience and primate cognition. Well, by and large those researchers don't have deep discipline in animal learning theory. Instead, they are coming in from anthropology, or if it is from psychology, it is cognitive psychology. One of the sad and pointless things about the 'cognitive revolution', in which behaviourism was allegedly 'overthrown', is that a lot of really useful skill in how to make animals learn, as well as how to design beautiful experiments, was lost in a kind of year-zero mentality. We need to build on the practical skill of behaviourist psychology, not throw it out with the cognitive bathwater.

To take a final example, in the last five years, I have begun to work in telomere biology, not really because I know much about telomeres, but because of the possibility they offer us to provide an integrative marker of the insults and damages inflicted by the world over the course of an individual's life to date—what is fashionably called 'the exposome'. We did four successive experiments where we showed that telomeres shorten very rapidly in the nestling starling, then a fifth experiment where they did not—they appeared to get longer, in fact. Had we discovered an extraordinary exception that would form the basis of a *Nature* paper? Reverse the ageing process! A group of Northumberland starlings holds the key to immortality!

I showed the data to a colleague I admire who is more skilled than me in the actual lab work. She didn't need to see the results of my multi-level model; she didn't even need to see the statistics on the technical replicates. She certainly didn't need to hear my elaborate theoretical interpretation. "Oh", she said, "your reaction hasn't worked. Look, those numbers are too high. Your primer concentration must be wrong." And that was it. She was like a choir-master, knowing in her bones where and why the basses had been misled. The basses have not stumbled on an interesting new direction for Western music; they are merely singing badly. The methods sections of the published papers on measurement of telomere length by quantitative PCR say all kinds of useful things, of course, but what they don't say is: if these numbers aren't quite a lot smaller than those other numbers, you've probably just done it wrong. And in the emerging trans-disciplinary field of telomere epidemiology, a number of the most exciting 'findings'—such as the idea that longer telomeres shorten faster, or that over short periods, the telomeres of about 50% of people get longer, probably reflect measurement error as much as anything else[9]. If we had all had a bit more discipline prior to publication, perhaps science could have proceeded more efficiently.

§

[8] This phenomenon becomes comprehensible once you appreciate that what you want to maximize is not the number of times that a stimulus has been paired with a reward, but the information an event carries about a reward. Rare events carry more information. See Ward R.D., Gallistel C.R. & Balsam P.D. (2013). It's the information! *Behavioural Processes* 95:3–7.

[9] See Steenstrup, T., Hjelmborg, J. V. B., Kark, J. D., Christensen, K., & Aviv, A. (2013). The telomere lengthening conundrum--artifact or biology? *Nucleic Acids Research* 41: e131; and Verhulst, S., Aviv, A., Benetos, A., Berenson, G. S., & Kark, J. D. (2013). Do leukocyte telomere length dynamics depend on baseline telomere length? An analysis that corrects for "regression to the mean." *European Journal of Epidemiology* 28: 859–866.

The conclusion so far seems to be that we should want to retain all the practical virtues of disciplines, namely having people with high levels of specific technical skill, but abolish the declarative distinctions between them. I don't know if this is possible; as Timothy Lenoir has argued, the connections between the practical and the declarative components of scientific programmes are intimate[10]. Indeed, I have often wondered to what extent the explicit propositional commitments of particular disciplines arise spontaneously from the types of practical activities that their work involves. If you spend all day in dealing with prices and purchases between anonymous actors in a fungible currency, maybe you start to think like an economist. Maybe you would do so even if you had not been indoctrinated in micro-economic theory; even if that theory did not already exist. If by contrast you spend all day in open-ended non-monetised interaction with a small group of people, maybe you either need to lie down in a very quiet room, or you start to think like a social anthropologist, or both. On this view, the declarative differences between disciplines would not be (just) historically contingent ideologies sustained by the dynamics of power and influence, but the inevitable *déformations professionelles* arising from individuals habitually working at different practical activities. Hence, there is no clear dividing line where the concern that the practical activities be done skilfully ends, and ideological boundary-maintenance begins.

So it's going to be hard work to overthrow the declarative balkanization of science without any loss of practical skill. And actually, allowing the steel man to be even steelier, there exists an interesting literature arguing that some declarative balkanization of science is a good thing. The argument (which I am interpreting in my own, slightly *misterioso*, manner) goes something like this: the progress of knowledge relies on variation, for exactly the same reason that adaptation by natural selection does. We need to be trying out a lot of different ways of thinking. It is difficult for individuals to implement more than one way of thinking simultaneously. Thus, the simultaneous existence of multiple *groups* of individuals, each thinking about human behaviour in a different way, is actually a strength. The internal coherence of the groups is not undermined by constant blending, and at the meta-level, competition between the groups for society's attention and support is an invisible hand propelling humanity towards a higher level of overall understanding[11].

