

To Dr Stephen Barker,

Please find enclosed a summary of the minor corrections made to my thesis – ‘The Better Angels of Our Nature’ – in response to the joint examiners’ report (received 27/09/23). In addition to the four mandatory corrections listed in the report, I have also made some further alterations. These further alterations are partly a case of me responding to various comments included in the report and partly a case of me rectifying various typos that I noticed upon re-reading the thesis in preparation for the viva.

Let me begin by restating the four mandatory corrections listed in the report:

- 1) We ask Mr Mitchell to say a little more about exactly what it is that is being constructed in his constructivism. Moral truth (itself, the metaphysical thing)? *Provable* moral truth? *Knowable* moral truth? Or any of these, with “practical” for “moral”? And what, bearing in mind the analogy with mathematical constructivism, in which the specification of a *particular* proof-procedure is everything, is the procedure that he favours whereby constructions are constructed?
- 2) What is the relationship between Mr Mitchell’s constructivism and his writing about peak experiences? Is the idea that those peak experiences give us a procedure for construction? If so, what is this procedure?
- 3) What is the relationship between what Mr Mitchell says about peak experiences, and what he says about “ethereal love”?
- 4) What is going on with Mr Mitchell’s deployment of Achilles? Is this seen as a peak experience? If not, what *is* it seen as? More clarity here would be useful.

In addition to the above, there was also mention of a possible direction that I might take the thesis in: “A fifth question is tempting, but too large and too nebulous to be one of the ones that we would like Mr Mitchell to address ... What about Iris Murdoch?”

I have politely declined to make this fifth, optional correction to the thesis; I feel it necessary to offer some brief explanation for this. I am not pursuing an academic career: I am not applying for any academic jobs, nor am I planning on trying to get any part of this thesis published (I am going to make it open access under a Creative Commons license). I will soon be starting a job in the Civil Service.

The reason I mention this is that most PhD students seem to use their thesis as a launchpad for future academic publications, such as by converting individual thesis

chapters into standalone papers and then submitting them to journals. If I were doing this, then it might be worth going back and making major substantive revisions to Chapter 4 of the thesis (the one dealing with peak experiences and meditation) with an eye towards establishing a solid foundation for future academic publications. Engaging with Iris Murdoch's work vis-à-vis peak experiences – and doing so in a way that was actually interesting and thoughtful and did justice to the nuances of her writing on this topic – would involve making major substantive changes of this kind. However, given that I am pursuing a career totally unrelated to academic philosophy, I think that investing time into substantially rewriting the thesis is not, at this point, the right option for me.

Below, I go through the changes that I have made to the thesis regarding the four mandatory corrections.

Correction (1):

I do not want this cover letter to become too verbose and long-winded. As Geoffrey Chaucer says in the *Canterbury Tales*, 'A windbag is to God abominable.' However, given that a considerable amount of time has passed since the date of the viva, I think it is worth offering a brief recap of Humean constructivism before specifying the page ranges where I have made alterations.

Humean constructivism is a thesis about the status of ethical and normative truth, not a decision-procedure. It offers a general view of normative claims about reasons for action. It explains what a creature must do in order to count as a valuer in the first place and how standards of correctness in the normative domain are generated by the attitude of valuing. It is concerned with the *nature* of normative truth, not with its *content*. (By contrast, the Rawlsian original position, Thomas Scanlon's moral contractualism, and Jürgen Habermas' ethics of discussion are, for example, all constructivist positions within the field of substantive normative ethical theory, and are concerned with the content of normative truth.)

Humean constructivism locates the moral within the domain of the normative, so an agent's moral reasons are just a subset of the full set of their normative practical reasons more generally. Under the view, moral reasons and normative reasons are *not* 'normatively independent' and do *not* belong to 'non-overlapping parts of the normative domain' (Forcehimes and Semrau 2018: 700).

In the thesis, I pitch Humean constructivism as a thesis about the conditions under which ascriptions of moral reasons are true. In focusing on *ascriptions* of reasons rather than *reasons themselves*, I am trying to reduce the need to commit to controversial assumptions regarding the metaphysical nature of reasons. Not much hinges on this, and the position can easily be reframed to focus on reasons themselves.

