Welfare and felt duration

Andreas Mogensen (Global Priorities Institute, University of Oxford)

Global Priorities Institute | November 2023

GPI Working Paper No . 14-2023

Please cite this working paper as: Mogensen, A. (2023) Welfare and felt duration. *Global Priorities Institute Working Paper Series, No. 14-2023* Available at: https://globalprioritiesinstitute.org/welfare-and-felt-duration-andreas-mogensen

GLOBAL PRIORITIES INSTITUTE

UNIVERSITY OF OXFORD

Welfare and Felt Duration

Andreas L. Mogensen*

Abstract

How should we understand the duration of a pleasant or unpleasant sensation, insofar as its duration modulates how good or bad the experience is overall? Given that we seem able to distinguish between subjective and objective duration and that how well or badly someone's life goes is naturally thought of as something to be assessed from her own perspective, it seems intuitive that it is subjective duration that modulates how good or bad an experience is from the perspective of an individual's welfare. However, I argue that we know of no way to make sense of what subjective duration consists in on which this claim turns out to be plausible. Moreover, some plausible theories of what subjective duration consists in strongly suggest that subjective duration is irrelevant in itself.

1 Introduction

An experience of pain is worse for you the longer it goes on. This much seems obvious. But how should we understand the duration of a pleasant or unpleasant sensation?

The question is worth raising because we seem able to distinguish between *subjective* and *objective time*. A minute sometimes feels much longer than a minute, and sometimes much shorter. It's possible that different kinds of minds – those of small, high-metabolism animals (Prosser 2016: 85–87; Schukraft 2020; Yong 2022: 74–76) or of digital

*For comments on previous drafts of this paper, I'd like to thank Adam Bales, Heather Browning, Hilary Greaves, Brad Saad, Jeff Sebo, Jobie Sheppard, Derek Shiller, Teru Thomas, and Hayden Wilkinson, as well as audiences at the 12th Oxford Workshop on Global Priorities Research in June 2023 and at the Department of Philosophy at the University of Bristol in October 2023.

persons of the not-too-distant future (Bostrom and Yudkowsky 2014; Hanson 2016; Shulman and Bostrom 2021) – might vary dramatically in their experience of time's passage, living through a much greater amount of subjectively experienced time within a given unit of objective time. To them, the experience of pain filling mere seconds or minutes might in some sense be more like our experience of a pain that lasts many hours or days.

How well or badly someone's life goes is naturally understood as something to be assessed from her perspective (Railton 1986; Rosati 1996; Sumner 1996; Hall and Tiberius 2016; Dorsey 2017). Therefore, it seems intuitive that a valenced experience that's subjectively experienced as longer makes a greater difference to your welfare, holding fixed its intensity, objective duration, and any other evaluatively significant properties, while a valenced experience that's objectively longer makes no greater difference to your welfare, holding fixed its intensity, subjectively experienced duration, and any other evaluatively significant properties (compare Lee 2013; Bostrom and Yudkowsky 2014; Schukraft 2020; Shulman and Bostrom 2021). As Terry Pratchett (1990: 10) writes: "the important thing is not how long your life is, but how long it seems."

I argue against the claim that the subjective duration of a valenced experience is the important thing. More exactly, I argue against the claim that a valenced experience that's subjectively experienced as longer makes a greater difference to your welfare, holding fixed its intensity, objective duration, and any other evaluatively significant properties. I do not also present a positive argument for the contrary claim that the extensive magnitude of a valenced experience should instead be measured in terms of its objective duration. As the natural alternative, I do think that position is a lot more plausible than it might initially appear. However, I also give some credence to the idea that perhaps neither subjective nor objective duration has any fundamental evaluative significance and that what makes longer pains worse ultimately has to be explained in terms that have nothing essentially to do with length of time or experience thereof (see section 4).

I start in section 2 by clarifying some basic conceptual issues and explaining the importance of getting clear on how, if at all, subjectively experienced duration modulates welfare. In section 3, I look at two analyses of the nature of subjective time experience in the recent philosophical literature that strike me as especially attractive. I argue that, al-

though each may be plausible as an account of what felt duration consists in, on neither is it plausible that felt duration *per se* modulates the contribution of valenced experience to individual welfare. In section 4, I rebut an intuition pump appealing to digital reproductions of conscious experiences that many people find persuasive as an argument for measuring the duration of valenced experiences in terms of subjective time. Section 5 provides a brief summary and conclusion.

2 Preliminaries

This section is designed to help set the terms of the debate. It also explains why the debate is worth having in the first place. I'll begin by explaining some core assumptions I do and don't make about welfare and valenced experiences. Then, I'm going to lay out some conceptual foundations for our thinking about objective and subjective duration. Finally, I discuss the importance of getting clear on how, if at all, subjective duration modulates welfare.

2.1 Welfare and Valenced Experience

By a 'valenced experience,' I mean a conscious experience that feels good or bad, pleasant or unpleasant. For simplicity and concreteness, I focus on pain.[1] Most of the discussion is therefore about how subjective duration contributes to pain's badness. Still, I conjecture that what I say about pain generalizes to valenced experiences more generally.

I assume no particular theory of welfare. In particular, I don't assume a hedonistic theory on which welfare supervenes on valenced experience. Among the constituents of welfare recognized by other theories, duration may well also be significant,[2] and so we can

[1] Some believe that pain without unpleasantness is possible in light of apparent dissociations in cases involving morphine analgesia, prefrontal lobotomy, or pain asymbolia (Hall 1989; Grahek 2007; Bain 2014). Assuming pain without unpleasantness to be possible, the discussion should be read as focused on experiences of unpleasant pain.

[2] For example, among philosophers who have recently defended the view that welfare consists in desire satisfaction, Dorsey (2013: 162) holds that "longer-held desires are more significant because they affect the welfare

ask the same sort of question I'm asking about pains. Is it objective duration that matters or subjective duration (or both or neither)? Nonetheless, pains and other valenced experiences seem especially like the sort of things whose contribution to welfare depends on felt duration. After all, they are naturally thought of as good or bad for you in virtue of how they feel.[3]

2.2 Subjective and Objective Duration

It may be thought that the theory of special relativity (Einstein 1905) entails that pains have no objective duration. The theory entails that observers at motion relative to one another needn't agree on the time-separation between events, and that the measured duration of a pain as determined by an observer's clock will depend on her state of motion relative to the person in pain. In that sense, the pain does indeed have no objective duration.

Taken in a different sense, however, relativity is compatible with attributing objective durations to pains. We can identify the objective duration of a pain experience with its *proper time interval*. The proper time elapsed along a timelike curve joining two events in

---

value of more times throughout a person's life." Heathwood (2005: 490) argues that a person is benefited by getting what she wants only if she gets it while she wants it, as a result of which he "does not take the *duration* of a desire to be as prudentially significant as some have taken it to be." But by that he just means that "[i]t doesn't matter, welfare-wise, how long I desire something before I get it." (ibid.) He thinks it matters "how long the concurrent desiring and getting last" (ibid.).

[3]While I don't assume any particular theory of welfare, I do assume the falsity of a certain kind of flat-footed subjectivist answer to our question, which says there's nothing to argue or inquire about: it's just a matter of taste. Do you prefer your pains to be subjectively shorter even if they're objectively longer or objectively shorter even if they're subjectively longer? There are no wrong answers, any more than there are wrong answers about whether to prefer chocolate or vanilla. I think subjectivists ought not to accept this flat-footed response. They characteristically hold that favourable attitudes explained by misunderstandings of the nature of their objects do not confer welfare significance on those objects (e.g., Rawls 1971: 416–417; Brandt 1979: 268; Railton 1986: 15-16), and key parts of my argument may be understood as trying to get us to see what the subjective experience of time might really consist in, and in making the case that once we do, we won't care about it. Note also that if the duration of a favourable attitude modulates the contribution of its object or the attitude-object pair to individual welfare (see the previous footnote), then subjectivists cannot in general escape the need to answer the kind of question I am raising. Thanks for Derek Shiller for the latter observation.

spacetime is the arc length of the curve (according to the Minkowski metric). This is the time measured by an idealized clock travelling along that path. Although timelike curves of different arc lengths can connect the same events, proper time along a given curve is invariant. Thus, in respect of pains, all observers will agree on how long the pain lasted according to the sufferer's own idealized clock.[4]

What, then, is subjective duration? Since core parts of my argument turn on exploring different answers to this question, a full analysis won't be given here. Instead, I'm going to lay out a handful of more fundamental concepts and distinctions that I think it's useful to have on the table already at this point.

Let's start off with the distinction between *judged duration* and *felt duration* (Arstila 2012: 3; Merino-Rajme 2014: 245–246; Wearden 2016: 131–141). Roughly, this is the distinction between (mere) beliefs about how much time has passed and the conscious experience of time as passing at a certain rate. Time can feel drawn out even if you know exactly how much time is passing and aren't at all disposed to overestimate duration. Conversely, your intuitive estimate of how much time has passed can overshoot or undershoot without any experienced feeling of time as having passed quickly or slowly.

This distinction seems especially important in the present context (Schukraft 2020). Someone in pain for six minutes might believe before, throughout, and after that it never went on for more than five. All else being equal, she is surely no better off for that. Whatever she believes, she had six minutes of pain. Mere beliefs about how much time has passed don't seem capable on their own of making an experience better or worse. Something about the character of the experience itself has to be different. Time needs not merely to be *believed* to have passed at a slower or a faster rate, it needs to have been *felt* as such.

