

# Once More, Without Feeling

Andreas Mogensen (Global Priorities Institute,  
University of Oxford)

Global Priorities Institute | January 2025

*GPI Working Paper No. 2-2025*

Please cite this working paper as: Mogensen, A. Once More, Without Feeling. *Global Priorities Institute Working Paper Series*, No. 2-2025. Available at <https://globalprioritiesinstitute.org/mogensen-once-more-without-feeling>



# Once More, Without Feeling

## Abstract

I argue for a pluralist theory of moral standing, on which both welfare subjectivity and autonomy can confer moral status. I argue that autonomy doesn't entail welfare subjectivity, but can ground moral standing in its absence. Although I highlight the existence of plausible views on which autonomy entails phenomenal consciousness, I primarily emphasize the need for philosophical debates about the relationship between phenomenal consciousness and moral standing to engage with neglected questions about the nature of autonomy and its possible links to consciousness, especially if we're to face up to the ethical challenges future AI systems may pose.

## 1 Introduction

Some things matter morally in their own right and for their own sake. That includes you and me. We aren't *mere* things. We merit concern and respect. If you prick us, not only do we bleed; we are wronged. We have *moral standing* (or *moral status*<sup>1</sup>).

What does it take to have moral standing? Many think the capacity for *phenomenal consciousness* is necessary (Singer 1993; Korsgaard 2018; Nussbaum 2022). To have moral standing, they think, there needs to be something it's like to be you. However, not everyone is convinced, and the dissenters appear to be growing their ranks (Levy 2014b; Kagan 2019; Bradford 2022; Shepherd 2024). Who is in the right?

Most of the recent discussion of this issue focuses on the relationship between phenomenal consciousness and being a *welfare subject*. Roughly speaking, a welfare subject is a being whose life can go better or worse for them. It's plausible that being a welfare subject is sufficient for moral

---

<sup>1</sup>I use these terms interchangeably.

standing. Is it also necessary? If not, we might be missing an important piece of the puzzle.

On its face, not all the obligations we owe to other people aim at promoting their welfare. Quinn (1984) distinguishes between the *morality of respect* and the *morality of humanity*. These aren't rival moral theories. Instead, they are partially overlapping systems of obligation that respond to different morally significant properties of humans and non-human animals. In this context, 'humanity' denotes the virtue of being humane or beneficent. The morality of humanity is thus concerned with promoting others' welfare; not just human welfare, but welfare more broadly. The morality of respect instead involves constraints on our behaviour that stem from recognition of the authority of rational agents to direct their own lives, even if they do so imprudently.

Suppose we grant that there are these two distinct dimensions to morality. It's plausible that there are morally statused beings, including many non-human animals, who fall outside the scope of the morality of respect and are protected only by the morality of humanity (Quinn 1984: 51; McMahan 2002: 245–246). Are there also metaphysically and/or nomologically possible beings who fall outside the scope of the morality of humanity and are protected only by the morality of respect – individuals whose autonomy merits respect but who are not welfare subjects? If so, what does this imply about the relationship between moral standing and consciousness?

These are the questions I'll address in this paper, arguing that there are indeed possible individuals who are protected only by the morality of respect, and exploring the potential implications for what we should think about the link between consciousness and moral status. Here's the plan. In section 2, I outline a collection of conditions that I take to be jointly sufficient for an agent to be autonomous. In section 3, I argue that being a welfare subject isn't necessary to satisfy those conditions. In section 4, I argue that welfare subjectivity is also unnecessary for someone's autonomy to merit respect. In section 5, I outline the potential implications of my argument for the relationship between phenomenal consciousness and moral standing. Finally, in section 6, I summarize my conclusions and discuss their practical significance.

## **2 Portrait of An Autonomous Agent**

The term 'autonomy' is highly polysemous (Feinberg 1989: 27–51; Arpaly 2003: 117–130). To a first approximation, it will be used here to denote a psychological capacity for self-government in virtue

of which an individual can be credited as the author of her own thoughts and actions, but which is vulnerable (at least in beings like us) to subversion by internal impediments like addiction, mania, and the like. My goal here won't be to defend a complete theory of autonomy, so understood. Instead, I'll outline a conception of what an autonomous agent might be like, couched in terms of a set of conditions that I take to be sufficient for autonomy. I make no assumptions as to their necessity.

Suppose, then, that we have an individual before us. Call her *Artemis*. We won't assume Artemis is a human being. Perhaps she is an alien or a robot. What conditions would suffice to make Artemis autonomous?

I begin from the well-worn idea that a person is autonomous if she is in a position to determine her behaviour based on rational reflection. Thus, we assume that Artemis has this ability. She isn't simply led about blindly by impulses. This may be understood to entail that Artemis formulates higher-order desires through which she identifies herself with some of her first-order desires and repudiates others, as on hierarchical theories of autonomy (Frankfurt 1971; Neely 1974; Dworkin 1976). Nonetheless, it seems that autonomy can also be exercised in the reflective evaluation of courses of action to which a person's first-order desires need not already incline her (Watson 1975: 218–219). Moreover, I assume that rational reflection involves forming normative and/or evaluative judgments that can guide behaviour (Watson 1975; Ekstrom 1993; Savulescu 2007). Formulating higher-order desires in principle requires no exercise of a capacity for normative or evaluative thought (Frankfurt 1971: 13 n.6).