I do not claim in the thesis that Humean constructivism is true. I do not claim that it is more plausible than theoretical rivals, such as naturalistic versions of meta-ethical moral realism. Rather, I try to assess whether the theory does a good job of accounting for one key facet of the precious gemstone that is ethical truth: the facet that concerns truths about our moral reasons. The idea is that *if* it does a good job of accounting for this facet, then it might also do a good job of accounting for other facets of the ethical gemstone, such as that which relates to truths about our moral obligations.

The comparison with mathematical constructivism is not one that I draw in the thesis. It is a comparison that Lenman and Shemmer (2012) draw in their seminal textbook on constructivism, but that is because they operate with a proceduralist characterisation of constructivism, according to which ‘the central, distinctively constructivist claim is the metaphysical claim about the order of determination moving from procedure to facts and not vice versa’ (ibid.: 3).

The novelty of *Humean* constructivism specifically is that it repudiates the idea of a hypothetical procedure for generating truths about what moral reasons one possesses. Under HC, there *is* no procedure. Any ‘procedure’ would just be an epistemic tool or heuristic device for discovering what the entailments of the practical point of view are, but *it is the entailments themselves, not our way of finding out about them, that constitute ethical/normative truth.*

The notion of ‘entailment’ that I’m appealing to here is ‘logical and instrumental entailment from within the practical point of view’. As shorthand, and to avoid confusion, I refer to this form of entailment as a ‘method’ for determining truths about what normative reasons one possesses (page 77 of the PDF of the thesis). So, under HC, there is a method for construction, not a procedure.

The bare-bones requirements that govern practical reasoning – e.g. the principles of logical entailment and means-end coherence – do not amount to substantive commitments. Rather, they are features that are constitutive of the very attitude of valuing, of making judgments with normative content. The task of specifying what is constitutive of the attitude of valuing is an exercise in descriptive philosophical analysis rather than substantive normative theorising. If one ignores the logical and instrumental requirements that govern practical reasoning in full consciousness of what one is doing, then one is not making a mistake about a normative matter; rather, one is simply failing to value. (I discuss this notion of logical and instrumental entailment on pages 21–22, 28–29).

I have added six new paragraphs towards the end of Chapter 1 clarifying what sort of moral truth the constructivist attempts to construct (pages 51–53 of the PDF, beginning with the sentence: ‘The constructivist maintains that moral judgments are attitudes of valuing rather than ordinary, robust beliefs ...’).

Given that HC does not involve a procedure, it is a bit tricky to specify something resembling a particular proof-procedure in mathematics. (It tends to be substantive normative constructivist views, such as Thomas Scanlon's moral contractualism, that claim that the correct application of a particular deliberative procedure leads to moral truth; Humean constructivism, as a meta-ethical/meta-normative view, does not do this.)

However, I have done my best to answer the question contained in the correction by specifying what, in the case of a *particular* moral judgment rendered by a *particular* agent, the substantive truth condition for that moral judgment is (pages 51–53). I think that §2.1 of Chapter 2 – especially my example of 'Hermann' – may also be relevant here as I go through the various ways in which, under constructivism, agents' moral judgments can be objectively wrong. I think my clearest statement of the Humean constructivist 'method' itself is on page 77.

Some of the points I discuss further down in this cover letter might also be of some relevance, especially regarding the distinction between normative and motivating reasons and the status of so-called 'external moral reasons' under a constructivist view of things.

Corrections (2) and (3):

I will discuss corrections (2) and (3) together as they are both structural points about the links between topics in the thesis. Just as a reminder, the thesis is structured as a collection of essays treating only loosely connected topics. I do try to identify some connecting threads running through the various topics, but these threads hang much more loosely than in the average PhD thesis.

I have to say that it would have been enormously helpful to have been forwarded the examiners' comments before the viva exam rather than after it. I don't know what the rules say about this – perhaps PhD candidates are not allowed to see the examiners' comments until after the exam. (I am not talking about the joint report; I just mean the comments on the thesis itself.)

During the viva, I did not grasp what the External Examiner was getting at in asking me about the links between constructivism and peak experiences. But now that I have read through Sophie's comments on the thesis, I think I understand the relevant point that I failed to communicate.