Unfortunately, almost none of the published psychological literature on time perception is about felt duration as distinct from judged duration. In a recent survey article, Matthews and Meck (2016: 870) claim that "we can never be sure whether a given variable really affects subjective experience rather than the participant's response strategy"

---

[4]Many thanks to Hilary Greaves for help in drafting this and the previous paragraph. Any remaining mistakes are entirely my own.

and propose to "follow widespread practice and use 'judged duration,' 'apparent duration,' and 'subjective duration' interchangeably." Nonetheless, research on reports of durational phenomenology (so-called *passage of time judgments* (Wearden 2015)) supports the view that judged duration and felt duration are psychologically distinct. For example, reported feelings of time as passing more slowly or quickly have been found to be unrelated to verbal estimates of duration for a target auditory stimulus or the time taken in producing a sound to match a target for tones lasting up to half a minute, although a relationship is observed for longer durations (Droit-Volet and Wearden 2016; Droit-Volet et al. 2017). Likewise, depression is associated with a feeling of time as passing slowly, but verbal estimates of duration and other measures of judged duration don't differ significantly between people with depression and controls (Thönes and Oberfeld 2015).

Passage of time judgments of the kind studied by psychologists have to do with experiences of felt duration as different from normal. They ask respondents to say whether time is felt as passing quickly or slowly. But we might be especially interested in differences in the way felt duration is normally experienced across individuals. For example, we might want to know whether small, high-metabolism animals or digital minds of the not-too-distant future live through more subjective time in a given day or year than we do. That thought is sometimes expressed in terms of the idea that the world as viewed through their eyes will seem to unfold as if in slow motion (Bostrom 2014: 53; Lee 2017: 157–158; Yong 2022: 76). But to each mind, perceiving the world the way it normally does, time's passage will surely seem neither fast nor slow, insofar as it seems any way at all.[5] In that sense, the kind of phenomenology captured by passage of time judgments arguably won't vary between ordinary flesh-and-blood humans, hummingbirds, or digital emulations of human brains run at high clock speeds. So in what sense could there be differences in the way felt duration is normally experienced across individuals like these?

One sense in which time might be said to pass more slowly for other minds than for me takes the form of the subjunctive claim that I would feel as if time were passing unusually

[5]This is not to suppose that it is impossible for there to exist minds that always have experiences with contents that represent time as passing unusually quickly or slowly, but simply that it is very implausible that something like this occurs in the cases under discussion. Thanks to Brad Saad for pressing me to clarify this point.

slowly if I imagined suddenly taking up their point of view. But that thought makes sense in only a limited range of cases. For example, I don't know what it would mean to take up the point of view of a hummingbird so as to be able to notice a difference in my subjective experience of time, and whatever I might imagine surely tells me next to nothing about the inner life of hummingbirds (compare Nagel 1974: 438–440).

The better approach, I think, is to ask whether there is variation across different kinds of minds in the normal values of those psychological variables whose variation within individuals is experienced as time passing unusually slowly or unusually quickly. That's what I'll mean when speaking in terms of the possibility that there are differences in the way felt duration is normally experienced across different kinds of minds. Note that there could in principle be different psychological changes experienced as time passing unusually slowly or quickly (compare Prosser 2016: 113–114; Lee 2017: 159). Furthermore, these could differ in whether and to what extent they matter for individual well-being. We shouldn't assume at the outset that felt duration is one thing and one thing only.

Here's one final issue to do with how we conceive of subjective duration. Imagine a mind that can't remember anything from as little as a second ago. It can feel pain, though, and its pains sometimes go on for a minute and sometimes for two. There's a certain obvious sense in which the two minutes of pain don't feel any longer than the one from its perspective. It can't tell the difference. Earlier, I claimed that it's intuitive to think that a valenced experience that's objectively longer makes no greater difference to your welfare, holding fixed its subjectively experienced duration. But surely it's worse for this being to have two minutes of pain rather than one. Does this refute the intuition?

No. The subjectively experienced duration of a given time interval needn't be understood in terms of how long the interval as a whole feels, where this involves some apprehension of the interval in its entirety. Instead, we can think of it as something constructed by aggregating the felt duration of the interval's proper parts. Arguably, even we don't really have anything that constitutes an experience of a minute as a whole. The orthodox view in the philosophy of time perception maintains that, at any given point in time, a person undergoing a conscious experience is aware of an extended time interval, but a very short one. This is the famous *specious present* (Clay 1882; James 1886; Husserl 1991

[1893-–1917]; Broad 1923; Dainton 2000). The assumption that the content of experience is an extended time-segment apprehended as an integrated whole, rather than an instantaneous time-slice, is supposed to allow us to account for the fact that we seem able to perceive motion - that is, change in position over time - with the same immediacy that we perceive shape and colour. But estimates of the extent of the specious present are typically at the sub-second level, and it seems clear that a minute is far too long to be experienced via a unitary percept. Psychological evidence supports the hypothesis that different processing mechanisms are involved in time perception for durations at the sub-second scale and for longer ranges (Rammsayer and Troche 2014).

Obviously, if we elect to understand the subjective duration of an event that extends beyond the specious present as built up from the felt duration of its parts, a question arises as to how to aggregate the felt duration of the different parts to determine the subjective duration of the whole. But that question is best reserved until the point at which we have a more concrete conception of what the felt duration of the parts consists in. For example, if we think of conscious experience as broken up across discrete experiential frames, each of which corresponds to a unit of subjective time experience, then we might want to just count the number of frames, whereas that approach seems to be unworkable if we take experience to be continuous and stream-like.

2.3 Importance

At this point, the nature of felt duration ought still to be a mystery to us. We know that mere judgments about how long a pain has gone on don't matter. Something about the way the pain is experienced must be in play. But what exactly? Until we know, it's hard to say whether subjective duration modulates welfare. But why should we care to say?

On the one hand, it might affect how we think about the relative weight of good and bad experiences. Positive affect is associated with an experience of time as passing quickly and negative affect with an experience of time as passing slowly (Droit-Volet and Wearden 2016; Droit-Volet et al. 2017). Pain leads to an experience of time's passage as drawn out in both clinical and experimental conditions (Somov 2000; Hellström and Carlsson 1997). Similar results are obtained for estimates of duration (Wearden 2016: 105–115). For ex-

ample, Rey et al. (2017) find that being in pain increases judged duration by about a third of a standard deviation in a temporal bisection task, where subjects have to say whether a sub-second stimulus is more similar to a short or long template.

Some philosophers claim that pains are worse for you than pleasures are good, holding fixed their intensity and duration (Moore 1903; Mayerfeld 1999; Hurka 2010). For example, Mayerfeld (1999: 133) claims that if we imagine "an episode of very intense suffering" and "an episode of happiness of equal intensity," then "the intense suffering would not be compensated by an episode of the intense happiness lasting for the same amount of time." Claims of this kind can seem mysterious, in that they appear to posit a kind of brute asymmetry between good and bad feelings. But suppose Mayerfeld has in mind objective duration when he talks about suffering and happiness lasting "the same amount of time".[6] Then, in light of the psychological facts noted above, we have a natural explanation for why the quoted claim comes out as true, assuming that felt duration is the proper measure of valenced experience. The suffering should fill more subjective time.

Thinking about how subjective time experience modulates welfare also has the potential to affect the weight we put on the valenced experiences of minds quite unlike our own who run at speeds quite unlike our own. For example, it's possible we underestimate the lifetime welfare of wide classes of non-human animals, because we measure the length of their lives in clock time, whereas those animals subjectively live through much more time within a given day or year than we do (Schukraft 2020). That's actually the point Terry Pratchett is driving at in the passage from which I quoted earlier. He writes: "One of the shortest-lived creatures on the planet Earth is the adult common mayfly. It lasts for one day. ... This may seem tough on mayflies. But the important thing is not how long your life is, but how long it seems. To a mayfly, a single hour may last as long as a century." (Pratchett 1990: 10)[7]

A number of different lines of evidence indicate that some animals sample percep-

[6] Insofar as Mayerfeld (1999) discusses durational extent, he speaks in terms of seconds, minutes, and hours, which to me suggests he must have clock time in mind. But he does not discuss how to understand duration in any depth, telling us there is "nothing mysterious" (130) about it.

[7] I'm grateful to Ian Phillips for bringing this quote to my attention.

tual information from the environment at much higher rates than we do and so might experience a greater number of subjective moments per unit objective time. For example, some animals are able to act and adjust their behaviour in response to stimuli at timescales that seem impossibly fast to us (Prosser 2016: 85–86; Schukraft 2020). Killer flies have the fastest known photoreceptors, requiring just 6–9 ms for the nervous system to process an image and prepare an action in response, compared to the 30–60 ms for human photoreceptors just to begin signalling the brain. The flies' hunts are over and done in about a quarter of a second and are nearly impossible for the human eye to follow, except when played in slow motion from recordings by high-speed cameras (Yong 2022: 74–75).

Further evidence comes from studies of *critical flicker-fusion frequency* (CFF), the frequency at which a flickering light source is perceived as continuously illuminated (Schukraft 2020; Yong 2022: 75–76). For human beings, that's at around 60 Hz. A small passerine bird known as the pied flycatcher has a CFF of 146 Hz, whereas honeybees, dragonflies, and flies have CFFs in the range of 200 to 350 Hz. These animals are able to perceive gaps in the flicker of rapidly alternating light sources that we could only detect from slow motion recordings.

As I've also hinted previously, even more dramatic speedups in processing and concomitant slowdowns in subjective time might be realized by non-organic minds run on digital hardware. That's due to the extraordinary speed advantages of electronic circuity. Electronic circuit boards can achieve signalling speeds millions of times greater than the speed of interneuronal communication and can modulate their internal states at rates billions of times faster than neuronal reaction times (Hanson 2016: 79–80). As a result, a digital emulation of a human brain could conceivably pack the experience of many lifetimes into mere hours of objective time.

This is one among a number of considerations highlighted by Shulman and Bostrom (2021) in arguing that digital minds we might one day create could exhibit superhuman capacities for welfare and capture the vast majority of total well-being across future time. In doing so, they rely on the intuition that felt duration is the proper measure of valenced experience. It's time I got round to explaining why I think we should demur from that belief.

