Hanging On To The Edges:

Waking up and going out to work in the uncanny valley

…giving up a compact disciplinary identity can be very risky.
- Rudolf Stichweh

Film folklore has it that, in Francois Truffaut's film *Tirez sur le pianiste* (1960), Charles Aznavour's central character never actually occupies the centre of the frame. This does not seem to be entirely true; it's what we might call a stylized fact. It is true though that for a protagonist, he certainly spends a lot of time round the edges, down the bottom, or out of shot entirely. It's an apt visual mirror: there's a gap between his great artistic aspirations and the reality of his achievements. He has an air of the current attention always having moved somewhere just slightly different from wherever he is.

My academic life is likewise permeated with a constant sense of slight marginality; of failure to ever get myself quite to the middle of the frame. Perhaps this is just my particularly volatile personal mixture of insecurity and self-importance. Or perhaps *all* academics out there feel that they are more peripheral than everyone else. For example, the Psychology degree programmes I have worked on are periodically audited by the relevant professional body. Part of the body's concern is with how much of the teaching is done by 'real' psychologists. I've always felt vulnerable here: I have slightly less than half a degree in Psychology, my PhD is in Anthropology, and I don't usually publish in journals with *Psychology* in the title. However, asking around, it seems like pretty much all of my colleagues also feel that, for one reason or another, they are not 'real' psychologists either. It turns out that almost nobody has the absolutely canonical history (or feels that they do). Everyone is an off-centre case. And what is more, plenty of my colleagues see *me* as the real psychologist!

Even if everyone feels a bit off-centre, though, I feel a long way off-centre sometimes. I have ended up, by random stumbling around as much as by judgement, living my whole life outside the comforting shelter of any single disciplinary or sub-disciplinary label. Two consequences of this strike me as non-obvious enough to dwell on. The first is the following: it is often easier to talk to colleagues about the components of my work that are further from their concerns, than the components that are rather near to their concerns. The second consequence is that parts of my work are consistently misinterpreted or misremembered as saying something that they really don't quite say. I offer these observations not (I hope) as mere whinges, but as reflections on something interesting about human cognition, about how it tries to impose simple order on a shifting and continuous landscape of information.

§

Back in the 1970s, the Japanese engineer Masahiro Mori introduced the concept of the uncanny valley. The uncanny valley originally described a reliable phenomenon that occurs when robots are made more human-like in appearance and behaviour. As robots move from very un-human-like to a bit more human-like, people's psychological response to them becomes more positive and empathetic; and when the robots are completely indistinguishable from humans, we respond to

them as humans. But there's a dodgy bit in between, in the place where the robots are getting really quite like humans, but occasionally leak cues that betray their artificiality. And in this gap—the uncanny valley—people don't like the robots much at all. They're ok with proper people, and they're ok with good old-fashioned droids with crazy LED eyes and wires hanging out. They really do not like the thing in the middle, 'the thing that should not be'[1].

Early accounts of the uncanny valley related it to our particular conception of 'the human', and our aversion to this key boundary being infiltrated or violated. But it turns out that the uncanny valley is a much more general phenomenon; you get one as you move continuously across many conceptual boundaries, not just the human/non-human one. So the explanation of the uncanny valley phenomenon needs to be rooted in more general ideas about how brains work[2].

Brains are prediction machines. You can't do perception or cognition purely inductively, allowing the information in the incoming sense data to impress knowledge about the world onto a flat blank canvass. You can't do this because there are too many gaps in the immediate data: objects that are partly occluded by other objects; patterns of luminance that could reflect either a change of colour or variation in illumination; retinal images that could represent various combinations of size, shape and distance; ellipses and ambiguities in people's utterances. So brains need to create high-level 'models' of what is out there. These models employ categories that are at least to some extent discrete. Perception and cognition are as much as a case of these models projecting downwards to funnel the sensory input into some kind of structured form, as they are of the incoming information driving upwards to determine what we believe. That's why we are famously susceptible to all kinds of perceptual and cognitive illusions: our internal models can be tricked into firing erroneously in various ways.