Under Humean constructivism, the strength of a particular normative practical reason for action that an agent possesses co-varies with the strength of whatever evaluations in that agent's subjective motivational set go towards determining the truth about that particular reason (some subjectivists, e.g. David Sobel in his book *From Valuing to Value*, conceptualise normative reasons in terms of weight rather than strength). So, other things being equal, the stronger one's relevant evaluations, the stronger one's reasons.

'Other things' include the relevant non-normative facts, which often have a decisive role to play in determining truths about what reasons one possesses (I delve into this point in §3.8). Another way of putting this is that constructivism involves proportionalism with respect to our normative practical reasons (of which our moral reasons for action are a subset).

This does not mean that just because an agent strongly values doing *X* they will automatically have a strong reason to *X*. And nor does it, I think, collapse the distinction between justifying or normative reasons (i.e. reasons that bear on what an agent *ought* to do) and motivating reasons (i.e. an agent's actual motivation for an action).

Under constructivism, an agent's actual mental states do not directly determine truths about what normative reasons they possess. Normative reasons are not simply equivalent to the existing motives of an agent. The link is indirect: *X* is a normative reason for *A* to *Y* if *X* is logically and instrumentally entailed from the full set of attitudes of valuing held by *A* – in either an occurrent or dispositional form – together with any non-normative facts that have a bearing on those evaluations.

Agents often suffer failures of coherence with respect to their own states of valuing. They can fail to notice how certain evaluations are in tension with one another. They can fail to grasp the full implications of certain evaluations. They can fail to keep in view all of the evaluations that they do, in fact, hold. Agents are often ignorant of the relevant non-normative facts. (I discuss these points in §2.1 of the thesis.)

So, under constructivism, the existence of a normative reason for an agent to do *X* may not at all *explain* their having done *X* (though that agent might have a reason for doing *X* that *both* justifies and motivates doing *X*).

The links that I sketch out at the beginning of Chapter 4 between constructivism and peak experiences don't make much sense without this co-variance piece in play. For example, in the introduction to Chapter 4, I outline the following motive for wanting to segue into the topic of peak experiences (this motive is *ad hoc*; I do also discuss a more theory-neutral one):

I want to explore whether introducing a peak-experience-based element of idealisation into the constructivist method is a plausible means of bolstering or 'shoring up' the strength of our core moral reasons. The rough idea is that peak experiences strengthen our evaluations that are more pro-social, cooperative, altruistic, and other-regarding in orientation, and if we draw on these 'peak' evaluations in the construction of ethical truth, then this would – depending on what the relevant non-normative facts are – probably translate into our having stronger moral reasons in general.

In response to this, Sophie writes, “what does it mean to speak of ‘stronger moral reasons’? *Reasons* don’t change their strength. It’s *motivations* that do that. This is a grammatical point, not a metaphysical one.”

I could be wrong about this, but I think that for the Humean constructivist it *is* a metaphysical point. The strength (or weight, depending on how you choose to conceptualise it) of a given agent’s moral reasons for action does co-vary, *other things being equal*, with the strength of whatever evaluations in that agent’s subjective motivational set go towards determining truths about what reasons they possess.

So, certain facts will count in favour of certain actions more strongly for a given agent depending (indirectly) on the strength of certain evaluations in that agent’s evaluative standpoint.

I ran multiple drafts of this chapter, Chapter 4, past both of my supervisors, and neither of them flagged to me that there was anything unusually controversial about the idea that, under a meta-ethical subjectivist theory, one’s moral or normative reasons for action might vary in strength or weight depending on the composition of one’s evaluative standpoint. I fully understand, however, that this view is open to many objections. Probably the most obvious of them is the amoralist objection, which I evaluate in §3.3, §3.4, and §3.5 of the thesis.

I think one potential point of confusion here is that what Sophie calls ‘external moral reasons’, many subjectivists would instead call ‘moral demands’. Under constructivism, all moral reasons are internal reasons in the sense that they are anchored to (but not directly determined by) one’s subjective motivational set.

Sophie writes: “Intuitively my moral reason to help to feed the victims of famine is entirely independent of any inclination or disinclination that I may feel to do so. To think that my reason to feed them is *constructed* out of my existing subjective motivational set (my S) seems to be simply wrong-headed.”