3 Against Felt Duration

'Felt duration' is obviously a philosopher's term of art. The phenomenon itself is elusive, especially once we keep in mind that it's supposed to be distinct from judged duration. Cases of mistaken identity may be widespread, with Wearden et al. (2014) claiming that reports of the experience of time passing quickly are often misleading, in that what people report as 'fast time' may involve nothing more than realizing at the end of some event that more clock time had elapsed than they thought. This means "the report is generated on the basis of an inference, often prompted by an external time marker, without any actual 'feel' of fast time during the event." (Wearden et al. 2014: 303) If what people reliably report as an experience of time passing quickly isn't really an experience of time passing quickly, what is? What exactly does the 'feel' of time's passage that Wearden et al. (2014) appeal to consist in? Until we can answer that question, we should be hesitant to draw conclusions about the contribution of felt duration to individual welfare.

I'm going to look at two kinds of analyses of the nature of subjective time experience that have appeared in the recent philosophical literature and that strike me as especially plausible.[8] Very roughly, the first of these can be thought of as appealing primarily to cognitive factors and the second as appealing primarily to aspects of perceptual processing in explaining our experience of time's passage. For reasons I've already noted, these needn't be thought of as rival theories. Each may be successful as an account of why time should sometimes be experienced as passing unusually slowly or quickly. But in neither case should it seem plausible to us that felt duration – as distinct from objective duration – matters for how well someone's life goes.

---

[8] Other analyses are, of course, available. Lee (2013) considers and rebuts a number of proposals that understand subjective durational experience as determined by the density of information processing. Prosser (2016) proposes an account of subjective time experience based on an enactivist theory of perception, which I discuss briefly in footnote 10, and which is similar in some respect to the views of Arstila (2012) and Phillips (2013), discussed in the next subsection, but distinguished by its singular focus on action. Prosser's view strikes me as similar enough to the view developed by Arstila and Phillips that I would expect that the key points I make about their view in the next subsection can be easily transposed to apply to Prosser's.

3.1 The Speed of Thought

If time really were passing more slowly or more quickly, we too would be sped-up or slowed-down accordingly, and so presumably wouldn't notice. On the other hand, in becoming aware of a mismatch between our own speed and the speed of external events, it would seem natural to expect a feeling of time as passing unusually quickly or slowly. This in turn gives rise to the idea that our experience of time's rate of passage may depend on our awareness of the speeds at which external events, as registered in perception, unfold relative to the speed of our internal cognitive processes. Call this the *cognitivist theory of felt duration* – or *cognitivism*, for short.

Arstila (2012) and Phillips (2013) both appeal to cognitivism to explain durational phenomenology in moments of life-threatening danger.[9] Dramatic experiences of time slowing down are reliably reported in extreme crisis situations (Noyes and Kletti 1977; Flaherty 1999: 50–56). For example, one student, whose steering gave out while driving at 60 miles per hour, reports: "Time seemed drawn out. It seemed like five minutes before the car came to a stop when, in reality, it was only a matter of a few seconds." (Quoted in Noyes and Kletti 1977: 376). The survivor of an earthquake in Armenia reports that the quake "was like a slow-motion movie ... There was a concrete panel slowly falling down." (Quoted in Flaherty 1999: 51)

In order to explain why time's passage seems to slow down in moments of life-threatening danger, Arstila and Phillips appeal to the idea that our conscious experience includes not only perceptions of events unfolding in the world, but also conscious thoughts and other cognitive processes unfolding in our minds. Not only that, but we're aware of temporal relationships between them. The subjective experience of time's passage can thus be accounted for in terms of the fact that we're "aware of the durations of environmental events relative to the non-perceptual conscious activity that occurs between their onset and offset." (Phillips 2013: 232)

The experience of time slowing down during moments of life-threatening danger can

[9] Phillips (2012) also applies the theory to explain the psychological effects of attention to time's passage on time perception.

then be explained in terms of a speed-up in the pace of non-perceptual conscious mental activity. Whereas Phillips seems to appeal to direct awareness of our thoughts in explaining how the speed of mental activity affects conscious experience, Arstila proposes that we become aware of our increased rate of cognitive activity indirectly, as a result of active engagement with the external world: e.g., by registering the heightened frequency with which we can shift attention from one stimulus to another, initiate new actions, or react to new stimuli.[10] A speed-up in mental activity has obvious adaptive value in life-threatening situations (see Phillips 2013: 245–246), and increased speed of thought is reported by a majority of subjects interviewed by Noyes and Kletti (1977). That's no coincidence, according to these authors. As Arstila (2012: 8) puts it: "our feeling that we are somehow faster than usual also amounts to the experience that things in the external world happen slower than usual." In fact, one of the accident victims interviewed by Noyes and Kletti (1977: 378) appeals to just this idea in making sense of their experience: "my thinking processes increased at an incredible rate so that my movements, in relation to them, seemed extremely slow."

I claim that insofar as variation in the subjective experience of time is explained in terms of the cognitivist theory of felt duration, variation in the subjective duration of a painful sensation shouldn't alter our evaluation of how bad the pain is, all else being equal. In other words, an experience of pain during an interval that feels as if it lasts for many minutes but is measured in mere seconds by the subject's clock is no worse than an otherwise exactly similar experience involving pain of the same intensity that seems to its subject to fill mere seconds, assuming that differences in the experience of time's passage are explained in terms of cognitivism. Why not? Because, intuitively, a pain is no better or worse merely in virtue of the fact that your thoughts, imaginings, and rememberings seem to move more slowly or quickly in relation to external events while you're in pain.

---

[10] Prosser (2016) offers a different view on which the subjective experience of time's passage is explained in terms of changes in our awareness of the actions available to us. In line with the enactivist theory of perception (Gibson 1979; Noë 2006), Prosser assumes that the intentional contents of perceptual states encode possibilities for action. He therefore interprets variation in the experienced rate of time's passage in terms of the variation in the possibilities for action perceptually encoded in relation to a given interval of time.

In and of itself, the relative speed of non-perceptual conscious mental activity is simply irrelevant to pain's badness.[11]

It's important to be clear that I say 'in and of itself'. As a matter of contingent fact, changes in the speed of non-perceptual conscious mental activity may bring about changes in pain's subjective disvalue. They might even do so reliably. It is well-established that felt pain depends not only on raw input to first-order nociceptive neurons, but on a range of social and psychological factors (Melzack and Wall 1988: 15–34; Ambron 2022: 140–199). Among relevant cognitive factors, the allocation of attention toward or away from pain sensations, anticipations of increases or decreases in felt intensity, and dispositions toward catastrophizing are all known to affect the severity of pain experience (Melzack and Wall 1988: 22–28; Ambron 2022: 171–184). People can even be trained via biofeedback to deliberately modulate pain-related activity in neurons in the anterior cingulate cortex, allowing them to dial pain's felt unpleasantness down by an exercise of will (de Charms et al. 2005). For these reasons, it's important to be clear that I don't mean at all to deny that changes in cognition can modulate pain's badness. All I mean to deny is that the speed of conscious thought might do so in and of itself.

Two additional points are worth bearing in mind.

---

[11] Someone might claim that how quickly one's thoughts seem to pass alongside a pain is relevant to pain's badness precisely because this is what the subjective experience of time consists in, and what matters in respect of a pain's duration is how long the pain is felt as lasting. They might note that we would probably also be left cold by a description of pain processing couched entirely in the language of neurobiology, if we do not realize that this just is what pain consists in. But that analogy is misleading, for two reasons. Firstly, there is a well-known inferential gap between physical-functional concepts and phenomenal concepts in light of which a characterization of pain in respect of how it feels remains cognitively isolated from any neurobiological description of the brain's inner workings (Loar 1990; Papineau 2002; Tye 2003b; Chalmers 2004; Goff 2011; Balog 2012). There is no similar conceptual gulf here, since we are in the business of giving a reductive explanation of a particular kind of phenomenology in phenomenological terms. No similar explanation can therefore be given for why the evaluative significance of that particular kind of phenomenology should be cognitively isolated from its reduction basis as characterized here. Secondly, we know all too well what it is like to be in pain. When it comes to felt duration, the phenomenon itself is elusive, and one on which we have a only a tenuous grasp, absent some analysis of what is meant by the 'feel' of time's passage. For this reason, our assessment of the welfare significance of felt duration ought to be guided by our appreciation of what felt duration consists in according to some preferred analysis, and not *vice versa*.

Firstly, to the extent that cognitive factors like those surveyed above affect the badness of pain experience, it's plausible they do so by altering experienced intensity, whereas my concern is with how we should think of the contribution of duration to pain's badness, holding intensity fixed. If that's right, then what we know of the ability of cognition to worsen pain turns out to be irrelevant to our inquiry, strictly speaking.

Secondly, there seems to me to be insufficient evidence for supposing that the particular changes in cognitive processing that underlie changes in the subjective experience of time according to the cognitivist theory of felt duration will tend to reliably increase or diminish pain's subjective disvalue, given our understanding of the ways in which cognitive factors alter pain experience. For example, even granting that attention modulates pain intensity, it seems to be an open question whether apparent increases in the speed with which you're able to shift attention from one thing to the next will exhibit a tendency to make pain experiences worse as opposed to better. Will this lead to increased overall attention directed toward pain? Or away from it? It would be valuable to have further empirical evidence that would allow us to say whether changes in the speed of non-perceptual conscious mental activity in fact tend to bring about worse pain experiences, even granting that how quickly your thoughts seem to pass alongside a pain of fixed intensity is irrelevant in itself to how much the experience detracts from your well-being.

