



24. 7. 2025

An Examiner's Report on Peter Tuck's PhD Thesis *Moral Particularism and the Non-Flatness Requirement*

This report has three parts. The first part briefly summarises of Tuck's thesis. The second part provides my evaluation of the quality of the thesis and whether is fit for defence. Finally, the third part outlines some of the philosophical questions I have about the thesis.

Summary

Tuck's thesis has five main chapters. **Chapter 1** first draws the distinction between generalism and particularism. Tuck's understanding is that, according to the former views, there are fixed and general moral truths that can be codified in learnable moral principles, whereas particularists deny such general ethical truths. This distinction is also given a metaphysical and metaethical flavour (as discussed below). After this, the chapter outlines four concerns generalists have had about the particularist position. Generalists have argued that particularists fail to explain important aspects of moral education and intrapersonal justification. They also argue that particularism leads to special pleading in a way that threatens social stability.

This chapter then makes a diagnosis. Whilst many of the previous objections threaten to beg the question against the particularists, they all point to a deeper objection which has become known in the literature as the *flattening objection*. Generalists have argued that there must be an external textured moral reality through which we can navigate ourselves with the help of general principles, whereas the landscape the particularists offer lacks sufficient landmarks. Behind this challenge is claimed to be the legitimate demand that the particularists must be able to explain what makes certain considerations more central for moral deliberation and others more trivial.

The end of the first chapter observes that there can be two kinds of responses to the previous challenge. The weak forms of particularism try to accommodate some forms of generality in the particularist framework, whereas the stronger forms of particularism argue that the local – what is true in particular cases – is enough to explain all moral phenomena.

Chapter 2 outlines and criticises a wide range of the former, weak type of responses. It argues that the particularists cannot incorporate global aspects of morality into their view by relying on default reasons, defeasible generalisations, evaluative generalisations, thick concepts,

virtues, or subjective contexts. These proposals are argued to be defective because, ultimately, they all seem to fail to explain why some considerations are morally central and others trivial in the way that these views suggest. This is because any attempts to ground these distinctions in statistical evidence or default conditions are bound to fail, or so it is argued. What is central ethically, after all, does not need to be more common or more normal in any recognisable sense.

Chapter 3, by contrast, provides a more foundational explanation for why all the responses to the flattening objection discussed in the previous chapter have failed. It argues that the general moral facts which the previous responses would require as a solution to the flattening objection are epistemically inaccessible to us. They would require us to know that some structure or pattern of how different considerations are reasons holds in every possible situation. We would need to know that some moral facts hold always or at least most of the time by default. Or, in other words, we would need to know how certain natural facts are always connected to certain moral facts. The argument then is that this kind of knowledge is impossible for us to obtain. This chapter also contains a discussion of different forms of quietism and of in what way the outlined sceptical view is a version of quietism.

The last two chapters finally argue that the concerns that motivated the generalists to make the flattening objection to particularism can be dissolved by a positive, constructive particularist account which explains how moral knowledge (including the knowledge of what is central and what is trivial) is possible in individual cases. Of these chapters, **Chapter 4** first focuses on so-called case-based reasoning where individuals can, through experience, acquire a skill of moral deliberation that extends also to new cases. This popular proposal, however, faces two constraints. It fails to explain how individuals can originally acquire knowledge in the very first cases so that they can develop their skill further, and also it is not always clear that the new cases will be comparable to the ones that the agent has already experienced. This means that case-based reasoning can be only a partial solution to the problem.

The final **Chapter 5** then attempts to provide the missing piece of the epistemic puzzle. It discusses moral perception and the question of whether we have some form of a ‘reason sense’ – a rational ability to see what reasons we have in particular situations especially assuming some form of nonnaturalism about the reasons relations. This chapter argues against the widely accepted idea that we can come to practical conclusions by weighing individual reasons that are present in the situation. The argument is that this view would commit particularists to objectionable forms of atomism and especially to the idea that the valence and strengths of reasons transfers from one situation to another. The end of the chapter also sketches another method according to which all we need to know about reasons can be acquired by comparing the comparative weights of reasons in the particular context where this does not require knowing the weight of the reason in isolation. The analogy used here is that we can similarly see who the tallest person in a room is without knowing the heights of the persons present. The claim then is that by explaining how we can compare the weights of the reasons in individual situations the flattening objection dissolves – we get a view of local centrality and triviality which is acceptable for the particularist and so the concerns that motivated the generalist objections go away.