The interaction of bottom-up sensory data and top-down internal model is a delicate one. Over the long run, our psychological models are built up from experience and continuously modified by it, so it is the incoming data that determines the model in the end. (Or at least, the incoming data in interaction with inbuilt priors, such as you can't have two objects in any one place at a same time, something can't be both plant and animal, and so on.) Over the short run, though, the internal model provides a lot more weight than any momentary piece of experience. A one-off anomalous scene or object is therefore apt to be reinterpreted as something else, something more compatible with existing model schemas (schemata?). The meeting up of incoming sensory data and internal model schema does not happen at any one stage in the neural hierarchy. Rather, there are many inter-connected processing levels in the brain. Each level passes down to the level below a prediction about what the world is currently like, and hence what data it should be receiving. The level below passes up one of two things: nothing, if the prediction is met and the world is as the model suggests, or an error signal, which effectively says 'No, it can't be that. What I'm getting deviates from that expectation in this particular way'. These prediction error signals, which on this account become a central plank of how brains compute stuff, do two things. Immediately, they cause the higher-level circuit to select another hypothesis about the world ('maybe that thing's not so close, it's just big; try this prediction'), and distally, they cause the model weights to be slightly adjusted so that next time, the circuit won't make the same mistake given the same cues.

How can we relate this to the uncanny valley? Well, all is going well for your brain when it can choose an interpretation of the world that produces no residual error signal at all. What's this? Is it a goat? No, big error signal. Is it a person? Error signal equals zero. My work here is done. But there are some things that are uncanny, and this means, precisely, that they give you a big error signal whichever way you interpret them. Consider the faun. Is it a goat? No, big error signal. Is it a person? No, now I've got a different but equally large error signal. Damn. This thing is just…yucky on my brain.

So, roughly speaking, people dis-prefer stimuli whose error signal can't be got down to a reasonable level, stimuli that defy sensible resolution into a model. Such stimuli are troublesome, incompressible: you can't wrap them into an economical, unified, higher-level conceptual category with zero prediction error and go on your way. You have, effortfully and unwillingly, to settle for lower-level representations of bits of the stimulus, and a kind of 'Warning: Failed to converge' at the higher conceptual level. The brain is nothing if not an economical beast. It doesn't want this kind of clutter hanging around. It needs to avoid it, ignore it, or tidy it up.

§

How does this relate to my experiences as an inter-discipliner, and particularly one who tries to bridge social science and biology? I have an interesting natural experiment to report here. For most of my career, I worked exclusively on humans, and not just any old humans, but (mostly) contemporary British humans. Much of my human work has to do with phenomena such as teenage pregnancy. I have a particular take on this phenomenon. I have argued that women from poor backgrounds who bear children at a young age are not necessarily 'making a mistake' or 'failing to exercise self-control'. They are implementing a behavioural strategy that makes a lot of sense given their relatively short healthy life expectancies (in the poorest English neighbourhoods, women can expect to be in good health only until they are just over 50—why would they wait until they were 40 to start a family?); and given the modest economic returns to delaying childbearing when only low-skilled, low-paid jobs are available[3].

My position is extremely congruent with some critical social-science perspectives. These have also argued that women who bear children young are not committing impulsive individual mistakes but responding, sometimes with some deliberation, to the circumstances in which they find themselves[4]. I'm making exactly the same argument, but I am prone to alluding to the evolutionary concepts of 'adaptive behaviour', 'lifetime reproductive success', 'fitness', and so on. I do this because evolutionary behavioural ecology, the source I drew inspiration from, provides rather useful general expectations (or methods for coming up with expectations) about how individual organisms should respond to their environments.

You might think, naively, that my teenage pregnancy work would fascinate and engage my social science colleagues, and make it easy to build academic bridges to the Faculty of Humanities and Social Sciences. I'm saying a lot of what you guys are saying, and then I am also relating it to a suite of more general concepts from behavioural biology that are already lined up and ready to be investigated. Isn't that exciting? I have given numerous talks of this kind to audiences from epidemiology, public health, sociology, and economics.