Before wrapping up this subsection, it's also worth dwelling on the distinctively cognitivist character of the explanation proposed by these authors and its associated limitations.[12]

Arstila and Phillips explain variation in the experience of time's passage by appealing to variation in the speed of internal cognitive processes that run alongside perception, as opposed to via changes in the character of perceptual processing per se. Therefore, no increase in the temporal resolution of perceptual experience need be expected as a result of experiencing time as slowed down (Arstila 2012: 4–5, 8; Phillips 2013: 230), and,

[12]To be clear, these limitations need not be understood as objections to cognitivism, considered as an explanation for why time is felt as passing quickly or slowly in the cases to which the theory has been applied. Rather, they emphasize limitations of the ability of cognitivism to explain variations in felt duration beyond those cases.

indeed, Stetson et al. (2007) find no evidence of increased temporal acuity for vision during frightening free-fall experiences (contrary to those authors' expectations). The latter represents no embarassment to the cognitivist: "subjects do not perceive more snapshots during frightening events (as Stetson et al. hypothesized), but they have the feeling that the snapshots they perceive last longer than usual. It is not that there was an abundance of experiences, but *abundance of time to scrutinize the perceptions*." (Arstila 2012: 8) But, as a result, cognitivism also isn't well-placed to account for potential differences in the way felt duration is normally experienced across individuals in cases where our evidence consists of differences in the temporal resolution of perceptual experience, such as evidence of variation in CFFs across the animal kingdom. Granting that Arstila and Phillips successfully explain the experience of time's passage slowing for human subjects undergoing moments of life-threatening danger, variation in the character of perceptual processing can give rise to differences in the way felt duration is normally experienced only on the assumption that felt duration is multiply realizable at the level of psychology.

Cognitivism has other limitations that become especially apparent when we try to apply it in explaining potential variation in the subjective experience of time beyond the boundaries of our species. Cognitivists claim that "what subjects are reporting in terms of 'time slowing down' are experiences in which unusually large amounts of non-perceptual mental activity occurs within a certainy objective period." (Phillips 2013: 233). But we aren't told what's meant by an *amount* of non-perceptual mental activity, and it's reasonable to worry that there aren't any natural units in which to measure the volume of conscious cognitive processing (Lee 2013: 14; Prosser 2016: 98–99). An influential programme of research in cognitive psychology argues that the bandwidth of human working memory should be understood not in terms of *bits* as defined in Shannon information theory (Shannon 1948a, 1948b) but rather in terms of *chunks* (Miller 1956; Cowan 2001). Roughly speaking, a chunk is a unit of information like the number '1492' or the word 'apple' that gets processed as a meaningful whole rather than as a random string. A more precise definition remains elusive.[13] When dealing with recognisably human minds, we might sup-

[13]Mathy and Feldman (2012: 347) write: "half a century after Miller's article [Miller (1956)], the definition of a chunk is still surprisingly tentative. ... Most attempts to define chunks are somewhat vague, ad hoc, or severely

pose that the character of cognition is sufficiently similar across individuals that we need not worry too much about this. A chunk is a chunk is a chunk. The more and more different are the minds under comparison, the less certain can we be that there is a meaningful common measure on which to base comparisons of the volume of conscious cognitive processing, and the more we seem to be in the position of someone asked to say whether more charge has flowed through a length of wire than water through a pipe. If combined with the assumption that subjective duration is the proper measure of how long a pleasant or unpleasant sensation lasts for the purposes of welfare assessment, cognitivism therefore threatens to derail the possibility of meaningful interspecies comparisons of welfare.

3.2 Frame Speed

In a well-known passage in his *Principles of Psychology*, William James (1890: 239) describes the character of conscious experience as follows: "It is nothing jointed; it flows. A 'river' or a 'stream' are the metaphor by which it is most naturally described." No doubt it seems that way to us.[14] Nonetheless, a wide range of evidence appears to support the contrary hypothesis that conscious experience actually consists of a succession of discrete experiential frames (VanRullen 2016; Herzog et al. 2016; White 2018;[15] Drissi-Daoudi et al. 2019).

For example, people are found to be able to experience the *wagon wheel illusion* when viewing a spinning disk through their own eyes under conditions of constant illumination (Schouten 1967; Purves et al. 1996). The wagon wheel illusion is familiar from film recordings of wheels that seem to spin backward, owing to the fact that the recording samples information discretely from the environment and can end up successively catching the wheel at moments just short of a full rotation. The fact that normal human vision gener-

---

limited in scope, especially when they apply only to verbally encoded material ..., making it difficult for them to explain the existence of chunking-like processes in animal learning".

[14] If not to Strawson (2009: 234–240) or to certain practitioners of Buddhist meditation techniques (Davis 2018).

[15] To be clear, White provides a critical summary of evidence supporting the hypothesis of discrete experiential frames. He finds the evidence unpersuasive.

ates the same illusion suggests that visual experience also relies on discretely generated representations of the external world.

Consider also a recent study by Herzog et al. (2020). They find that two opposite-direction offsets in a stream of vertical lines integrate and cancel out any effects on visual experience if they occur within 290-450 ms of one another. However, if three offsets are present at 0, 330, and 490 ms, this cancelling effect occurs only among the first and second, even though the second and third are much closer together. This suggests that conscious visual information is updated at discrete intervals based on processing windows of around 450 ms, as opposed to being a continuously updated stream derived by integrating information within a sliding window.[16]

Some authors relate the integration of information within discrete perceptual frames to oscillatory cycles in neural activity (VanRullen 2016). For example, the scalp-recorded alpha rhythm (8-12 Hz) has been hypothesized as related to the 'frame rate' of perception. Slower alpha rhythms degrade the temporal resolution of perception, an effect hypothesized as due to the greater likelihood that successive stimuli are processed within the same perceptual frame and so can't be distinguished in time (Samaha and Postle 2015). Strikingly, Mioni et al. (2020) report that by experimentally speeding up or slowing down alpha rhythms using transcranial alternating current stimulation, they were able to increase or decrease judged duration in a time generalization task that required subjects to say whether a probe stimulus was of the same duration as a previously learned target stimulus.

Results of this kind seem to vindicate a hypothesis about time perception originally proposed in the middle of the 19th century by von Baer (1864). Von Baer hypothesized that conscious experience consists of a succession of discrete experiential frames and that the speed at which these frames succeed one another provides "the basic standard

[16]Granting that conscious experience is structured in terms of windows of integration, it's actually plausible that there are multiple, layered frames on the basis of which perceptual information is integrated at different levels of temporal resolution (Montemayor 2012; Wilson 2022). At the highest grain might be very brief elementary frames within which stimuli aren't registered as temporally ordered. Layered on top of those might be windows with lower temporal granularity that integrate information over longer intervals and allow stimuli within them to be discriminated on the basis of time.

by which we measure time in observing nature." (von Baer 1864: 258)[17] Among contemporary philosophers, the clearest defense of a theory of experienced duration based on the assumption that conscious experience consists of a succession of discrete perceptual frames comes from Merino-Rajme (2014). She argues that conscious experience is divided into *quanta*, each presenting a tightly unified and temporally bounded arrangement of experienced qualities, corresponding to the specious present. In her view, "whatever features are present in a quantum will be experienced as taking one subjective unit of duration." (Merino-Rajme 2014: 255) The felt duration of an event is then to be understood in terms of the number of quanta experienced as making up one's overall perception of the event. She calls this the *quantum theory of felt duration*.

Before looking at what a theory of felt duration based on the number of experienced frames that make up the perception of an event might tell us about the way felt duration modulates welfare, I want to stop and clarify two issues. The first has to do with what it means to say that conscious experience is discrete as opposed to continuous.[18] The second is to do with how the number of experiential frames experienced within a given interval relates to felt duration.

I take the evidence presented by authors like Herzog et al. (2020) to support the hypothesis that normal waking experience involves discontinuous changes in the content of conscious experience. It's important to be clear that this needn't entail that there are gaps in consciousness itself. In other words, it needn't entail that we aren't actually having experiences throughout our waking lives and that consciousness instead arises in short-lived bursts, surrounded by periods of unconsciousness, flickering on and off like the image on an old cathode ray tube monitor, as suggested by currents in Buddhist philosophy of mind (Davis 2018) and, more recently, by Strawson (2009: 243–244). Discontinuities in the content of conscious experience need not entail gaps in consciousness itself. Conscious-

[17] Strikingly, von Baer (1864) also hypothesized that the frame rate of conscious experience varies significantly across the animal kingdom and that extreme variations could allow a person to watch the flight of a shotgun shell through the air or see mountains being born.

[18] For prior discussion of the many different ways this idea may be understood, see Strawson (2009) and Rashbrook (2013).

ness might operate on a 'sample and hold' principle, like modern liquid crystal display screens, continuously maintaining a given content until it undergoes an instantaneous or near-instantaneous refresh at discrete intervals. Nonetheless, Herzog et al. (2020: 833) write as if their results imply the existence of "gaps in between the conscious percepts". That hypothesis seems to me to be unsupported by their evidence.

On the question of how the number of experiential frames experienced within a given interval relates to felt duration, I want to comment on what I think marks a wrong turn by Merino-Rajme in her treatment of this issue. The wrong turn, as I see it, involves explaining felt duration in terms of our awareness of the frames themselves. In her view, "The number of quanta 'counted' while experiencing a long-lived event determines the amount of duration that the long-lived event is experienced as having." (Merino-Rajme 2014: 256) She's clear that she doesn't imagine that we literally successively enumerate experienced quanta in order to form a subjective impression of duration. Rather, we're supposed to construct an approximate representation of their number, as we might do when glancing over a room of people and estimating their number without counting each in turn. But even this seems to me quite implausible. If in fact conscious experience consists of a succession of discrete experiential frames, that's not something of which most of us have any awareness in the course of our waking lives. As James (1890: 239) says, consciousness "does not appear to itself chopped up in bits."