Autonomy may be thought to require not only that an individual exercises a capacity for normative and/or evaluative thought, but that she does so competently. In addition to tending to satisfy coherence constraints associated with structural rationality, an autonomous individual may be required to be substantively rational to the point of being at least moderately reasons-responsive (Wolf 1990: 67–93; Fischer and Ravizza 1998: 62–91). I'll therefore assume that Artemis not only has a capacity for normative and/or evaluative thought, but that she is reliable in forming reasonably accurate and coherent normative and/or evaluative judgments and in responding appropriately. I also assume that she has a suitable degree of epistemic competence when it comes to gaining relevant empirical information about the world she inhabits, the options she confronts, and their possible consequences (Mele 1995: 179–182; Killmister 2013; Pugh 2020: 131–136).

To ensure that an autonomous agent is able to make up her mind independent of controlling influences external to herself, a number of authors treat autonomy as a property that doesn't supervene on an individual's current state of mind and instead requires the absence of certain forms of manipulation in her past (Christman 1991; Mele 1995; Fischer and Ravizza 1998; Haji 1998). We can stipulate that Artemis has the right kind of history. This might entail, for example, that there are no unsheddable values determinative of her behaviour whose presence is explained by processes that have bypassed her capacity for control over her mental life (Mele 1995: 144–176).

This completes my outline of the conception of autonomy on which I'll rely. Before I move on, let me quickly address the following concern. The conception of an autonomous agent outlined above might strike you as overly rationalist. It might seem to put too much emphasis on reason and rational reflection, and not enough – arguably, none at all – on emotion. This might seem overly fusty. Unlike Plato or Kant, contemporary writers are likely to see emotions as capable of making a positive contribution to autonomy and not merely as impulses threatening to divest command of the will from the self (Frankfurt 1998: 129–141; Friedman 2003: 3–29; compare Damasio 1994).

Keep in mind, however, that the conception of an autonomous agent outlined above doesn't involve the proposal of any necessary conditions. If you think that acts governed by passions that people experience as moving them against their own better judgment can be autonomous (Watson 1987; Frankfurt 1998: 129–141; Shoemaker 2003), that's completely compatible with everything I've said. All I claim is that the conditions I've outlined suffice to make an agent autonomous.

### **3 Autonomy and Welfare Subjectivity**

Granting that a person satisfies the sufficient conditions for autonomy noted in the previous section, does it follow that she is a welfare subject? I'll assume that to be a welfare subject is to be an entity that can accrue welfare goods (or bads) (Bradford 2022). Given this assumption, does it follow that Artemis must be a welfare subject?

It does not. Here's my main argument. A capacity for affective states is necessary for accruing welfare goods (and bads), but nothing in Artemis's portrait entails that she's capable of undergoing affective states. While I have slightly more to say in defence of a negative answer than just this argument, most of this section will focus on unpacking and defending its premises.

### 3.1 Affect as necessary for welfare subjectivity

First off, what do I mean by an *affective* state? Affective states are a class of psychological states of which emotions are the paradigm instance. Other commonly recognized affective states include moods and valent bodily sensations, like itches and pains. I assume that affective states are defined at least in part by their degree of positive or negative valence and their constituent level of arousal (Russell 1980, 2003).

Why think that states of this kind are necessary for the ability to accrue welfare goods (and bads<sup>2</sup>)? One major reason is simply that a very wide range of theories of welfare tell us so.

Most obviously, there is *hedonism*, according to which a person's welfare level is determined entirely by the extent to which her experiences are pleasant as opposed to unpleasant (Crisp 2006; Bradley 2009; Bramble 2016). There are also hybrid views on which pleasure or enjoyment is necessary for anything to constitute a welfare good, but the welfare values of enjoyments depend on more than just the degree of felt pleasantness (Darwall 2002: 73–104; Feldman 2004; Kagan 2009).

We can also count the view that welfare consists in *happiness* (Sumner 1996: 138–183; Feldman 2010). Even if we reject the equation of happiness with pleasure often seen in the welfare literature (e.g., Parfit 1984: 493–494), it's highly plausible that happiness necessarily has an affective component (Sumner 1996: 145–146; Haybron 2008: 105–151).