Evaluation

First of all, I should say that I strongly believe that this thesis is fit for defence – I am fairly certain that will pass the examination process. This is because the thesis has several important virtues:

- Firstly, Tuck has identified an important, focused, and central topic in moral philosophy. The debate between generalists and particularists matters for our understanding of morality and the flattening objection is a fundamental issue for trying to come to a reasoned conclusion within this debate.
- Secondly, Tuck shows both comprehensive and detailed knowledge of the literature on the topic and very good understanding of the subject and the researched materials. He shows mastery and thorough knowledge of the central writings on the topic, and he also covers the most recent literature on the topic too. All this literature is also explained in a thorough way by both paraphrasing and providing the key quotations. The thesis also shows some wider knowledge of the field of metaethics more broadly.
- Thirdly, and relatedly, in many cases the way in which Tuck brings together different arguments on a given topic is illuminating. It provides a nice sense of the whole logical space, and also of the similarities and differences between different views and arguments. This was especially evident, for example, in the explanation of the four objections in Chapter 1, in the discussion of the eight responses in chapter 2, in the discussion of the different forms of quietism in chapter 3, and in the discussions of case-based reasoning, moral perception, and weighing reasons in the last two chapters. As overviews of these topics, the relevant parts of these chapters are of a publishable quality.
- Fourthly, I also thought that the thesis as a whole had a nice unified and coherent narrative going through it. A key objection to particularism has been identified, existing responses to it are explained and criticised, and finally particularism is defended against the objection by trying to dissolve the concerns behind the objection with a positive account of local moral knowledge.
- Fifthly, the thesis also illustrated an ability for sophisticated independent philosophical theorising and for critically engaging with the discussed materials. It showed very good analytical skills in the way that it broke down arguments from the source texts into premises and conclusions, and oftentimes good critical points were made against various views and responses to the objection.
- Sixthly, the thesis was well presented. It had a clear structure that was transparently signposted with section headers and guidance in the text too. The use of relatively short and focused sub-sections was helpful in this respect. Overall, the writing was clear and easily readable as well, and the essay did use the Harvard system of referencing correctly.

Thus, on the basis of these virtues, I do think that the overall quality of the thesis is pretty high - comparable to many PhD theses in moral philosophy that in my experience have passed with

praise in the UK, the US, and Scandinavia. I also believe that several parts of the thesis contain material that is, with some work, suitable for publication in high quality philosophical journals. For all these reasons, I do believe that the thesis is fit for defence.

Before I go to some of the philosophical questions the thesis raises in my mind, I do want to raise one feature of the thesis due to which the thesis is perhaps not of the very highest quality. I think that there is a slight difference between the quality of the exposition and the original philosophical argumentation (including both the objections to other views and the explanation of the positive proposal at the end). This also concerns the balance of how much space is dedicated to both. I slightly worry that, whilst the exposition of other people's views and arguments is often detailed, careful and even insightful, by comparison the places where original critical arguments are made are shorter and sparser. Here I have in mind, for example, the objections to the initial objections to particularism in sections 1.2-1.4, the objections to the accommodating solutions in chapter 2, the objections to atomism and weighing reasons in section 5.6 and especially the positive proposal in section 5.7.

These more original critical discussions tend to be shorter than the expositions of the targets, and the arguments denser and more abstract. Even the writing at times gets less clear here. Thus, overall, one thing that would have improved the thesis would have been a bit more patient and clearer at these critical junctures. Using even more concrete examples at these points, for example, would have helped, as would have at times breaking down the arguments in syllogisms. In addition to helping the reader to evaluate the original arguments, this would have shifted the balance more towards original thinking.

Philosophical Questions

In this last part of this report, I want to explain some of the more substantial philosophical concerns I have about the thesis. I have focused here on the main questions I have about each chapter of the thesis. These are intended for discussion in the defence of the thesis (and they by no means are intended to downplay the quality of the thesis).

Chapter 1

In the way things are set up in chapter 1, the distinction between particularism and generalism is presented as a metaethical difference that is a metaphysical and ontological disagreement about the existence of something. I was not certain, however, about why we should accept this metaphysical way of drawing the distinction. For one, as it emerges throughout the thesis, both particularists and many generalists too accept non-naturalist realism and so the disagreement cannot be over the existence of properties. It also cannot be, as it is sometimes claimed, about the existence of principles as all kinds of principles exist (the principle 'Do a jumping jack at noon!' exists even if it might not be a true or a correct moral principle nor one adopted by anyone). So, it is not clear to me the existence of what the particularists and generalists could disagree about.