This kind of argument comes up in domains other than science. The economic dynamism of early modern Europe (in contrast to China, say) has been attributed to the existence of many independent city states, in competition with one another, where different things could be tried out: the good things could spread, and the bad things didn't drag the whole continent down. Political devolution in the United Kingdom has had the virtue that new policy ideas are tried out independently in Scotland or Wales, with adoption more broadly dependent on the results of those innovations, which are effectively natural experiments. In science, this kind of progress-through-variation depends on the disciplines being somewhat informationally isolated (or all the variation would be rapidly washed out), but nonetheless a little bit leaky. For all of the United Kingdom to benefit, the idea that works well in Scotland does have to find its way to England in the end. Similarly, the ideas of behavioural ecology need to be able get to sociology (and vice versa), but without sociology just becoming behavioural ecology.

What we seek then, if this line of reasoning is correct, is a goldilocks level of trans-disciplinary integration: not too little, not too much, but just right. The optimal level would be something that could be modelled. And much as people like me moan on that the current level is too little, the solid evidence that this is the case (for example, that science progresses faster in periods or areas where disciplinary integration is greater) is currently lacking. Indeed, bibliometric studies have generally

[10] Lenoir, T. (1997). *Instituting Science: The Cultural Production of Scientific Disciplines.* Stanford: Stanford University Press.

[11] For versions of this argument, see Stichweh, R. (1992). The sociology of scientific disciplines: on the genesis and stability of the disciplinary structure of modern science. *Science in Context* 5: 3–15; and Jacobs, J. A., & Frickel, S. (2009). Interdisciplinarity: A critical assessment. *Annual Review of Sociology* 35: 43–65.

concluded that inter-disciplinary leakage is already quite substantial and ubiquitous under the status quo[12].

§

Having spent a few days with the steel man argument for disciplines, has my commitment to greater trans-disciplinary integration changed? In one sense, no: there is still so much that has not yet been done, but could be if we could draw our resources and assets together. But I have perhaps tempered my views. Not everyone should be bludgeoned into inter-disciplinarity all the time; and when we undertake it, we should undertake it constructively. The thinking styles of other disciplines, arbitrary though they can appear, have arisen for particular reasons, often to do with the reality of the practical activities in that area, and these reasons need to be understood. It's perhaps inevitable that there will always be cavaliers and roundheads, and that the first to rush into a new area will be more cavalier than those who come along later to tidy up. Nonetheless, we can try not to be *too* cavalier. The disciplinary specialists have good practical reasons for apparently niggling concerns.

Modish desires for inter-disciplinarity should not trump our commitment to training people to high levels of specialist skill. Some of the key skills that we need to train—statistical modelling, for example—are already trans-disciplinary anyway, but others are specific to particular kinds of data. The glitter of inter-disciplinary declarative statements can, for certain types of personality, be more attractive than the grit of intra-disciplinary practical skill, but at least some of the latter is indispensable if one wants to get anywhere. I wish I had spent more of my post-doc years learning how to do something hard from some gadgies who really knew. It would have been more sustaining than the showier things I did instead. I was too young and too arrogant to understand what I was missing. I have spent more time trying to learn methods in latter years, which I don't regret at all. It is salutary to continue to be a learner.

Rather than viewing other disciplines as competitors to be rhetorically trashed, ailing companies susceptible to a hostile merger, or sources of revealed truth, we should try to understand what skills those disciplines embody, and see if we can get access to those skills to raise the level and refinement of our research. This is something that can be done by training, and/or by collaboration. The best way of understanding what a discipline has to offer is not just to read its theoretical end-products (though one might be interested in these), but to attend to the detail of its practical methods and processes. Ideally, it's good to gather and handle the kind of primary data that particular discipline deals with; this might illuminate why its practitioners have the particular concerns that they do. If you want to build deep cross-discipline links, my hunch is that you will get further by spending time shadowing your collaborator working in their lab (or field site, or archive, or computer programme), than you will giving formal talks about each others' high-level belief systems. The latter activity rapidly becomes polarising: it entrenches the apparent incommensurability, the self-justification, the stereotypes. Trying to do work together, on the other hand creates common ground.

What should the academic landscape look like at the macroscopic scale? What we seek is some kind of limited and provisional modularity. Students of social life need to be able to pursue the phenomena of social life without having to wait for the physicists to reconcile the strong and the weak nuclear forces first. They need to be allowed to build up and stabilize the best possible skills for doing so. On the other hand, the modules should never become completely hermetic. If our social theories are incompatible with the laws of physics or biology, we can't just shrug and say, "not our problem" indefinitely. How can we realize this paradoxical, unity-in-plurality, world?