What about the *desire-fulfilment* theory? Very roughly, this is the view on which something is a welfare good for a person if and only if she desires it under the right conditions (Rawls 1971; Brandt 1979; Griffin 1986; Heathwood 2005). Like 'autonomy', 'desire' is polysemous. There is a sense in which a person can be said to desire any outcome that they intentionally act to bring about (Davidson 1963; Nagel 1970: 29–30). There is another sense of 'desire' according to which the desired object has to evoke some kind positive affective response in the desirer in the right conditions (Vadas 1984; Lewis 1988; Sumner 1996: 113–122; Heathwood 2019). I think there are strong arguments for thinking that we should interpret the desire-fulfilment theory as working with the second, affective conception of desire (Sumner 1996: 120–121; Heathwood 2019; Pallies

---

<sup>2</sup>As Kagan (2014) notes, philosophical theories of welfare bads are significantly underdeveloped by comparison with theories of welfare goods. For this reason, most of the discussion focuses on welfare goods and parenthetical allusions to welfare bads will become more sparse as the argument gets going.

2022). Moreover, those arguments generalize to hybrid theories of welfare on which what's good for someone is having some objectively desirable thing, like achievement, while strongly desiring it (Parfit 1984: 501–502).<sup>3</sup>

*Value-fulfilment* theories are closely related to desire-fulfilment theories, and tend to explicitly endorse a necessary role for affect. Very roughly, these are views on which something is a welfare good for a person if and only if she values it under the right conditions (Raibley 2010; Yelle 2014; Tiberius 2018; Dorsey 2021). Accounts of valuing typically say that it has both cognitive and affective dimensions (Raibley 2010; Yelle 2014; Tiberius 2018: 37–46).<sup>4</sup> For example, Tiberius (2018: 50) claims that “values are comprised partly of emotional dispositions”.

Finally, we come to *objective list* theories, on which there are a plurality of welfare goods, at least some of which are welfare goods for a person independent of any pro-attitude she might have toward them (Arneson 1999; Hurka 2011; Fletcher 2013; Hooker 2015). While theories of this kind can provide a basis for rejecting my claim that affective states are necessary to accrue welfare goods, they needn't. Many of the goods recognized by objective list theories entail the occurrence of positive emotions or other affective states. Such goods include pleasure, happiness, loving relationships, virtue, and the aesthetic appreciation of beauty. Assuming that achievement requires the exertion of *effort* (Bradford 2015: 26–63), it's plausible that achievement requires a capacity for *negative* affect (Bermúdez forthcoming). Any list theory that mentions only subsets of the goods noted here therefore accords with the idea that a capacity for affect is necessary for accruing welfare goods.

That completes my survey. If you accept any of the theories I've highlighted, you should already be convinced that a capacity for affective states is necessary for the ability to accrue welfare goods. Moreover, all the arguments that have been marshalled to support any element in this big array of theories indirectly support that claim.

Setting aside arguments tailored to support any of the particular theories covered in my survey,

---

<sup>3</sup>Here's one such argument (Heathwood 2019: 677–678), stated very succinctly. There are some things we do because we feel we morally ought to, although they don't inspire any positive feelings in us. Intuitively, cases of this sort involve choosing against our self-interest. That's easy to explain if the welfare good of desire satisfaction requires the desired object to evoke positive emotions in the desirer and difficult to explain otherwise.

<sup>4</sup>Notably, Dorsey (2021) dissents, accepting instead a purely cognitive account of valuing. I discuss Dorsey's view later on in this section.

can anything be said in general about why a theory of welfare ought to impose a capacity for affective experience as a necessary condition for accruing welfare goods? I think so. In particular, we can appeal to the popular idea that what's good for someone has got to *resonate* with them and can't simply *leave them cold* (Railton 1986; Rosati 1996).

Intuitively, we're left cold when something fails to elicit any positive emotion or other affective response in us. Affect is what gives heat to our mental lives. If you aren't able to respond affectively to anything at all, then everything leaves you cold. The assumption that welfare goods have to resonate with us and can't leave us cold therefore provides intuitive support for the claim that you need affective states to accrue welfare goods (compare Velleman 1998; Heathwood 2019; Fanciullo 2023; Smithies forthcoming). A welfare good (*qua* welfare good) must entail the occurrence of some corresponding affective response in its beneficiary.

Dorsey (2021) will contest this claim, since he appeals to the resonance constraint to motivate a value-fulfilment theory of welfare, but rejects the claim that valuing has any essential affective dimension. It will help my case to explain why I think Dorsey is in error.

Why does Dorsey think that valuing has no essential affective dimension? He asks us to imagine that Stan is a successful lawyer in the big city. Stan judges that it would be good for him to move to a small town that's close to his family and friends, but he's so burnt out from work that the prospect of moving doesn't elicit any positive affective response in him. In spite of his burn-out, Dorsey claims, Stan values living in a small town.

This is not a counter-example to Tiberius's claim that "values are comprised partly of emotional dispositions". Dispositions can be masked and finked, as a result of which they fail to manifest under normal stimulus conditions (Choi and Fara 2021). Stan is most naturally understood to have a positive sentiment toward small-town living that's masked by burn-out. If we imagine otherwise, the intuition that small-town living is a welfare good for Stan fades.