I have to tell you that it has not, in the main, gone very well (though there have been a few enjoyable exceptions). People have always been very polite and friendly. The overt hostility to evolutionary approaches that people attribute to social scientists is not, in my experience,

---

[3] Nettle, D. 2010 Dying young and living fast: variation in life history across English neighborhoods. *Behavioral Ecology* 21: 387–395.

[4] See for example Arai, L. (2009). *Teenage Pregnancy: The Making and Unmaking of a Problem.* Bristol: Policy Press.

widespread. What I do experience is simply a slightly embarrassed awkwardness, and no return phone call. Paul Feyerband said that scholars who introduce transgressive ideas find themselves faced, "not with arguments, which they could most likely answer, but with an impenetrable stone wall of entrenched *reactions*"[5]. And the reaction in Britain is, most often, quiet puzzlement, a polite but slightly strained question or two, checking of the train timetable home, and perhaps—I am remarkably neurotic but, yes, I think it might sometimes be there—a very fleeting moment of well-concealed *disgust*[6].

I accept that my rather broad level of analysis simply may not answer the fine-grained questions that social scientists want to answer. But there may be something deeper going on. Perhaps the central argument of my research on humans comes across as 'the thing that cannot be'. The central dogma of evolutionary biology is (allegedly) that genes make bodies and therefore the successful replication of genes is, in the final analysis, the determinant of what goes on. The central dogma of social science is (allegedly) the idea that context determines behaviour[7]. These two dogmas (dogmata?) seem a long way apart, like you would need to choose one or the other. And it may not be clear to my audience which one I have chosen; hence the reaction.

In fact, though, you don't need to choose. Let's move a little way from the evolutionary biology end towards the social science end: *because* the replication of genes is so important, and because the best way of surviving and reproducing is very different in different local environments, evolution has produced creatures that are highly sensitive to the contexts they get put in. Evolution, instead of making a train that can only go down fixed tracks, has made a self-driving autonomous vehicle that can go to a wide range of places according to the landscape it finds itself on. Genes have done this, if you will, because it is in ultimately in their replicatory interest. Thus, the sensitivity to context that they have given us is not a random one, but one structured toward certain needs or goals. On the other hand, it's not like all the possible consequences that could emerge once a whole fleet of self-driving autonomous vehicles start driving around a town were already present in the heads of the engineers who designed the vehicles; of course they weren't. Likewise with genes. So you don't need to choose between genes and context, any more than you need to choose between brake pads and traffic jams.

You might think this kind of position would make everyone happy, and we could all get on famously. But all too often, it seems to fall into the uncanny valley. What does this guy really think? He talks like a social scientist about these contextual factors, then he starts mentioning genetic fitness, and it is as if we suddenly see his battery pack and realize he is really a replicant. If we stick to his actual claims about context—for example that poverty and emotional trauma predict teenage pregnancy—then we already knew those facts, and don't feel any great need for any explanation for them beyond those we already have (indeed, these facts *constitute* a certain kind of explanation for the behaviour). If we focus on all that evolutionary baggage, then we end up with something that gives a rather large prediction error signal whether you try to think of it as a duck *or* as a rabbit. So the view that social science and evolutionary biology can be productively integrated into a synthetic position incorporating information from both is hard to hang on to. This applies in professional academic discussions, but also in interactions with funders, and with the general public.

---

[5] Feyerband, P. (2010 edition). *Against Method: Outline of an Anarchist Theory of Knowledge*. New York: Verso, p. 59.

[6] The anthropologist Mary Douglas importantly linked the emotion of disgust to the violation of category boundaries, and hence, though she did not put it this way, to prediction error signals in the brain and the uncanny valley phenomenon: Douglas, M. (1966). *Purity and Danger: An Analysis of Concepts of Pollution and Taboo.* London: Routledge.