The better approach seems to me to simply rely on how many frames make up your experience over a given length of time, setting aside any awareness of each window of integration succeeding the one before. Recall that according to Merino-Rajme, "whatever features are present in a quantum will be experienced as taking one subjective unit of duration." (Merino-Rajme 2014: 255) In a similar vein, Strawson (2009: 253–253) writes that "the subjective feel of the lived present of conscious experience, the temporal feel of the temporal 'window' within which one lives as subject of experience ... stays relatively fixed at all times". A view of this kind suggests a different and quite straightforward way in which changes in the frame rate of experience can be said to modulate experienced duration. Assume that the specious present corresponds to a given window of integration. Since every such window is experienced as one subjective unit of duration, more units of subjective

time are experienced in a given interval of proper time for subjects for whom more such windows of integration are processed, regardless of any awareness on the subject's part of how many windows go by.

One might worry that this doesn't get us all that far, in that it amounts to explaining the felt duration of experiences that stretch beyond the specious present in terms of the felt duration of the specious present itself, without explaining what the felt duration of the specious present consists in. But the natural reply, I take is, is that the specious present represents the basic unit in terms of which we experience the flow of time and occupies one unit of subjective duration for us in just the same sense that the standard metre bar occupies a length of one metre. This may be just what von Baer had in mind in describing the subjective moment as "the basic standard by which we measure time in observing nature."

Keeping those points in mind, why should the number of discrete frames that divide up one's experience over a given objective time interval matter to the badness of a pain experienced throughout that time?

It may seem tempting to respond that since a discrete experience as of pain filling the experienced present simply is the fundamental unit of which experiences of pain over time are built, it has to follow that the extensive magnitude of such experiences should be measured by the number of such units they encompass. A discrete experience as of pain filling the specious present functions simply as the fundamental unit of accounting – the atom out of which pain experiences are composed.

I think it's clear we shouldn't say that. Even if our waking lives are composed of discrete experiential frames, there appears to be nothing in the nature of consciousness itself that requires that kind of discretization. Continuous consciousness seems to be possible, even if in fact our own experiences are discrete. For continuous minds, we seem forced to say either that there is no number of discrete experiences that make up the experience of a minute of pain or that there are uncountably many or that minds like that undergo exactly one discretely demarcated experience during any period of uninterrupted consciousness (compare Tye 2003a: 97). But each of these claims yields absurd results when making welfare comparisons across discrete and continuous minds, as well as among the experiences

of continuous minds, if we insist that the right way to measure pain's extent is in terms of the number of discrete experiences that comprise an experience of pain. For this reason, we can't insist on a discrete experience as of pain filling the specious present as a fundamental unit of accounting in assessing the subjective badness of pain. We need to be able to explain the evaluative significance of the compartmentalization of pain experiences into more or fewer discrete frames.

I can think of two possible explanations of the evaluative significance of the compartmentalization of pain experiences into varying numbers of discrete frames. Nonetheless, each comes with its own unique limitations, and each is based on a bold conjecture about how the mind works.

Here's the first. As noted earlier, theories of discrete perception typically link the frame speed of experience with temporal resolution: the faster we integrate perceptual information within experiential frames, the shorter is the represented interval that fills the specious present, and the better are we at discriminating events in time (see Lockwood 2005: 372–374). It's conceivable that pain experiences that admit of a higher degree of temporal resolution are typically worse because higher temporal resolution in nociceptive perception leads to the experience of fine-grained modulations in pain's intensity that would be averaged out at a lower temporal resolution.

Why should that be worse? Some philosophers believe that the badness of pain is a strictly convex function of intensity (Mayerfeld 1999: 134–135; Hurka 2010).[19] In other words, the extent to which you're made worse off by an increase of a fixed size in the intensity of pain is an increasing function of your pain's intensity. Pain intensity exhibits increasing marginal disvalue. It follows that an experience of a very intense pain followed by a relatively mild pain is worse overall than an experience of pain whose intensity is the average of the two, all else being equal. As a result, the ability to perceive more of the detail of how pain unfolds in time might end up being worse for you insofar as this means you experience fine-grained modulations in pain's intensity that would be averaged out at a

[19] Something like this may also follow in respect of pain's moral disvalue from a particular interpretation of the prioritarian hypothesis that gains in well-being matter more the worse off you are (Parfit 1991; Holtug 2010; Adler 2012). See Mayerfeld (1999: 149–158) and Hurka (2010) for further discussion.

lower temporal resolution.

The thing that's pure conjecture here is that pain intensity experienced at a coarser temporal grain is (no greater than) the (unweighted) average of the pain intensities that would be experienced at a finer temporal grain. A further major limitation is as follows. Consider an experience of pain that doesn't vary over time in those properties that translate as experienced intensity and so will be experienced as of the same intensity regardless of the temporal grain of experience. Appeal to the badness of pain as a supralinear function of intensity and the modulation of experienced intensity by the frame speed of perception provides no reason to suppose that increases in the temporal resolution of perception should worsen a pain like that. As a result, we haven't been given a general case for treating a pain that unfolds across a higher number of discrete subjective moments as worse, all else being equal.

Here's a second way in which we might explain the importance of frame speed for welfare, and one that might be able to achieve the level of generality we found lacking in the first.

As I noted earlier, it's important to be clear on what we mean if we say that conscious experience is discrete as opposed to continuous. In particular, the existence of discontinuous changes in the content of conscious experience doesn't entail the existence of discontinuities in consciousness itself. Let's suppose, nonetheless, that there are such discontinuities: conscious percepts don't fill our waking lives but instead arise briefly out of extended periods of surrounding unconsciousness. We can then imagine that a pain that feels longer is made up of percepts that are packed more densely in time due to the increased frame rate of conscious experience. If the duration of a percept itself changes little if at all with frame speed, the periods of surrounding unconsciousness that separate conscious percepts become briefer the slower time is felt as passing, as a result of which more time overall is filled with pain experience, as opposed to unconsciousness. That surely must be worse.

For reasons I've already noted, the postulate of unnoticeable intermittent on/off discontinuities in ordinary waking experience is conjectural. However, it is also not as farfetched as it might initially strike us as being. In particular, I think it's harder than we

might think to claim that we have introspective evidence against the existence of intermittent on/off discontinuities in ordinary conscious experience. That's even granting that introspection in the course of ordinary waking life presents us with nothing but gapless continuity.

We should remember the *content-vehicle* distinction. Just as an experience of a round, red patch is arguably not itself either red or round, so, we might think, an experience of something as continuous and uninterrupted need itself be neither continuous nor uninterrupted. Combined with the hypothesis that experience is 'transparent,' in the sense that the properties on which we focus in introspecting an experience are merely the represented properties of the experience's intentional objects, this delivers a picture on which the gappy character of conscious experience is compatible with our ordinary introspective awareness of persistent gapless continuity in the course of ordinary waking life (compare Tye 2003a: 95–97).[20] We can imagine that an experience as of continuous pain arises due to the fact that pain experiences are perceptual states with bodily injuries as intentional contents (Armstrong 1968: 306–322; Pitcher 1970; Tye 1995) and that individual experiences of pain, while discrete and instantaneous or near-instantaneous in themselves, have as their intentional contents the state of the body in relation to the site of injury throughout continuous, overlapping, just-passed intervals of time, in line with the so-called *retentionalist* conception of the specious present (Husserl 1991 [1893-–1917]; Broad 1923; Tye 2003a; Paul 2010; Lee 2014). (See Fig. 1.)

In any case, the postulate of unnoticeable intermittent on/off discontinuities in ordinary waking experience is only the first of the key conjectures we're relying on. Here is the second. Even assuming that the subjective experience of time slowing down involves a higher density of discrete percepts per unit time, I know of no reason to suppose, as we have assumed, that the duration of those percepts remains roughly constant, as opposed

[20] Arguably, we need not here assume *strong transparency*, on which it is *impossible* to attend directly to our experience, but only *weak transparency*, on which it is merely *very difficult* to focus our attention directly on features of our experience in introspection, as opposed to instead attending directly to the represented properties of the intentional objects of experience (Kind 2003). Thus, we can allow that the the gappy character of conscious experience might in principle be revealed through carefully practiced introspection of the kind supposedly available to experienced practitioners of certain Buddhist meditation techniques: see Davis (2018).

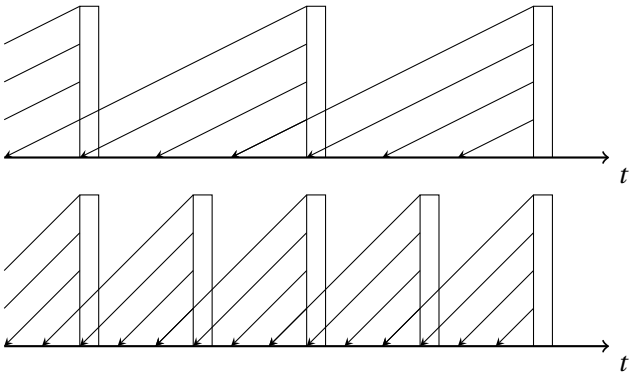

Fig. 1. Each vertical bar is a discrete moment of pain experience whose intentional content is the state of the body in relation to the site of injury throughout the time interval marked out by the descending arrows. In the lower panel, time is experienced as passing more slowly in light of a larger number of experienced moments per unit of objective time. The extent of the specious present is represented as contracting, but the duration of each discrete moment of pain experience is imagined as held fixed.

to contracting in proportion to their frequency and so filling roughly the same amount of time overall. And whatever doubts we might have about our warrant to suppose that discrete pain percepts have a roughly constant objective duration when dealing with intrapersonal variation in the subjective experience of time should clearly be amplified considerably when we compare across species.

Let me also note two important limitations of the proposed explanation for the importance of frame speed when it comes to pain's badness.