Suppose, for example, that Stan isn't burnt out at all. He's energized by work and by the noise and chaos of the big city. He sincerely believes that life in a small town, close to family and friends, would be best for him. His beliefs are coherent and he has an accurate picture of what small-town living is like. But whenever he visits, something of the life goes out of him. He's glum. He never smiles. He finds getting out of bed more and more effortful. Intuitions may differ, but, as I see it, Stan's belief about the prudential value for him of small-town living is simply in error. Small-town



life just doesn't resonate with him (compare Haybron 2008: 199–200).

That completes my case for thinking that a capacity for affective states is necessary for accruing welfare goods (and bads). Obviously, I won't have convinced everyone. Some philosophers accept objectivist theories of welfare that entail the falsity of the resonance constraint under any reasonable interpretation. While I personally feel confident in rejecting those theories, I can afford to be more ecumenical. Once I've finished defending the second key premise of the main argument, I'll give reasons to think that welfare goods recognized by these theories might still be inaccessible to Artemis.

### **3.2 Affect as unnecessary for autonomy**

The second premise of my main argument says that the conditions identified in section 2 as sufficient to make a person autonomous do not entail a capacity for affective states. Since affective states weren't mentioned explicitly in defining any of the criteria, I assume the burden of proof is on those who want to reject the second premise. In what follows, I'll consider three different philosophical positions that might lead you to reject that premise and explain why I think that would be a mistake.

Let's start with the *Humean theory of motivation* (Hume 1739-40 [1978]: 413–418; Lewis 1988, 1996; Smith 1994: 92–129; Sinhababu 2009, 2017). Humeans believe that any motivating reason consists of a desire and an appropriately related means-end belief. They also typically believe that any complete psychological explanation for the acquisition of a desire requires a desire to be mentioned in the explanans (Sinhababu 2009). Moreover, they typically believe that desires and beliefs are 'distinct existences' (Lewis 1988, 1996; Smith 1994: 119–121). This may be understood to mean that there are no propositions,  $p$ ,  $q$ , such that believing  $p$  and desiring  $q$  are type-identical. An even stronger version of this idea would be the claim that no mental state is a token of both a type of belief and a type of desire.

I think the view that Artemis is an autonomous agent but incapable of experiencing affective states probably precludes us from accepting the full package of these views. However, it might not be obvious why we have to reject any part of it. Granted, the behaviour of an autonomous agent needs to be explicable in terms of motivating reasons. It follows that we have to reject the Humean theory of motivation when understood as generalizing over the space of all metaphysically and/or

nomologically possible minds and formulated using an affective sense of ‘desire.’ But why think the theory should be understood that way? Its most prominent contemporary defenders do not.<sup>5</sup> So what’s the problem?

Here’s the problem. Not every kind of state we call a ‘desire’ can serve as a motivating reason when paired with a corresponding means-end belief (Quinn 1993; Smithies and Weiss 2019; Smithies forthcoming). Motivating reasons do not merely explain behaviour. They rationalize it; they make action intelligible (Davidson 1963; Smith 1994: 94–98). Arguably, not everything we call a ‘desire’ can do that when conjoined with a suitable means-end belief.

Consider the *radio man*, introduced by Quinn (1993). Suppose I am disposed to turn on any radios that I learn are not currently on. I derive no enjoyment from switching on radios. I don’t think anything good will come of it. In some sense of ‘desire’, I can be said to desire that radios are switched on, purely in virtue of the fact that I’m disposed to bring it about that they are. Nonetheless, my disposition does not seem to contribute to rationalizing my behaviour. As Quinn (1993: 237) notes: “It may help explain, causally, why I turn on a particular radio, but it does not make the act sensible”.

What would make the act sensible? According to Quinn, what the radio man is missing is a normative or evaluative belief: something like “the *thought* that the direction in which I am psychologically pointed leads to something good” (Quinn 1993: 242; emphasis in text). According to Smithies (forthcoming), what the radio man is missing is a disposition to experience affectively valenced feelings of desire to turn on radios. Both Quinn and Smithies claim to identify conditions that are both necessary and sufficient to rationalize behaviour. I think they’re wrong about that.<sup>6</sup> I think it’s more plausible that both normative or evaluative thought and affect can rationalize

---

<sup>5</sup>Lewis (1988, 1996) and Smith (1994) make no appeal to an affective conception of desire. Sinhababu (2009) does, but restricts the Humean theory of motivation to a claim about human psychology.

<sup>6</sup>Smithies (forthcoming) makes a compelling case that normative or evaluative beliefs are unnecessary. His argument that evaluative beliefs cannot rationalize action is unconvincing, in my view. Smithies claims that beliefs stand in need of justification, and, when justified, merely transmit justification from the reasons on which they’re based. On the one hand, this is a strange objection to raise given that Smithies elsewhere defends *doxastic conservatism* and maintains that “believing a proposition is defeasible evidence for its truth.” (Smithies 2019: 118). Even if we reject doxastic conservatism, we can accommodate Smithies’ objection given only modest revisions to the claim that behaviour can be rationalized by evaluative beliefs. We just add the stipulation that those beliefs must be epistemically justified. Smithies gives no argument that justifying reasons for evaluative beliefs can only be provided by affective states, and work in moral epistemology suggests a wide range of plausible alternatives, including perception (McGrath 2004), intuitions understood as non-perceptual seemings (Huemer 2005), and testimony (Sliwa 2012).

action. But either way, it ends up being hard to defend full-blown Humeanism alongside the claim that some possible autonomous agents completely lack affect.