What is the disagreement then about? It could be about whether any learnable moral principles could be true or correct. This, however, to me seems like a first-order ethical question that concerns which actions are right and wrong or what reasons we have. It could also be a disagreement about what role learnable principles play in moral education or moral cognition. This would be a metaethical question, but it would be a question in moral psychology and epistemology rather than in metaphysics. So, it would be nice to get some clarity of why we should think of particularism in the ontological way.

The second, related thing I would like clarity in this chapter is the diagnosis that is drawn concerning the first set of objections to particularism (based on moral education, social order, the practice of justification, and the like). The diagnosis seems to be that, even if these objections in themselves fail mainly due to concerns about question-begging, the real concern behind these objections is the worry that the moral reality must have some kind of a structure, which we need to be able to understand to be moral agents.

But, I am not sure I am yet convinced of this diagnosis, and I did find the discussion of, for example, Hooker's objection on pages 12-14 somewhat difficult to follow. I am not sure that Hooker's concern at this point is that moral reality ultimately must have some structure but rather. Rather, his concern seems to be merely that it is really difficult to see how moral education could work without some reliance on some kind of principles no matter how rough and approximate they are (given how such preliminary principles play a role in learning of most skills). This point seems to be a genuine point about the role of principles in moral education and moral psychology rather than a deep point about moral ontology. I think the same goes for the discussion of social stability and the role of principle-based justification plays in it. I would also like to hear more about how these arguments beg the question as it is often claim.

Chapter 2

In this chapter, my main concern is about the objections to the default reasons solution (though something similar goes also for some of the other responses too). The default reasons view itself is explained very well – the idea that the default reasons are considerations that are in some way switched on by default to count in favour of the relevant action whereas when these considerations are not reasons we need an explanation in terms of some disabling conditions. The claim was that this machinery could help the particularists to explain the intuition behind the flattening objection that some considerations are more central than others in a way that at least seems principled.

The chapter nicely explains how this raises an explanatory challenge – we need some explanation of why some considerations are default reasons and others are not. I also like the explanation of the arguments to the conclusion that in giving this explanation the particularists cannot rely on statistical generalisations or defeasible generalisations based on privileged conditions. I, however, worry that the particularists taking this option would think that the question of what considerations are default reasons is an internal first-order normative question.

Thus, in case of pain being a default reason against an action, we need to think about the pain and essence of pain to see that it is by default a reason-giving property. As far as we understand what pain is – what the essence of pain is, we can arguably understand that it by default gives us a reason. Here it would help to look at the work of Stephanie Leary.

Chapter 3

In this chapter, there is an original argument to the conclusion that the responses covered in the previous chapter could have never worked. The crux of this argument is that that type of responses would need to rely on facts that are epistemically inaccessible for us. The core of this argument seems to be that in order to find the kind of structure generalists are defending (that certain considerations are default-reasons and thus more central, whereas others are not and thus they are more trivial) we would need to know every single situation anyone has ever been in, or could ever be in, which obviously is impossible (76).

There are several reasons why I am not sure that this argument is successful, but the main reason is that I am not sure that the generalist position would require such knowledge of all possible situations. First, generalists do not seem to be committed to having exhaustive lists of the morally central considerations (Ross himself thought that his list of prima facie duties is open ended, and the quote from Hooker does not seem to support that idea either). Second, according to generalists the default reasons need not be always genuine reasons and the non-default reasons can sometimes be strong reasons (depending on the disablers, enablers, attenuators, and intensifiers) as long as there is some hope of this complexity being captured with learnable principles. Finally, we seem to be justified in accepting some universal principles for example in science even if we do not know all cases. So, I am not quite sure I see where the high epistemic threshold for generalists comes from in this chapter.

Chapter 4

I did find the discussion of case-based machine-learning fascinating. I also agreed with the concern that the model relies on new cases being sufficient similar to the old ones and also that there is a problem on how we get the process started. There is also a quicker way to see this in the case of machine-learning. In that case, in the training data (when we teach the machine, for example, to detect cancerous growths from scans) the machine that is building its model is rewarded and punished based on getting the diagnosis right which leads to the machine refining the model in order to achieve maximum rewards. But, this process assumes that there are already known cases of getting it right that can be used to train the machine.