The first is that it imposes a hard limit on how much worse a pain can be made by increasing its subjective duration. If a pain percept ordinarily lasts for x milliseconds and the gaps between percepts are y milliseconds, then the experience of time slowing down during moments of pain at most magnifies the pain by a factor of $(1 + y/x)$. How so? Well, even if it's possible in principle for the subjectively experienced duration of a pain to increase by a greater factor by becoming divided over more than $(1 + y/x)$ times as many discrete percepts, each representing a unit of subjective time, the objective time available to be filled with additional percepts will have been maxed out, unless the average duration of each percept contracts.

Since we have no idea what values x and y might actually take, it's difficult to say anything concrete about the significance of this observation. Nonetheless, it gives us some

reason to be skeptical of the idea that if possible minds of the not-too-distant future that run on digital hardware process experiences at millions or billions of times the speed of grey and white matter, then the badness of a pain of fixed intensity experienced for the same objective time interval by a mind like that is millions or billions of times worse by virtue of stretching out through millions or billions of times the amount of subjective time. For that to be possible on the present model, the time occupied by conscious percepts would have to be mere millionths or billionths of our waking lives.

This leads me to my final point – and the most important limitation of the current proposal. Ultimately, it doesn't actually constitute an explanation for why felt duration – as distinct from objective duration – matters for how well someone's life goes. Really, it's a story on which objective duration is the proper temporal measure of pain experience. What's bad for you is to have more clock time overall filled with pain experience, as opposed to unconsciousness. It's just that what we ordinarily think of as a pain's objective duration is based on the false assumption that a pain fills out the entire interval between (what we call) its onset and its end. That false assumption in how we normally reckon objective duration means that subjective time experience can function as a better guide to the amount of clock time that the pain experience really fills. But that's still what ultimately matters.

4 The Simulation Argument

For many people, the intuition that felt duration matters – and objective duration doesn't – appears to receive its strongest support from thinking about simulated minds that have the same experiences but run through those same experiences at different objective speeds.[21] In this penultimate section, I want to to explain why we shouldn't be drawn in by intuition pumps of that kind.

---

[21] Cases involving biological minds run through the same physical transitions at different objective speeds might also be imagined and used to mount similar arguments, but are harder to imagine and describe, I find. Thanks to Brad Saad for this observation.

4.1 Setting Out the Argument

Let's start by making the key thought experiment fully concrete. Suppose that nanoprobes are scattered throughout Simone's nervous system when she's born. Throughout her life, the probes exhaustively record her neural activity, right down to the molecular level. Nothing is missed. After her death, these recordings allow people to reconstruct a richly detailed digital simulation of Simone's neural activity throughout her life (compare Kurzweil 2005: 198–204).

Assuming for the sake of argument that the physical basis of consciousness is substrate neutral (Chalmers 1996: 247–275) and that the recordings are sufficiently detailed, each simulation of Simone's life may then be expected to generate a new person who undergoes the very same experiences Simone had while alive. Suppose multiple simulations are created and the simulations are run at different speeds by adjusting the clock speed of the simulation hardware. Each time, exactly the same neural processing is simulated via exactly the same computations. Exactly the same experiences occur. Each of these lives feels exactly alike from the inside. On some runs, the simulation completes in less than an hour. On others, it runs for decades.

Assume that we can isolate that aspect of a person's lifetime welfare that depends only on her phenomenal states. Call this the *phenomenal component of lifetime welfare*. Thus, the contributions valenced experiences make to lifetime welfare in virtue of how they feel belong to the phenomenal component of lifetime welfare, but need not exhaust it.

In respect of Simone and her simulated copies, many find it extremely natural to think that the speed at which we run the simulation doesn't affect the phenomenal component of lifetime welfare. The subject herself can't tell the difference. It's the same experiences that play out every time, and those experiences feel exactly the same from the inside. Surely, then, lifetime welfare should remain fixed, at least insofar as we have in mind those features of lifetime welfare that supervene on the subject's phenomenal states. Given the assumption that the objective time occupied by the good and bad experiences of the person changes dramatically across the different runs,[22] objective duration can't be

[22]This assumption might conceivably be rejected if we accept the kind of view discussed at the end of the previ-

what matters. What counts must instead be the subjective experience of time's passage, which doesn't vary overall across the different runs, since each feels just like every other.

We can think of the argument as an abductive inference that relies on two central claims: 1. The phenomenal component of lifetime welfare is the same for Simone and each of her copies; 2. This is best explained by the assumption that the durations of pains and pleasures and other valenced experiences are to be equated with their felt duration for the purposes of welfare assessment.

I can think of two sub-arguments supporting 1, which I'll try to rebut in what follows. But I also think 2 can be challenged on the basis of what I've argued already in this paper.

Granting 1, we may conclude that the objective duration of a valenced experience is irrelevant of itself to the phenomenal component of lifetime welfare for Simone and her copies, assuming that the objective time occupied by the good and bad experiences of the person changes dramatically across the different runs. But in light of what I've argued in the previous section, it remains hard to say why felt duration should matter. We should be open to the possibility that neither subjective nor objective duration has any fundamental ethical significance. It may be that the duration of a pain normally points to deeper underlying properties of the experience that determine how good or bad it is, but that have nothing essentially to do with time's passage. These properties might be shared across experiences that are subjectively indistinguishable or that involve the same computations, but need not otherwise be shared between experiences that have the same felt duration.

4.2 Subjective Indistinguishability

Here's the first argument that might be used to support the claim that the phenomenal component of lifetime welfare is the same for Simone and each of her copies. The argument appeals to the fact that the experiences of these people are subjectively indistinguishable and the claim that subjectively indistinguishable experiences contribute equally to the phenomenal component of lifetime welfare.[23]

---

ous section.

[23] I'm grateful to Hayden Wilkinson for pressing me to respond to this kind of argument.

Fans of Williamson's *anti-luminosity argument* (Williamson 2000: 93–113) might raise their hackles at the second assumption, since they think there is virtually nothing - not even pain - such that you can always distinguish between its presence and absence. I want to set aside the voluminous controversy surrounding that argument (see McGlynn 2014: 145–166). Rather than attacking the argument from subjective indistinguishability as unsound, I'm going to show that it's question-begging.

What does it mean to say that two experiences that extend in time are subjectively indistinguishable? In some sense, an experience of six minutes of pain is subjectively indistinguishable from an experience of a minute of pain that comes equipped with false memories of having been in pain for the preceding five minutes also. But the latter is clearly not as bad. On the other hand, those experiences are only *really* subjectively indistinguishable in their final minute. In their first five minutes, it would probably be easy to tell which you're having. I'm therefore going to interpret subjective indistinguishability in terms of what a given subject is in a position to know based on evidence that supervenes on her internal states *at any time* during either experience.

Here is the proposed interpretation of subjective indistinguishability. Suppose that we have two experiences, x and y, had by subjects S_1 and S_2, respectively. These experiences will be said to be subjectively indistinguishable just in case the following conditions are satisfied. At any point in time during S_1's experience of x, based on evidence supervening on her internal states at that time, S_1 is not in a position to know that her experience contemporaneously satisfies any property, F, where F is satisfied by x but not y at the given point in time; and, similarly, at any point in time during S_2's experience of y, based on evidence supervening on her internal states at that time, S_2 is not in a position to know that her experience contemporaneously satisfies any property, G, where G is satisfied by y but not x at the given point in time.

When subjective indistinguishability is interpreted this way, in order to say whether two experiences are subjectively indistinguishable, we need to determine what the subject is in a position to know and what her experience is like at times that we define to be 'the same point in time' across the two experiences. We therefore need to rely on some kind of mapping between the temporal parts of each experience.

That's where the problem arises. We need to make a choice about what kind of mapping among temporal parts we're going to rely on in defining indistinguishability. One possible proposal is that a mapping of the relevant sort exists among the sets of times associated with two different experiences just in case those times are identically marked by two identically constructed, synchronized clocks that begin recording the proper time interval at the onset of each experience. That is not going to deliver the result that the experiences of Simone and her copies are subjectively indistinguishable, given that those experiences will often be wildly different at points that are recorded as the same based on clock time in the experiencing subject's frame of reference. On the other hand, the proposal to rely on a mapping of that kind might seem to beg the question in favour of a view on which objective duration is what really matters. That seems right to me. But that point cuts both ways. To get the result that the experiences of Simone and her copies are subjectively indistinguishable, you seem to need to rely on a mapping among temporal parts based on their location in time as subjectively experienced, which seems to beg the question in favour of views on which subjective duration is the thing that's ethically significant.

Rather than appealing to a privileged mapping among the parts of an experience, in light of which we need to take a stand on whether the mapping represents similarity in objective or subjective time, maybe we should say that two experiences are subjectively indistinguishable just in case there exists *some* bijection relative to which the subject(s) can't tell apart times mapped to one another. Since we make no prior assumption that the mapping is a mapping in terms of equivalence of subjective time, the charge of the begging the question might look like it's been defeated.

However, I think it's clear we shouldn't say that two experiences are subjectively indistinguishable just in case there's some bijection relative to which the subject(s) can't tell apart any times mapped to one another. At the least, we shouldn't say that if we also want to say that subjectively indistinguishable experiences contribute equally to the phenomenal component of lifetime welfare.

Here's why. Consider again a kind of thought experiment discussed at the end of 2.2 (and compare Lee 2013: 18–19). Imagine a mind that continuously experiences a pain of a constant intensity throughout a one second time period, measured in terms of the inter-

val $[0, 1]$. It has no memory outside of an ultra-short-term memory buffer required for its experiences to have a specious present of a few milliseconds as their intentional objects. At every point throughout the one second interval, its experience is as of having a pain of the given intensity and of that pain filling the corresponding specious present. The mind begins to exist at the start of the one second interval and ceases to exist abruptly at its end. Compare the same mind who continuously experiences a pain of the same intensity occupying the same specious present at every point in time throughout a two second period, measured in terms of the interval $[0, 2]$. Because the experience is uniform in character throughout the time it fills and is so across the two cases, the monotone increasing bijection $f : [0, 1] \rightarrow [0, 2]$ defined as $f(t) \coloneqq 2t$ makes the two experiences subjectively indistinguishable on the current proposal. But, intuitively, the second experience has to be worse because it's exactly like the first at every point in time, but twice as long.