The conflict is obvious if we appeal to the capacity of affect to rationalize behaviour. Suppose, on the other hand, that we think a person's behaviour can also be rendered intelligible by the presence of appropriate evaluative beliefs. The Humean theory of motivation entails that all motivating reasons have only two constituents: a desire and a suitably related means-end belief. Where do evaluative beliefs fit in?

We can render the hypothesis that some possible motivating reasons have evaluative beliefs as constituents consistent with the core of the Humean theory of motivation via the claim that it is possible for one and the same token mental state to be both a desire and an evaluative belief, and so to play the role of a desire in constituting some possible agent's motivating reason (compare McDowell 1978; McNaughton 1988: 106–117; Gregory 2021). This is consistent in principle with all claims associated with the Humean theory of motivation except, of course, the claim that no mental state is a token of both a type of belief and a type of desire. But so far as I know, no one who defends the Humean theory of motivation accepts that claim, at least not about the space of all metaphysically and/or nomologically possible minds.<sup>7</sup> I conclude that the hypothesis that autonomous agents need not have affective states ought to be acceptable in principle to actual Humeans.

The foregoing discussion assumed that some possible normative or evaluative judgments neither constitute nor are constituted by affective states. Some theories about the nature of evaluative judgments or affects might be thought to give us reason to reject that claim.

Consider, firstly, *non-cognitivism* (Ayer 1936; Gibbard 1990, 2003; Blackburn 1993, 1998). According to non-cognitivists, normative or evaluative judgments are not robustly representational states. Unlike garden-variety beliefs, they do not purport to represent the world out there as being a certain way. When it comes to giving a positive characterization of normative or evaluative judgments, it may be thought that non-cognitivists are committed to analyzing them as affective

---

<sup>7</sup>Smith (1994: 119) argues that Humeans ought to affirm “that it is always possible for agents who are in some particular belief-like state not to be in some particular desire-like state” Similarly, Lewis (1988: 323) attacks the claim that “some beliefs are, at least, necessarily conjoined with corresponding desires.” These remarks entail that no type of belief is also a type of desire, but are consistent with the possibility that some mental state is a token of both a type of belief and a type of desire. As noted previously, Sinhababu (2009, 2017) restricts the Humean theory of motivation to a theory of human psychology.

states.<sup>8</sup>

They are not. Consider the analysis of normative judgments about rationality put forward by Gibbard (1990: 55–82). According to Gibbard, to think that something is rational is to accept a norm that permits it. Roughly speaking, accepting a norm requires having a motivational tendency to act in accordance with the norm and a disposition to sincerely avow it in normative discussion. So understood, norm acceptance doesn't appear to have any essential affective component, and Gibbard (1990: 80) notes that the psychological capacities involved “fit a traditional conception of the faculty of reason.”

A very different view might incline us to reject the possibility of normative or evaluative judgments that do not constitute affective states. According to *pure cognitivism* (Solomon 1973; Nussbaum 2001), emotions just are normative and/or evaluative judgments. They are not psychological complexes in which judgments feature alongside other mental states. The judgment itself is the emotion in its entirety. Since we stipulate that Artemis is able to guide her behaviour through the exercise of normative and/or evaluative judgment, it may be thought to follow from pure cognitivism that Artemis must be able to exhibit emotions, since emotions are nothing over and above normative or evaluative judgments of a certain kind.

This doesn't follow. The qualifier ‘of a certain kind’ is crucial. Pure cognitivism is not the view that all normative or evaluative judgments are emotions. Emotions are supposed to have a particular content or subject-matter, being necessarily concerned with personal status (Solomon 1976) or with the importance to us of factors beyond our control (Nussbaum 2001). Moreover, emotions are judgments made in a particular way (Solomon 1988): for example, they must be hasty and dogmatic (Solomon 1973), exhibiting kinetic properties such as *rushing toward* a particular content (Nussbaum 2001: 44-45), and perhaps even imbued with phenomenal qualities associated with bodily experiences (Nussbaum 2001: 60; Solomon 2003).

By insisting that emotions aren't merely judgments individuated by their contents, but also by the distinctive manner in which they're made, pure cognitivists have the beginnings of a response to the stock objection that any judgment associated with emotional agitation can also be made dispassionately (Griffiths 1997: 29–30; Roberts 2003: 97–102). At the same time, they undercut any

---

<sup>8</sup>In his *Stanford Encyclopedia* article on moral sentimentalism, Kauppinen (2022) actually defines non-cognitivism as the claim that “moral thoughts are constituted by sentiments.”

inference from pure cognitivism to the conclusion that Artemis must be able to exhibit emotions. Even if she can formulate normative or evaluative judgments with the distinctive contents characteristic of emotions, she need not be able to make those judgments in the distinctively emotional way required to realize an affective state.