[7] "The idea that context determines behavior is the 'central dogma' of all social sciences from anthropology to sociology, political science, psychology and economics." Glass, T. A. & Bilal, U. (2016). Are neighborhoods causal? Complications arising from the 'stickiness' of ZNA. *Social Science and Medicine* 166: 244–253.

The uncanny valley is steep on both sides. The same talks that produce polite puzzlement in social science departments produce just as much puzzlement, or perhaps even more puzzlement, in Zoology departments. A recent survey of working biologists found only 60% agreeing that what we learn from humans is relevant to understanding other evolved creatures, and the survey probably over-sampled biologists working on humans[8]. The explicit reasons evolutionary biologists give for a queasiness about humans vary from the sensible and practical—like the long generation time and the difficulty of performing true experiments—to the completely bizarre and question-begging—like "humans are influenced by social factors"[9].

Five years ago now, I began to work on European starlings (sensitivity to context thereof, as it happens; the starling is 75 grams of very cool and very sophisticated autonomous vehicle). One consequence of this shift, unsurprisingly, has been greater interest in my work from zoologists and evolutionary biologists. This is the payoff for having climbed out of the uncanny valley up the biological escarpment, and hence generating no troublesome error signals from the fact that our study species wears clothes and watches cable TV.

Strikingly, colleagues from the social sciences and humanities *also* engage with my starling work much more enthusiastically than they do with my human work. You're that starling guy who shows that what you get to eat early in life affects your behaviour when you grow up. How fascinating! I *love* what you do! Actually, I was watching the starlings in my garden, and I was wondering: do they communicate with each other about the location of food/ do they have leaders and followers / do they recognise their siblings / do they return to the place they were hatched? I am continuously pumped with starlingy questions, questions that one might quite comfortably ask of an exotic human society (and to which I usually do not know the answer). This curiosity extends to the general public too. When we are working in the field, people stop their cars to ask what we are doing, why the starling has declined so much, whether starlings can feel pain, and whether it is true that the male starling must offer a nuptial gift of aromatic greenery, placed ceremoniously on the nest, before the female will begin to lay[10].

There is no unease or edge in any of this questioning, just delight. People grasp that there is a form of life over on the other side of the uncanny valley in birdland, and they know that they don't know what it is like. Hence they love meeting someone who has tried to go there and can talk to them about it. That person generates no prediction error signal and poses no embarrassment. Even the insight that the form of life in birdland has features recognisably akin to our own—parenthood and nuptial gifts—is just fine, as long as we are talking about two *different* worlds with some parallels, not a mixed world. The two sides of the uncanny valley can echo each other's landscapes in a kind of aesthetic way, and each can contemplate the other admiringly from afar. Waving from the opposite hillside is a lot easier than living down in the uncanny valley.

§

The same processes that give us the uncanny valley produce a continual loss of the nuanced middle ground in the behavioural sciences. I am prepared to bet that despite my rather careful explanations about autonomous vehicles, the compatibility of genes with sensitivity to context, and so forth, there have been times when someone has said: "We had that Nettle here yesterday. He believes teenage pregnancy is caused by genes rather than the environment!". And it is not just me that has this problem. Mischaracterization of evolutionary approaches to human mind and behaviour—particularly, the claim that such approaches must deny the importance of context—is pervasive,

---

[8] Briga, M. et al. (2017) What have humans done for evolutionary biology? Contributions from genes to populations. Proceedings of the Royal Society B: Biological Sciences 284: 0171164.
[9] See *What we talk about when we talk about biology.*
[10] It is true.

despite repeated and explicit statements to the contrary by their proponents[11]. One has to ask oneself: is it our fault for not being clear, their fault for not listening, or is something more general going on?