4.3 Computational Equivalence

Here's a second argument for the claim that the phenomenal component of lifetime welfare is the same for Simone and each of her copies.

The argument appeals to the fact that the same computation occurs each time. In addition, it relies on the idea that if what goes on in the head and gives rise to the mind is Turing-style computation, then the phenomenal component of lifetime welfare must be the same whenever the same underlying computational processes are reproduced, regardless of elapsed time. A Turing machine model of computation, after all, has nothing in it corresponding to the flow of time. The machine's computation is defined in terms of the sequence of configurations yielded by the starting configuration. There's nothing in the model corresponding to the amount of time the machine spends in a given configuration or requires when transitioning from one configuration to another. If the mind is essentially Turing machine-style computation, the time the computation needs in order to complete when physically instantiated ought to be irrelevant to the character of mind, and so to the phenomenal component of lifetime welfare.

I think we should reject this argument, not because it assumes the computational theory of mind, but because it fetishizes a well-known limitation of abstract computational

models of the mind.

The observation that Turing machine models of computation are atemporal is considered an objection to the hypothesis that mental activity is Turing-style computation (see Van Gelder 1995; Van Gelder and Port 1995; Clark 1998; Eliasmith 2003; Weiskopf 2004; Piccinini 2010). To quote Van Gelder and Port (1995: 19): "Since cognitive processes unfold in real time, any framework for description of cognitive processes that hopes to be fully adequate to the nature of the phenomena must be able to describe not merely *what* processes occur but *how* those processes unfold in time." More controversially, Van Gelder and Port (1995: 21) go on to assert that "it is futile to attempt to weld temporal considerations onto an essentially atemporal kind of model." Defenders of the computational theory of mind roundly reject this claim, arguing that computational models can be successfully supplemented with auxiliary assumptions about the time required for the brain to implement a given computational operation in order to bring the model closer to reality (Clark 1998; Weiskopf 2004; Piccinini 2010).

Clearly, conscious experience unfolds in time, and any fully adequate account of consciousness as an empirical phenomenon needs to be able to account for this. On some views, temporal properties of conscious experience end up playing an essential role in determining the representational content of consciousness, because experience represents temporal properties in the world based on a mirroring principle: an experience of change requires a change in experience, and, more generally, any experience representing any temporal property must itself instantiate the property it represents (Phillips 2014). However, we need not agree with this controversial position in order to recognize the more basic point that atemporal models of the basis of consciousness should be presumed to be incomplete.

5 Conclusion

Because individual welfare is naturally thought of as relative to a person's own perspective on the world, it feels intuitive that the temporal extent of a valenced experience should be measured according to how long it feels, and not how long it really is, at least for the

purposes of reckoning how well someone's life goes. I don't accept that view. I can't see any way to make sense of what felt duration consists in on which that idea turns out to be plausible. It's not that I find myself unable to make sense of the nature of felt duration.[24] Rather, it's that insofar as I'm able to make sense of what subjective duration consists in, it seems to me irrelevant in itself for how well someone's life goes.

Of course, we shouldn't mistake a failure of imagination for an insight into necessity, and I admit the intuition is hard to shake off. It's also worth emphasizing that it's possible in principle and consistent with measuring valenced experience by its objective duration that we should pay attention to the subjective experience of time's passage and attach more weight to the pains and pleasures of those who experience more subjective time within a given hour or day. In relation to the cognitivist theory of felt duration, it seems to be an open question whether increases in the speed at which an individual's thoughts, imaginings, and rememberings seem to move in relation to external events while they are in pain may tend to bring about increases in the intensity of pain's felt unpleasantness. When it comes to theories of felt duration based on perceptual frame speed, we should also keep in mind that it could be the case that our ordinary conception of objective duration relies on a false assumption in treating conscious experience as continuously 'on'. If so, pains that fill more subjective time might fill more objective time also, although the empirical assumptions required to support that conclusion strike me as extremely speculative on current evidence.

---

[24] In this respect, my argument contrasts with that of Lee (2013), whose skepticism toward the idea that subjective time is the proper measure of a pain's duration for ethical purposes derives from skepticism about the existence of a meaningful way of accounting for subjective duration (see also Lee 2017). So far as I can tell, I haven't committed myself to the existence of a subjective time metric of the kind Lee denies. His central argument is that the existence of an appropriate subjective time metric is ruled out by his preferred 'atomist' theory of time experience (Lee 2014), which is roughly equivalent to the retentionalist account of time experience, noted at the bottom of page twenty-two. On this view, he argues, "there is no such thing as phenomenal duration, understood as a global phenomenal metric that gives extended portions of experience an intrinsic felt duration; only atoms have intrinsic felt duration." (12) As noted at the end of section 2.2, I reject the idea that the subjectively experienced duration of a given time interval must be understood in terms of how long the interval as a whole feels, where this involves some apprehension of the interval in its entirety, as opposed to something constructed by aggregating the felt duration of the interval's proper parts.

References

Adler, M. 2012. *Well-Being and Fair Distribution: Beyond Cost-Benefit Analysis*. Oxford: Oxford University Press.

Ambron, R. 2022. *The Brain and Pain: Breakthroughs in Neuroscience*. New York, NY: Columbia University Press.

Armstrong, D. M. 1968. *A Materialist Theory of the Mind*. London: Routledge Kegan Paul.

Arstila, V. 2012. Time slows down during accidents. *Frontiers in Psychology 3*(196).

Bain, D. 2014. Pains that don't hurt. *Australasian Journal of Philosophy 92*(2): 305–320.

Balog, K. 2012. In defense of the phenomenal concept strategy. *Philosophy and Phenomenological Research 84*(1): 1–23.

Bostrom, N. 2014. *Superintelligence: Paths, Dangers, Strategies*. Oxford: Oxford University Press.

Bostrom, N. and E. Yudkowsky. 2014. The ethics of artificial intelligence. In *The Cambridge Handbook of Artificial Intelligence*, eds. K. Frankish and W. M. Ramsey, 316–334. Cambridge: Cambridge University Press.

Brandt, R. B. 1979. *A Theory of the Good and the Right*. Oxford: Oxford University Press.

Broad, C. D. 1923. *Scientific Thought*. London: Routledge and Kegan Paul.

Chalmers, D. J. 1996. *The Conscious Mind: In Search of a Fundamental Theory*.

Chalmers, D. J. 2004. Phenomenal concepts and the knowledge argument. In *There's Something About Mary: Essays on Phenomenal Consciousness and Frank Jackson's Knowledge Argument*, eds. P. Ludlow, Y. Nagasawa, and D. Stoljar, 269. MIT Press.

Clark, A. 1998. Time and mind. *Journal of Philosophy 95*(7): 354.

Clay, E. R. 1882. *The Alternative: A Study in Psychology*. London: Macmillan & Co.

Cowan, N. 2001. The magical number 4 in short-term memory: A reconsideration of mental storage capacity. *Behavioral and Brain Sciences 24*(1): 87–114.

Dainton, B. 2000. *Stream of Consciousness: Unity and Continuity in Conscious Experience*. London: Routledge.

Davis, J. H. 2018. Meditation and consciousness: Can we experience experience as broken? In *Routledge Handbook of Consciousness*, ed. R. J. Gennaro, 436–438. Abingdon: Routledge.

de Charms, R. C., F. Maeda, G. H. Glover, D. Ludlow, J. M. Pauly, D. Soneji, J. D. E. Gabrieli, and S. C. Mackey. 2005. Control over brain activation and pain learned by using real-time functional MRI. *Proceedings of the National Academy of Sciences 102*(51): 18626–18631.

Dorsey, D. 2013. Desire-satisfaction and welfare as temporal. *Ethical Theory and Moral Practice 16*(1): 151–171.

Dorsey, D. 2017. Why should welfare 'fit'? *Philosophical Quarterly 67*(269): 685–24.

Drissi-Daoudi, L., A. Doerig, and M. Herzog. 2019. Feature integration within discrete time windows. *Nature Communications* 10.

Droit-Volet, S., P. Trahanias, and M. Maniadakis. 2017. Passage of time judgments in everyday life are not related to duration judgments except for long durations of several minutes. *Acta Psychologica* 173: 116–121.

Droit-Volet, S. and J. A. Wearden. 2016. Passage of time judgments are not duration judgments: Evidence from a study using experience sampling methodology. *Frontiers in Psychology* 7: 176.

Einstein, A. 1905. Zur elektrodynamik bewegter körper. *Annalen der Physik 322*(10): 891–921.

Eliasmith, C. 2003. Moving beyond metaphors: Understanding the mind for what it is. *Journal of Philosophy 100*(10): 493–520.

Flaherty, M. G. 1999. *A Watched Pot: How We Experiene Time*. New York, NY: New York University Press.

Gibson, J. J. 1979. *The Ecological Approach to Visual Perception*. Boston, MA: Houghton Mifflin.

Goff, P. 2011. A posteriori physicalists get our phenomenal concepts wrong. *Australasian Journal of Philosophy 89*(2): 191–209.

Grahek, N. 2007. *Feeling Pain and Being in Pain*. Cambridge, MA: MIT Press.

Hall, A. and V. Tiberius. 2016. Well-being and subject dependence. In *The Routledge Handbook of Philosophy of Well-Being*, ed. G. Fletcher, 175–186. Abingdon: Routledge.

Hall, R. 1989. Are pains necessarily unpleasant? *Philosophy and Phenomenological Research 49*(June): 643–59.

Hanson, R. 2016. *The Age of Em: Work, Love, and Life When Robots Rule the Earth*. Oxford: Oxford University Press.