Fair enough. Most people probably share that intuition, to a greater or lesser degree. But the constructivist denies that there are any stance-independent facts about what you have moral reason to do that are *entirely* disconnected from your natural sources of motivation: your conative states. And positing such facts raises its own set of problems – e.g. Christine Korsgaard’s practical open question argument, which I discuss in §2.2.

To run with Sophie’s example, the constructivist could agree that everyone is subject to a moral *demand* to aid the victims of famine, irrespective of the composition of their motivational sets. But whether or not a given agent actually possesses a normative practical reason to aid the victims of famine depends – indirectly and depending on what

the relevant non-normative facts are – on the composition of their particular motivational set.

For the constructivist, moral demands such as ‘everyone has a duty to aid the victims of famine’ are always rendered with tacit reference to a particular evaluative standpoint or combination of standpoints – e.g. the standpoints of those of us who have some moral feeling (which, as I argue in §3.7, is virtually all of us).

One cannot step outside of one’s own evaluative standpoint – or of *any* evaluative standpoint – and occupy some Archimedean point from which to ‘objectively’ pass such judgments. Such judgments implicitly rely on other evaluative judgments; it is other evaluative judgments which set the standard against which such judgments can be passed.

Of course, one is free to think that this view is completely crazy. Perhaps it is completely crazy. But even crazy views can be interesting and worth exploring – and quite elegant, in their own way.

Anyway, I did a poor job of verbally communicating the preceding points during the viva. I think that when these points are borne in mind, the motives that I set out at the beginning of Chapter 4 for segueing to the topic of peak experiences make more sense. I have edited and significantly lengthened the introduction of Chapter 4 (**pages 127–131**) in order to make things clearer and provide relevant context.

Regarding the connection between the topic of peak experiences and the topic of ethereal love – I have rewritten §4.5 and §4.8 of the thesis.

In **§4.5 (‘Whither Enlightened Constructivism? The Relevance of Meditation Practice to the Constructivist Method’)**, I try to make the general case for segueing from a discussion of peak experiences to a discussion of meditative experiences. (Page 167 onwards).

In **§4.8 (‘Experiments in Constructivism: Introducing ‘Meditative Constructivism’)**, I try to give a sense of why I am particularly interested in meditative experiences of ‘ethereal love’, which is a term I use to refer to the state of mind that is cultivated through loving-kindness meditation. (Page 201 onwards).

I think one potential point of confusion here is that I was not clear enough about my argumentative stance going into the discussion of meditative experiences. My stance is exploratory rather than evangelising. I am trying to explore *whether* meditative experiences are ethically insightful and relevant to the theory of value that is Humean constructivism; I am not claiming that they definitely are. Indeed, the conclusions that I come to at the end

of the thesis are decidedly pessimistic: peak experiences and meditative experiences do not have a role to play in the construction of truths about our moral reasons.

Anyway, this cover letter is becoming frightfully long, so let me move on to the final correction ... I'm reminded of a quote from *Don Quixote*: "'Proceed friend,' said the ecclesiastic, at this period; 'for you are going the way with your tale not to stop till you come to the other world.'"

Correction (4):

It was never my intention to suggest that Achilles undergoes a peak experience during his immortal quarrel with Agamemnon in Book One of the *Iliad*. The purpose of the example was to illustrate the far narrower point about why we should defer to the wiser ethical perspective of our peak selves. I was trying to give a sense, through the use of a colourful example, of why we should care what our peak selves have to say when it comes to moral matters. I was not attempting to derive, in an exegetical manner, any moral truths from the text of the *Iliad*.

In the example, the grey-eyed goddess Athena is a stand-in for Achilles' peak self. This ties in with my discussion of Firth's (1952) Ideal Observer theory, where one concern is that the perspective of the Ideal Observers is so far removed from our own as to be profoundly alienating. I liken Firth's Ideal Observers to Spock from *Star Trek*; I try to show how the hypothetical beings that I posit, our *peak selves*, occupy a perspective that is far more relevant and familiar to us (e.g. I suggest that there is an experiential bridge connecting one's actual self to one's hypothetical, peak self).