Failure of scientists to correctly represent one another's positions is not surprising. Classic work by Frederick Bartlett showed that if you give someone an irreducibly complex shape, have them copy it, than have a second person copy the first, a third person copy the second and so on, the shape soon becomes less complex, and closer to something simply nameable like a cartoon cat or a letter 'A'. Monica Tamariz and Simon Kirby showed that all you need to do to make this loss of nuance happen is to have each participant need to store the shape in memory for a short time[12]. It's the compressive processes of cognition—specifically, to be stored, something has to be expressed in terms of an internal model, the higher-level the better—that inexorably drive the distortion and polarisation of complex ideas.

In another revealing experiment, participants were trained on a mathematical function: they saw one x-value at a time, represented as the width of a bar, and they had to select a corresponding y value[13]. They were given feedback until they got it right. They were then given a set of test trials with no feedback: the x were given, they proposed a corresponding value for y, and this was recorded. The next participant was then trained, not on the original function's x-y pairings, but on x-y pairings that the previous participant had offered during their test trials; their version of the function.

The results are some of the most remarkable I have ever seen. Within a depth of about seven participants, all the functions had become positive linear ones, regardless of what the starting point was*. Curvilinear functions became positive linear ones. Randomly generated patterns became positive linear functions. Remarkably, even negative linear functions became positive linear functions. In short, within a few rememberings and retellings, the image we were left with carried *no information at all* about the stimulus we started out with. It carried information *only* about the kind of pre-existing schemas that people find easiest to learn and remember. This is a very sobering result, given that the ideas worth devoting one's life to are nuanced, layered, and don't fit into convenient pigeon-holes like nature or nurture, genes or environment, individual or society.

§

What happens to ideas that don't fit neatly into any existing schematic paradigm? Mostly, they get turned, like the functions in the function-learning experiment, into something more black-and-white, and less accommodating. (This can give a competitive advantage to ideas that are black-and-white and not very accommodating to start with.) Occasionally, though, a new idea gets a foothold. In effect, the community's internal models get updated enough for there to be a new recognisable category of argument out there. The idea becomes sufficiently stabilised to hold its identity. Often this event is associated with one or two key authors who are gifted, clear and dogged communicators. Even in these cases, though, I suspect the key to success is not so much that the pioneers change the minds of the field at large through their papers and talks, but rather that the pioneers train some good students, and train them well. It's easier to bring about conceptual change by close acquaintance and practical labour together than it is through verbal arguments aimed at strangers. However exactly it happens though, sometimes collective conceptual remodelling does

---

[11] See Kurban, R. & M.G. Haselton (2010). Making hay out of straw: Real and imagined controversies in evolutionary psychology. In J. Barkow (ed.), *Missing the Revolution: Darwinism for Social Scientists*, pp. 149-66. Oxford: Oxford University Press.

[12] Tamariz, M. & Kirby, S. (2015) Culture: Copying, compression, and conventionality. *Cognitive Science* 39: 171–183.

[13] See Griffiths, T. L., Kalish, M. L. & Lewandowsky, S. (2008) Theoretical and empirical evidence for the impact of inductive biases on cultural evolution. *Philosophical Transactions of the Royal Society B: Biological Sciences* 363, 3503–3514.

occur, and a new nameable idea is born. In my field, human evolutionary behavioural science, there were some key examples of this back in the 1980s and 1990s: for example, the idea that our current minds (and bodies) are specialised for life in small-scale Pleistocene societies; or the idea that culture is an inheritance system, parallel to genes. For better or worse, we are still living with the consequences of these new schemas getting created.

Once new schemas get a foothold in middle of uncanny valley, what happens next is predictable but somewhat ironic. They recruit paradigmatic adherents, who are often more dogmatic than the founders, and form mini-cliques with one another. This shows us that most people are unable or unwilling to live out there in the blazing sun of complex and ambiguous phenomena with just their bodies and their native wits to protect them. They crave the epistemic shade provided by a nameable paradigm: a comforting system of assumptions, fellows and sacred texts. They crave this for their own cognitive ease, but also because of the social processes involved. It's easier at an academic party to say 'I am an X' than 'I look at A with a bit of Y and a bit of Z'.