### 3.3 Other welfare goods

I noted at the end of section 3.1 that some philosophers are likely to remain unconvinced of the first premise of the main argument of this section, because they recognize objective welfare goods that have no necessary connection with affect. I promised then to give some reasons to think that those welfare goods might anyway be inaccessible to Artemis. I'll focus on three putative welfare goods: *nature fulfilment*, *knowledge*, and (something we might call) *autonomy*.

By 'nature fulfilment', I mean, roughly, the development and exercise of those capacities central to an individual's nature. According to a *perfectionist* theory of welfare (Aristotle 350 BCE [1980]; Kraut 2007; compare Hurka 1993), nature fulfilment is the one and only welfare good. Plausibly, nature fulfilment in general has no essential connection with affect, being a putative welfare good accessible to plants, protists, and other mindless organisms (Kraut 2007: 6–8).

Must nature fulfilment be realizable in principle by any autonomous agent? It depends what it means to have a 'nature'. Not every property that serves to characterize an individual's essence is a plausible candidate for the kind of property whose full development is a welfare good. Instead, Hurka (1993: 16) settles on the idea that our nature consists in those "essential properties that humans could not have if they were not living; they presuppose life, or are necessarily distinctive of living things." Similarly, Kraut (2007: 131) defends the view that "flourishing is a biological phenomenon" and that living things flourish "by developing properly and fully". But if only organisms have natures of the relevant kind, it follows that the good of nature fulfilment need not be realizable by Artemis, since Artemis need not be a living thing, as opposed to, say, a robot.<sup>9</sup>

Consider, next, the putative welfare good of knowledge (Hurka 2011; Hooker 2015). Must this good be accessible to Artemis? I actually don't think that's at all obvious. Knowledge requires

---

<sup>9</sup>If we adhere to a form of perfectionism on which the welfare value of nature fulfilment isn't basic and instead admits of explanation in terms of some more fundamental welfare value, then Artemis might be able to accrue other forms of this deeper kind of value. On the question of whether the welfare value of nature fulfilment admits of this kind of explanation, see Bradford (2017: 11–21).

belief. We might well think it requires *full* or *categorical* belief. However, Artemis might be so constituted that her doxastic attitudes are restricted to (non-extremal) *credences* or *partial beliefs*. If so, we either have to give up the idea that knowledge is a welfare good accessible to Artemis or follow Moss (2016) in allowing that states involving (non-extremal) credences can constitute knowledge of non-standard contents.<sup>10</sup>

Last but not least, some philosophers say that something they call ‘autonomy’ is a welfare good (Sher 1997: 176–179; Wall 1998: 129–130; Hooker 2015; Pugh 2020: 235–258). It’s important to be clear that this doesn’t automatically commit them to thinking that Artemis is a welfare subject. Recall that ‘autonomy’ is polysemous. At least some authors who assert that autonomy is a welfare good seem to have in mind a different conception of autonomy – one involving a suitably high degree of independence of mind, for example. Thus, Wall (1998: 128) writes that autonomous people “neither drift through life ..., nor adopt projects and pursuits wholesale from others.” Independence of mind, so understood, might be inaccessible to Artemis, who might be so constituted that her own reflective judgment instructs her to defer to others on questions about what sort projects and pursuits should structure her life (compare Arpaly 2003: 120).

I don’t claim that we can always raise credible doubts about the availability of any putative welfare good to some possible autonomous agent. Some of the philosophers who view autonomy as a welfare good do have in mind the power of rational self-government I’ve assumed that Artemis has (e.g., Hooker 2015; Pugh 2020). Nonetheless, I hope I’ve shown that there’s more room to raise such doubts than we might initially have expected.

## 4 Moral Standing Without Welfare Subjectivity

Let’s assume Artemis is able to guide her actions based on rational judgments about what to do, but she never feels joy or sorrow, nor pleasure, nor pain. Let’s assume that, as a result, she’s incapable of being benefited or harmed. Can she nonetheless be wronged?

A number of philosophers treat a capacity to be benefited or harmed as necessary for moral standing. According to DeGrazia and Millum (2021: 176–177), “only beings with interests have

---

<sup>10</sup>It might also be argued that there is some other kind of epistemic good that is a welfare good and to which Artemis has access in spite of the fact that the welfare good of knowledge is inaccessible to her.

moral status”. Similarly, Chalmers (2022: 339) defines an individual with moral status as one “whose welfare we need to take into account in our moral deliberations.” These claims suggest that we can’t have any obligations to individuals who aren’t welfare subjects.

In this section, I’ll present two arguments against this claim. The first is an argument by analogy that draws on assumptions about what we owe to other human beings. The second is a more direct argument that appeals to intuitions about the ethics of killing agents like Artemis.