I did have one general question about this model though which I was not sure was mentioned. The concern I have is that, even if many particularists have found this model appealing and compatible with their view, I am not sure how confident we should be that the machines (and perhaps people too when they learn in this way) are not in this case building a set of fine-grained principles with weights to ground their decision-making. There seems to be no reason why the kinds of principles that a machine—or we—internalize through the previous kind of

learning process couldn't remain outside our conscious awareness, even if those principles are physically embedded in the system's structure (like nodes in a machine or neurons in the brain)

Chapter 5

With regards to this chapter, I first had a few minor questions. For example, it is not clear to me that particularists need to or should be buck-passers and thus reduce right and wrong, and good and bad to reasons. For example, I am not sure that Dancy is a buck-passer (he seems to think that the same natural grounds ground both reasons and right/wrong and good/bad even if he does not want to reduce the latter to the former). I think the same particularist and epistemic claims could in any case be made of whatever normative properties and relations particularists take to be basic. Second, at times, there is a bit of a slide between metaphysical questions (if anything explains that something is a reason) and epistemic questions (how we know whether something is a reason).

I also found the explanation of why the epistemic challenge is so difficult for the particularist (on page 119-121) especially insightful. It really does seem to follow from the fact that many particularists tend to be nonnaturalists but not transcendental realists. This seems to push them towards moral or reason perception, but given nonnaturalism the perceptual model becomes more questionable (given that perception seems to require causation but nonnaturalism seems to rule it out). I was also very sympathetic to many of the objections to the process of weighing reasons when this is understood atomistically and combined with the assumption that the reasons retain their salience and strength across contexts.

What I was slightly disappointed by and what I would like to hear much more about is the alternative model provided in the final section 5.7. Firstly, the view that comparative weights are in some way an essential element of reasons has been developed and defended most recently by Justin Snedegar in his book *Contrastive Reasons*, which would have provided more positive resources for developing the view here in a particularist friendly way.

Secondly, and most importantly, I'm not sure I see how this section addresses the epistemic problem that was promised to be solved and how it also thus could dissolve the flattening objection. The positive view here relies on an analogy. If I see three men who are 180cm, 185cm, and 190cm, I can just see who is the tallest amongst them without knowing their individual heights – I can just compare how tall they are in perception directly. The claim seems to be that, likewise, I can just compare what salience reasons have and come to see which reason is the strongest/most demanding one without knowing the individual strengths of the reasons in the case. And, this local knowledge about the case is all the basic moral knowledge there is or is ever needed. In a way it also shows that particularists can locally accept that some considerations are more central than others.

My worry is that the sparse presentation of this positive proposal ends up raising more questions than it answers. Here are just some of these questions:

- Before we can compare the strengths of any reasons, presumably we need to know which considerations are reasons in the given case. For particularists, any consideration whatsoever can be a reason and presumably there are millions of considerations present in any given case. So, how do we know which considerations are reasons to be compared?
- The height analogy seems to be different in significant ways. The tallness ordering between the men is an internal relation – it only depends on the heights of the men and nothing else at all (there cannot be external disablers or anything like that). Because of this, the tallness ordering supervenes on the heights. By contrast, the reasons relations are not internal relations according to particularists and so they depend also on other things than the considerations that are the reasons. External factors can, after all, be disablers, enablers, attenuators, and intensifiers. This means that the strength-ordering of the reasons does not supervene only on the considerations that are reasons – the context makes a difference too. This means that the comparison requires a lot of additional extra knowledge of the situation (and there can always be a disabler lurking about). So, again, how can we come to have all this knowledge?
- Relatedly, according to the non-naturalist and anti-transcendental picture, in order to know the relative weights of the reasons it is not enough to perceive and know all the natural features of the situation, the grounding factors, but rather we also need to have some kind of an access to the grounded non-natural reason-relations. Yet, the mere fact that we can compare strengths of reasons does not yet seem to suffice as an explanation of how we could ever do this if non-naturalism were true. Here, for example, Matthew Bedke's accidentality worries seem to be relevant.
- Finally, even if the picture was an accurate representation of local moral knowledge, it is still not quite clear to me that this dissolves the concerns behind the flattening objection fully.

I should say that the candidate probably has many things to say in response to these concerns – or, maybe there is an explanation of why we do not need to respond to these questions. However, it would have been important to address at least some of these questions more explicitly in the thesis itself. As it stands, we only really get a page and a half of the most interesting positive proposal whereas we get much more depth in the discussions of many of the views that are rejected.

These are my main substantial philosophical concerns, though I also have some more detailed smaller questions that we may get to in the defence. But, as I emphasised above, despite these philosophical concerns I do think that the thesis is of a good quality and so is fit for a defence.

Jussi Suikkanen

Department of Philosophy, University of Birmingham