So once a mighty tree takes root out there in the uncanny valley, soon there's a little copse of saplings underneath. And these copses come to have some of the unfortunate properties of the big disciplines they initially hoped to bridge. They have mini-uncanny valleys all around them, between them and neighbouring copses. They acquire purity to be defended. The boundaries of the mini-paradigms are in some ways more troublesome than the macro-boundaries like the social science/biology boundary. They are troublesome because the copses, though verdant, are small enough to have limited resources and odd founder effects. The adherents defend their copse. To switch metaphors, they want to build the brand and increase its market share as much as to find out the truth about the world. They try to suppress competitors, and, as any good capitalist knows, the competitors you most need to suppress are those whose offer is closest to your own. I have been struck, as I have watched mini-paradigms take root in my community, how easily the new adherents abandon the responsibility of pluralism. They don't often feel the need to visit the other copses, except maybe to dismiss them. They don't even engage much in an ongoing way with the broader intellectual sources—ethology, cognitive science, social science—from which their inter-disciplinary mini-discipline springs, and on which its future flourishing depends.

All this means trouble if you want to operate in the uncanny valley between evolutionary biology and social science, but *don't* want to commit exclusively to one of these copses. (You may see the value in all of them, but also recognize the incompleteness of all.) On the macro-scale of engaging distant colleagues, it's hopeless. It's hard enough to get them to understand that there is *one* way of combining aspects of evolutionary biology and aspects of social science. To get them understand that there are *several different* ways of combining them, and these are not interchangeable, that's hard. One risks, as it were: "I've heard about one attempt to be evolutionary about modern humans and I didn't like that, so I assume you're just the same" (and if not that, then: "you evolutionary people don't even agree amongst yourselves"—it's hard to win at this game). Within the tiny community of human evolutionary behavioural scientists, too, it's hard to be an in-betweener. People need you to either be one of their own copse flag-bearers, or else a straw man. You can get castigated for deviations from schemas you never intended to adopt. Many of the most interesting empirical findings—observations that would extend the copses, but are not immediately recognisable as central exemplars of any of them—simply languish. They fill badly needed gaps in the literature.

To distort Oscar Wilde, the only thing worse than having people not cite you is having them cite you. When your work does get picked up in the literature, it's salutary to look carefully and try to identify the claims that those citations are used to support. When I do this exercise, what I mainly see is that those claims are not really the claims I made; they are *somewhat* similar claims that are either more familiar, or more obviously ridiculous, when looked at from the centre of one of the copses. The argument that Ian Rickard, Willem Frankenhuis and I have been developing about why childhood

family conditions have an effect on a wide range of adult outcomes in humans *really isn't quite* the argument that the childhood household furnishes cues about the harshness of the adult environment (and it really isn't quite the argument that childhood stress just messes up your brain, either)[14]. I really have *never claimed* that the reason some people behave less pro-socially and more anti-socially than others is because they are following a 'fast life-history strategy'[15]. And my favourite recent example: Melissa, Clare and I wrote a paper giving an evolutionary take on human obesity[16]. The first draft started with a long background section in which we were very critical of the widespread idea that contemporary humans are obese because fats and sugars were rare in ancestral environments, and thus we have not evolved control mechanisms for saying 'stop' when these things are abundant. That background section didn't make the final edit, because that idea was not the main point of our paper anyway. Within six months of the paper being out, guess what I saw in a draft manuscript I was reviewing? "Sugars and fats were rare in Pleistocene environments, and so humans have not evolved restraint mechanisms to stop them over-eating when these are available (Nettle, Andrews and Bateson 2017)." If you didn't laugh, you would probably cry.