Heathwood, C. 2005. The problem of defective desires. *Australasian Journal of Philosophy 83*(4): 487–504.

Hellström, C. and S. Carlsson. 1997. Busy with pain: Disorganization in subjective time in experimental pain. *European Journal of Pain 1*(2): 133–139.

Herzog, M., L. Drissi-Daoudi, and A. Doerig. 2020. All in good time: Long-lasting postdictive effects reveal discrete perception. *Trends in Cognitive Sciences 24*(10): 826–837.

Herzog, M. H., T. Kammer, and F. Scharnowski. 2016. Time slices: What is the duration of a percept? *PLOS Biology 14*(4): 1–12.

Holtug, N. 2010. *Persons, Interests, and Justice*. Oxford: Oxford University Press.

Hurka, T. 2010. Asymmetries in value. *Noûs 44*(2): 199–223.

Husserl, E. 1991 [1893-–1917]. *On the Phenomenology of the Consciousness of Internal Time*, ed. and trans. J.B.Brough. Dordrecht: Kluwer.

James, W. 1886. The perception of time. *Journal of Speculative Philosophy 20*(4): 374–407.

James, W. 1890. *The Principles of Psychology, Volume 1*. New York, NY: Henry Holt and Company.

Kind, A. 2003. What's so transparent about transparency? *Philosophical Studies 115*(3): 225–244.

Kurzweil, R. 2005. *The Singularity Is Near: When Humans Transcend Biology*. New York, NY: Viking.

Lee, G. 2013. Subjective duration. https://philpapers.org/archive/LEESD.pdf.

Lee, G. 2014. Temporal experience and the temporal structure of experience. *Philosophers' Imprint* 14.

Lee, G. 2017. Making sense of subjective time. In *The Routledge Handbook of Philosophy of Temporal Experience*, ed. I. Phillips, 157–168. London: Routledge.

Loar, B. 1990. Phenomenal states. *Philosophical Perspectives* 4: 81–108.

Lockwood, M. 2005. *The Labyrinth of Time: Introducing the Universe*. Oxford: Oxford University Press.

Mathy, F. and J. Feldman. 2012. What's magic about magic numbers? Chunking and data compression in short-term memory. *Cognition 122*(3): 346–362.

Matthews, W. J. and W. H. Meck. 2016. Temporal cognition: Connecting subjective time to perception, attention, and memory. *Psychological Bulletin 142*(8): 865–907.

Mayerfeld, J. 1999. *Suffering and Moral Responsibility*. Oxford: Oxford University Press.

McGlynn, A. 2014. *Knowledge First?* New York, NY: Palgrave Macmillian.

Melzack, R. and P. D. Wall. 1988. *The Challenge of Pain: Updated Second Edition*. London: Penguin Books.

Merino-Rajme, C. 2014. A quantum theory of felt duration. *Analytic Philosophy 55*(3): 239–275.

Miller, G. A. 1956. The magical number seven, plus or minus two: Some limits on our capacity for processing information. *The Psychological Review* 63: 81–97.

Mioni, G., A. Shelp, C. T. Stanfield-Wiswell, K. A. Gladhill, F. Bader, and M. Wiener. 2020. Modulation of individual alpha frequency with tACS shifts time erception. *Cerebral Cortex Communications 1*(1): 1–9.

Montemayor, C. 2012. *Minding Time: A Philosophical and Theoretical Approach to the Psychology of Time*. Boston, MA: Brill.

Moore, G. E. 1903. *Principia Ethica*. Cambridge: Cambridge University Press.

Nagel, T. 1974. What is it like to be a bat? *Philosophical Review* (4): 435–450.

Noë, A. 2006. *Action in Perception*. Cambridge, MA: MIT Press.

Noyes, R. and R. Kletti. 1977. Depersonalization in response to life-threatening danger. *Comprehensive Psychiatry 17*(4): 375–384.

Papineau, D. 2002. *Thinking About Consciousness*. Oxford: Oxford University Press.

Parfit, D. 1991. *Equality or Priority?* The University of Kansas: The Lindley Lecture.

Paul, L. A. 2010. Temporal experience. *Journal of Philosophy 107*(7): 333–359.

Phillips, I. 2012. Attention to the passage of time. *Philosophical Perspectives 26*(1): 277–308.

Phillips, I. 2013. Perceiving the passing of time. *Proceedings of the Aristotelian Society 113*(3): 225–252.

Phillips, I. 2014. The temporal structure of experience. In *Subjective Time: The Philosophy, Psychology, and Neuroscience of Temporality*, eds. V. Arstila and D. Lloyd. Cambridge, MA: MIT Press.

Piccinini, G. 2010. The resilience of computationalism. *Philosophy of Science 77*(5): 852–861.

Pitcher, G. 1970. Pain perception. *Philosophical Review 79*(3): 368.

Pratchett, T. 1990. *Truckers: The First Book of the Nomes*. London: Corgi.

Prosser, S. 2016. *Experiencing Time*. Oxford: Oxford University Press.

Purves, D., J. Paydarfar, and T. Andrews. 1996. The wagon wheel illusion in movies and reality. *Proceedings of the National Academy of Sciences of the United States of America 93*(8): 3693–3697.

Railton, P. 1986. Facts and values. *Philosophical Topics 14*(2): 5–31.

Rammsayer, T. H. and S. J. Troche. 2014. In search of the internal structure of the processes underlying interval timing in the sub-second and the second range: A confirmatory factor analysis approach. *Acta Psychologica* 147: 68–74.

Rashbrook, O. 2013. The continuity of consciousness. *European Journal of Philosophy 21*(4): 611–640.

Rawls, J. 1971. *A Theory of Justice*. Cambridge, MA: Harvard University Press.

Rey, A. E., G. A. Michael, C. Dondas, M. Thar, L. Garcia-Larrea, and S. Mazza. 2017. Pain dilates time perception. *Nature Scientific Reports 7*(1): 15682.

Rosati, C. S. 1996. Internalism and the good for a person. *Ethics 106*(2): 297–326.

Samaha, J. and B. Postle. 2015. The speed of alpha-band oscillations predicts the temporal resolution of visual perception. *Current biology 25*(22): 2985–2990.

Schouten, J. F. 1967. Subjective stroboscopy and a model of visual movement detectors. 44–55. Cambridge, MA: MIT Press.

Schukraft, J. 2020. The subjective experience of time: Welfare implications. https://forum.effectivealtruism.org/posts/qEsDhFL8mQARFw6Fj/the-subjective-experience-of-time-welfare-implications.

Shannon, C. 1948a. A mathematical theory of communication. *Bell System Technical Journal 27*(3): 379—-423.

Shannon, C. 1948b. A mathematical theory of communication. *Bell System Technical Journal 27*(4): 623–666.

Shulman, C. and N. Bostrom. 2021. Sharing the world with digital minds. In *Rethinking Moral Status*, eds. S. Clarke, H. Zohny, and J. Savulescu, 306–326. Oxford: Oxford University Press.

Somov, P. G. 2000. Time perception as a measure of pain intensity and pain type. *Journal of Back and Musculoskeletal Rehabilitation 14*(3): 111–121.

Stetson, C., M. P. Fiesta, and D. M. Eagleman. 2007. Does time really slow down during a frightening event? *PLOS ONE 2*(12): 1–3.

Strawson, G. 2009. *Selves: An Essay in Revisionary Metaphysics*. Oxford: Oxford University Press.

Sumner, L. W. 1996. *Welfare, Happiness, and Ethics*. Oxford: Oxford University Press.

Thönes, S. and D. Oberfeld. 2015. Time perception in depression: A meta-analysis. *Journal of Affective Disorders* 175: 359–372.

Tye, M. 1995. A representational theory of pains and their phenomenal character. *Philosophical Perspectives* 9: 223–39.

Tye, M. 2003a. *Consciousness and Persons: Unity and Identity*. Cambridge, MA: MIT Press.

Tye, M. 2003b. A theory of phenomenal concepts. *Royal Institute of Philosophy Supplement* 53: 91–105.

Van Gelder, T. 1995. What might cognition be, if not computation? *Journal of Philosophy 92*(7): 345–381.

Van Gelder, T. and R. F. Port. 1995. It's about time: An overview of the dynamical approach to cognition. In *Mind as Motion: Explorations in the Dynamics of Cognition*, eds. T. van Gelder and R. Port. Cambridge, MA: MIT Press.

VanRullen, R. 2016. Perceptual cycles. *Trends in Cognitive Sciences 20*(10): 723–735.

von Baer, K. E. 1864. *Reden gehalten in wissenschaftlichen Versammlungen und kleinere Aufsätze vermischten Inhalts. Erster Theil, Reden*. St. Petersburg: H. Schmissdorf.

Wearden, J. H. 2015. Passage of time judgements. *Consciousness and Cognition* 38: 165–171.

Wearden, J. H. 2016. *The Psychology of Time Perception*. London: Palgrave Macmillan.

Wearden, J. H., A. O'Donoghue, R. Ogden, and C. Montgomery. 2014. Subjective duration in the laboratory and the world outside. In *Subjective Time: THe Philosophy, Psychology, and Neuroscience of Temporality*, eds. V. Arstila and D. Lloyd, 287–306. Cambridge, MA: MIT Press.

Weiskopf, D. A. 2004. The place of time in cognition. *British Journal for the Philosophy of Science 55*(1): 87–105.

White, P. A. 2018. Is conscious perception a series of discrete temporal frames? *Consciousness and Cognition* 60: 98–126.

Williamson, T. 2000. *Knowledge and its Limits*. Oxford: Oxford University Press.

Wilson, K. A. 2022. Windows on time: Unlocking the temporal microstructure of experience. *Review of Philosophy and Psychology*. Online Preprint.

Yong, E. 2022. *An Immense World: How Animal Senses Reveal the Hidden Realms Around Us*. London: The Bodley Head.