I have edited and reworded this section to clarify the purpose of the Achilles example, which is what the correction asks for. I have also updated the translation of the *Iliad* that I use. In hindsight, I'm not quite sure what possessed me to use the dusty old museum-piece of a translation by E.V. Rieu. I have updated it to a more recent verse translation by the Homeric scholar Professor Barry Powell. The relevant changes are to be found on pages 150–153.

Some further, minor alterations:

Page 2. Rewording of the following sentence: "it will involve a strange fusion of topics: meta-ethics, moral psychology, the phenomenology of peak experiences and meditative experiences, etc." The revised sentence in the corrected version of the thesis now reads: "it will involve a strange fusion of topics: meta-ethics, moral psychology, and the phenomenology of peak experiences and meditative experiences."

Page 6. Rewording of the following sentence: "The cave system stands for the network of views in ethics that fly under the name of 'constructivism'". The revised sentence now

reads: “The cave system stands for the network of views in ethics that fall under the heading of ‘constructivism’”. Similar change made to a nearly identical sentence on page 218.

*The above changes are very minor. The other changes listed on this page are similarly minor. As such, **I will only quote the revised sentences as they now appear in the corrected version of the thesis** – otherwise, I fear that this cover letter will grow pointlessly long and will become confusing to navigate.*

Page 8. Reworded sentence: “Broadly speaking, the first half of this thesis focuses on traditional, centre-of-the-fairway meta-ethical issues: laying out a particular theory of ethical truth, contrasting it with rival positions, and evaluating its costs and benefits.”

Page 11. Reworded sentence: “An example of a hypothetical amoralist would be a hypothetical agent who is stipulated to be ideally coherent, informed, and imaginative, say, but who, in virtue of his perverse evaluations, has moral reasons that seem utterly perverse and deranged from our perspective.”

Page 39. Deletion of unnecessary comma: “He thinks that by living this way he is expunging not only his own sins but the sins of all the saucy hedonists and libertines out there.”

Page 40. Insertion of missing comma: “But, as Street (2012: 49) points out, ‘as a matter of fact we do have an option here ... the proper answer is not to say yes, but rather to reject the question at hand as ill-formulated’”.

Page 73. Addition of missing paragraph indentation (final paragraph of Chapter 2).

Page 81. Reworded sentence: “Socrates argued that the latter option seemed to make divine approval arbitrary, thereby undermining the goodness of the gods.”

Pages 92–92. I have replaced the term ‘ibid.’ with the name and date of the relevant author in two citations on this page.

Page 149. Deletion of two unnecessary commas: “Instead of asking what one *in fact* valued during one’s most intense peak experience ...”

Page 162. Insertion of missing comma: “I argued in §4.2 that there is a higher-order stability involved in changes of this kind during a peak experience: it is not just that such changes take place, but that we *approve* of their taking place ...”

Page 166. I have removed a section of a long quote from Steven Pinker (2012: 157) that is displayed in a footnote at the bottom of page 144 of the thesis. The reason for this is that

part of the quote, as the External Examiner points out, contains a historical inaccuracy – a rare slip on Pinker’s part, given that he is usually a meticulous and careful scholar.

Page 206. I have removed a sentence that contained a colloquial phrase that was inappropriate for the context. The sentence used to read: “Are meditative experiences more ethically insightful, in general, than regular, bog-standard peak experiences?”

Page 175. Rewording of a sentence to remove the misleading use of a technical term (“unconditioned reality”). It now reads (I am quoting the paragraph for context): “As he sits there on the park bench, looking down at the black tangled mass of the root, he is suddenly shunted into experiencing the world as a raw, immediate, amorphous, Heraclitean flow of phenomena. The familiar, comfortable concepts of everyday life have melted away. Roquentin’s internal self-chatter has died down to a whisper. He has been left with an unsettling vision of reality unmediated by discursive thought.”

Anyway, to paraphrase from a review article of James Joyce’s book *Finnegans Wake*, I feel like my Incamination has been thoroughly Exagminated and Factified by this process of doing a PhD. To the Internal and External Exagminator, my thanks, and may our paths cross again someday.

Warm wishes,

Max