My case is not unusual or worthy of any special consideration. It's just the one on which I have the richest data. It illustrates a more general problem: what people actually say is not what we remember them as having said; and not what it would be more convenient from the point of view of our agendas if they had said. There are two commonplaces from the history of science that fit with these observations. The first is that the troublesome data that will eventually necessitate a change of view often exist in plain sight for decades or even centuries before the change of view happens. The problem is not that the critical observations have not been made. The problem is that the community does not know where to put them in its mental models, so it either ignores them, or mis-represents them as something different from what they really are. The final work of getting them accepted is often a case of compelling communication and social mobilization as much as original discovery. The other commonplace is that new ideas get dismissed as wrong until exactly the point where people say that they are obvious and they always knew them anyway. From the point of view of mental models and prediction errors, we can see what is happening here. Initially, I hear idea X one or two times. I can't assign it to any mental model. It just makes error signal. It must be wrong. But later, I have heard it dozens of times, enough for it to have changed the representational options available in my internal mental model. Here's that thing X again! I've already got a mental model of it, so it seems obvious!

§

The human brain is often described as one of the greatest remaining scientific problems. I think this is true, and not just in the way it is usually meant. One lesson is the need to be extremely clear and do a lot of very patient, cheerful, and if necessary, bloody repetitive signposting. Some of our most successful innovators have been prepared to do this, year in year out, writing the same paper for different audiences, or even the same paper for the same audience, until the penny begins to drop and the idea gets recognised for at least approximately what it is. If like me you are prone to constant shifts of views, banging the drum for the same idea year after year is not something that comes easily. Clear repetition also hardens all too easily into dogmatism and parochialism. Nonetheless, some modest insistence is often necessary.

I would like to close by proposing a scientific innovation: the anti-abstract. This is a short summary of what a paper does not say. I think all papers should have these, published immediately after the

---

[14] Rickard, I. J., Frankenhuis, W. E. & Nettle, D. (2014). Why are childhood family factors associated with timing of maturation? A role for internal prediction. *Perspectives on Psychological Science* 9, 3–15.

[15] One of the places I really don't say this is in: Nettle, D., Colléony, A. & Cockerill, M. (2011). Variation in cooperative behaviour within a single city. *PLoS ONE* 6, e26922.

[16] Nettle, D., Andrews, C. & Bateson, M. 2017 Food insecurity as a driver of obesity in humans: The insurance hypothesis. *Behavioral and Brain Sciences* 40, e105.

conventional abstract. Casual readers could read the paper, spend thirty seconds counting backwards in 7s from 116, and then read the anti-abstract. If they get a big error signal from comparing the anti-abstract to their memory of what the paper said, then they know they need to read the paper again paying special attention. They need to put the effort into creating a finer-grained representation of the paper's claims: they are not what they thought at first pass. The anti-abstract would be very handy as a one-stop-shop for all the likely peer reviewer misinterpretations. Instead of having to take paragraphs of precious Introduction and Discussion laboriously setting out a load of ideas that are not in fact the ideas you want to test, you could simply mention them in the anti-abstract as claims you are not advancing or even considering, and for which your work should not in any circumstances be mistaken.

I am really looking forward to writing anti-abstracts. In fact I might do so and keep them in the file drawer for the day journals start asking for them. I can imagine beginning with the broad theoretical anti-sweep: "Researchers have argued that individual differences in many behaviours can be mapped onto a single underlying continuum of fast versus slow life history strategies. In this paper, we have nothing whatever to say about these arguments". Then there's the anti-summary of methods: "Methods we did not use in this study include the public goods game". And of course, the anti-implications: "Results are not interpreted in terms of the poor lacking self-control"; or "The implication that fertility behaviour is controlled by specific genes in no way arises from our findings". The best thing of all about the anti-abstract is that it gives the perfect chapter and verse defence when you get mischaracterized. Stronger as a defence than "I never said that" is "Look, I actually anti-said that in the anti-abstract". Undermining misapprehensions at source is surely a worthwhile goal. Mutual misconceptions comfort and simplify—the inside of one's prejudices is an easy place to live, after all—but are not, in the end, very useful. The more we clear out misconceptions about what other groups are saying, the more connected our conversations might become to the world itself.