#### **4.1 The Argument from Humanity**

The first argument relies on the following observation, to which I appealed already in the introduction. At least some of the obligations we owe to our fellow human beings don’t seem to be grounded in concern for their welfare. Instead, they seem to be grounded in respect for their autonomy and their status as separate, self-governing agents with their own lives to lead. For example, we can be obligated not to engage in paternalistic acts that we know will benefit others to the greatest extent possible, because acts of that kind may fail to show adequate respect for others’ autonomous preferences (Feinberg 1989: 52–97; Shiffrin 2000; Enoch 2016). On a view we might attribute to Kant (1785 [1998]), all moral obligations are grounded in this way. I assume only the weaker claim that some are.

Assuming that our reasons to respect others’ autonomy are independent in principle of concern for their welfare, it stands to reason that autonomous agents should be able to be wronged by acts that interfere with the exercise of their capacity for rational self-government, even if they can’t be harmed or benefited at all. What reason could we have to think otherwise?

We might deny that there are any reasons to respect others’ autonomy that cannot be explained in terms of reasons to promote others’ welfare if we endorse some form of welfarist consequentialism. However, I take myself to have set such views aside already, since I’ve helped myself to the assumption that at least some of our obligations are grounded in respect for autonomy, rather than concern for others’ welfare. In doing so, I have assumed that a certain kind of deontological moral theory gives the correct account of at least some of our moral obligations.

Even if we grant that the morality of respect isn’t straightforwardly reducible to the morality of humanity, we might still think that our reasons to respect others’ autonomy presuppose that they are welfare subjects. Pugh (2020: 239) claims that “[a]utonomy itself ... is conditional on the agent’s

beliefs about what constitutes the good life, since autonomous choices must be grounded in part by these beliefs.” We might think that our reasons to respect others’ autonomy ultimately reduce to reasons to respect their ability to live in accordance with their own views about what would make their lives go best. Therefore, we might claim, there is no moral reason to respect others’ autonomy except insofar as their choices express self-regarding prudential preferences. Together with the epistemic competence condition on autonomy set out in section 2, this may be taken to imply that only autonomous welfare subjects are owed a duty of respect.

This argument is unsound. We have reasons to respect others’ autonomy over and above our reasons to respect their self-regarding prudential preferences. For example, we have reason to respect their ability to live in accordance with their ethical convictions. Thus, it’s wrong of you to paternalistically add meat to your vegan housemate’s lentil stew in small quantities she can’t detect, whether she foregoes meat for prudential or moral reasons, and whether you do so because you’re worried about her well-being or because you think she needs more protein to be able to care adequately for her ailing mother (compare Shiffrin 2000: 216–217).

Respect for autonomy is therefore not just a matter of respecting each person’s ability to pursue their own conception of the good life for them. There is no sound argument from that premise to the conclusion that only autonomous agents who are welfare subjects merit respect.

## 4.2 The Argument from Vulcanity

Do we think that it would be permissible to kill or destroy a being like Artemis without her consent for the sake of some trivial good? Suppose there were large numbers of agents just like her. Do we think it would be permissible to kill arbitrarily many of them in order to provide a hamster with a juicy huckleberry?

The questions I’ve just posed parallel questions Chalmers’ poses about the ethics of killing *Vulcans* (Chalmers 2022: 343–345). A Vulcan is here understood as a conscious creature quite like you or me, but who is incapable of experiencing any affective states. They might have a rich conscious life and various intellectual and moral goals, but they never feel pain or pleasure, nor happiness, nor sorrow. Artemis might well be imagined as a Vulcan, so understood.

Chalmers appeals to the possibility of Vulcans in arguing against the the view that affective consciousness or *sentience* is necessary for moral standing (Singer 1977, 1993; DeGrazia 2021). He



notes that this would make it acceptable to kill arbitrarily many Vulcans in order to provide small benefits to sentient individuals. Chalmers thinks this would be monstrous. I'm inclined to agree.

That intuition favours the claim for which I'm arguing. Vulcans are autonomous beings who aren't welfare subjects, at least not if the argument of section 3 is to be believed. Positing that agents like that have moral standing in virtue of their autonomy allows us to explain why it's nonetheless wrong to kill them. Moreover, that explanation is continuous with an independently plausible account of the ethics of killing persons (McMahan 2002: 232-265).

Intuitively, the prohibition against killing is not exhaustively explained by the need to take others' welfare into account in moral decision-making. It's seriously wrong to painlessly kill a person against their wishes even if they're so badly off that their lifetime welfare isn't increased by continuing to live. Moreover, variation in how much a person is harmed in death through deprivation of the goods of further life doesn't appear to give rise to similar variation in the wrongness of killing them. Except perhaps at the extremes, it seems no more or less wrong to kill a person regardless of whether she is young or old, happy or depressed.<sup>11</sup>

A plausible explanation for these observations is that the wrongness of killing persons is explained primarily in terms of the moral standing conferred by autonomy, as opposed to by considerations of non-maleficence. As McMahan (2002: 245) puts it, the capacity for rational self-determination confers on persons "a form of moral standing that demands respect but is entirely independent of the value their lives have, or of the goods their lives might contain." A view of this kind straightforwardly accommodates the intuition that it is wrong to kill Vulcans – including the sense that "a Vulcan matters about as much as an ordinary human" (Chalmers 2022: 344) – without requiring us to suppose that Vulcans are welfare subjects.