[12] See Jacobs, J. A., & Frickel, S. (2009). Interdisciplinarity: A critical assessment. *Annual Review of Sociology* 35: 43–65 for review.

Something we can draw on here is the idea of a small-world network[13]. Small-world networks manage to have two interesting properties simultaneously. The first is a high degree of clustering or cliquishness: most interactions are local, and most of a focal individual's interaction partners are also interaction partners of one another. Thus, small-world network architectures would allow for domain-specific transmission of skill—for guilds, if you will. The second property is a surprisingly short path length connecting any two nodes (the famous six degrees of separation/Kevin Bacon/Paul Erdős). Thus, if you put information into a small-world network at any point, it's really not long before it shows up everywhere (if it's spreadable information that is). Small-world-ness turns out to crop up in evolved biological systems, like nervous systems, presumably because these have similar requirements as knowledge does: specific sub-parts need to deal with specific problems through local specialization, but the whole thing also has to function somewhat coherently at the macroscopic scale.

How do you make a small-world network? It's simple. You set up a lot of short connections between neighbours, and a smaller number of long-distance ones to a random point elsewhere. From a small-world network point of view, debates about inter-disciplinarity are simply debates about whether a parameter *p* of the network architecture (the proportion of connections that are long-distance rather than local) is currently too low (too much hermeticism) or too high (insufficient epistemic specialization and competition).

We would make science a small-world network if we trained all people deeply in a disciplinary tradition, but encouraged many of them to spend at least one year, at some point in their careers, in a completely different discipline, getting research training and actually doing some stuff. I don't think this teleported year should be restricted to students starting out; full professors could benefit from it too. The disciplinary combinations could be as bizarre as you like: the way you get small-world-ness in a network is exactly by any long-distance link whatever being possible. We need to provide mechanisms, within academic careers, for these long-distance links to be made, and the extra skills acquired, but not at a cost to the depth of formation in whatever people initially do.

Another mechanism we could employ to maintain small-world-ness is complementary peer review. I don't mean your reviewers should compliment you on how great your paper is, though that would indeed be refreshing. I mean that every paper, during its development, should receive two complementary inputs. The first is a detailed assessment of the methods. You are really only going to get this from people who have practical skills in the right domain. These reviewers are going to be responsible for driving the methodological rigour, in a purely local sense, higher. The second input is from someone in a distant discipline. The point of this review is to say: did you know people have already thought about this kind of problem in this other literature, and they tend to think about it in this way; or, we've actually got a method over here for testing that – why don't you incorporate it? The local review would hold the work to account against the practices of the discipline; but the distant review would also hold the whole paradigm to account against the rest of knowledge. Both of these are important functions. The distant review might also have the by-product benefit of reigning in bad writing and jargon, since authors would have to describe their work comprehensibly enough that the distant reviewer could understand what they were talking about. I think both of these reviews should be published alongside the revised paper, essentially tying in both the disciplinary experts and the neighbouring fields in the common cause of trying to understand the world better.

At present, you tend to get only the local review in specialist journals, which is why these journals are filled with papers that are locally adequate but conceptually derivative, and sometimes only the distant review in inter-disciplinary journals like *Nature* and *Science*, which is why these journals contain more than their fair share of methodologically unsound research. I personally think we should all publish in some big online archive that spans all subject areas, and features both local and distant

[13] Watts, D.J. and Strogatz, S.H. (1998). Collective dynamics of 'small-world' networks. *Nature* 393: 440-2.

open peer-review[14]. And I think, more generally, the onus of peer-review should be moved away from *rejecting* pieces work, and more to collaborating constructively on their development. If reviewers could make the conceptual leap from anonymous controllers of access to a scarce and hence zero-sum resource, to critical co-producers of shared information about the world, surely we would all be better off.

A final point about small-world networks. Although they have clusters, they have no abrupt cluster boundaries. One dense region gives way continuously to another. And this was of course true of early modern Europe too. It was only relatively recently that passports, borders, and border guards existed[15]. Prior to that, there were certainly centres of influence and innovation, each with the ability to generate its own laws and norms, but their influence shaded off with distance like gravitational fields. Many people lived in marches or borderlands with access to more than one political centre. Political centres influenced one another and were involved in higher-level federations. Travel, though physically hard, was not administratively policed. Perhaps there are ways of maintaining discipline, but without boundaries.

[14] Inevitably I think this should be controlled by a self-organizing not-for-profit collective, much as the programming language R is.

[15] Carr, M. (2013). Beyond the border. *History Today,* January.