How else could we explain our intuitions about the wrongness of killing Vulcans? Chalmers (2022: 340–345) could be read as claiming that a complete explanation need only appeal to the fact that Vulcans are conscious. He writes: "My own view (shared with many others) is that what bestows moral status is *consciousness*." (341) He doesn't explicitly point to any other property of Vulcans as relevant to their possession of moral standing.

It seems implausible that consciousness alone does all the work. We can imagine very simple conscious minds that don't obviously have any degree of moral standing. For example, we can

---

<sup>11</sup>See Lippert-Rasmussen (2007) for critical discussion of this claim.

imagine a sessile filter feeder whose mental life is exhausted by a simple, unvarying, neutral quality, perhaps an auditory experience of a low, muffled tone.<sup>12</sup> Does this animal have higher moral status than an otherwise similar unconscious organism growing next to it on the sea-bed? I'm not sure it does; but even if it does, it's surely not by much. There's got to be something more going on with Vulcans. Autonomous agency remains an extremely plausible candidate for what that something more might be.<sup>13</sup>

Smithies (forthcoming) suggests that Vulcans might matter morally in a purely derivative sense by virtue of being valued highly enough by others who have moral status, and this may explain why it is wrong to destroy them. I don't find this convincing either. We can stipulate that there aren't any sentient individuals who value Vulcans in the relevant possible world. My intuitions remain the same. If someone goes around killing Vulcans in order to provide treats to hamsters, what they do is wrong.

## 5 Autonomy and Phenomenal Consciousness

In the previous section, I argued that there can be autonomous agents who aren't welfare subjects but who are nonetheless owed a duty of respect and non-interference. However, I didn't argue that *any* metaphysically possible autonomous agent must have moral standing.

Even if we grant that being a welfare subject isn't necessary for an autonomous agent to be in a position where they can be wronged by acts that interfere with the exercise of their capacity for self-government, we might think that *phenomenal consciousness* is. Moreover, we might find it unobvious that the sufficient conditions for autonomy set out in section 2 entail the presence of phenomenal consciousness.

What exactly is the relationship between autonomy and phenomenal consciousness? And what

---

<sup>12</sup>Similar creatures are imagined by Kagan (2019: 13) and Lee (2018: 663).

<sup>13</sup>At one point, Chalmers (2022: 340) writes that "systems with a minimal degree of consciousness (ants?) have only a minimal degree of moral status." He might therefore argue that in fact, possession of consciousness *can* do all the work, because although both Vulcans and the sessile filter feeder are conscious to some degree, Vulcans are conscious to a higher degree, and so have a higher moral status. However, it's controversial whether phenomenal consciousness really admits of degrees. What would it even mean to say that there's more something it's like to be a Vulcan than an ant? Many philosophers claim that the idea makes no sense (Bayne et al. 2016; Carruthers 2019; Birch 2020). Even philosophers who think it's coherent to suppose that consciousness comes in degrees aren't sure exactly what these degrees would consist in (Lee 2023). Therefore, it seems best to reserve judgment on whether degrees of consciousness have any moral significance.

does this tell us about the relationship between phenomenal consciousness and moral status? These questions are addressed in an exploratory fashion in this section.

## 5.1 Autonomy without consciousness

One possible view is that phenomenal consciousness is unnecessary for autonomy: there are some metaphysically possible autonomous agents for whom everything within is dark. Since the particular cognitive capacities that I identified as sufficient for autonomy make no explicit mention of phenomenal consciousness, this doesn't seem especially implausible on its face, assuming that cognitive states in general can occur in the absence of any capacity for consciousness (Block 1978; Chalmers 1996).

If you're otherwise convinced by my argument so far and take the view that some possible autonomous agents aren't conscious, you might nonetheless try to hold on to the claim that consciousness is the key to moral standing. You might claim that autonomy confers moral standing only in the presence of consciousness. Stated otherwise, phenomenal consciousness is an enabling condition for the fact that someone is autonomous to constitute a reason not to interfere with their self-governing activities.<sup>14</sup> This is analogous to a view on which there are welfare subjects who aren't conscious, but only the interests of conscious welfare subjects ground moral reasons (Basl 2019), or on which individuals who aren't conscious can accrue welfare goods, but only conscious individuals are welfare subjects (Lee forthcoming).

One major worry about views of this kind is that they resemble a kind of prejudice against the unconscious (compare Goodpaster 1978; Bradford 2022). They tell us that when something happens to a special group, it matters morally in and of itself. However, when the very same thing happens to a different group, it doesn't matter at all, because the members of the second group lack a special property, although this doesn't affect their ability to share in the insults that demand our concern when visited on members of the first group. A moral view that enjoins this kind of discrimination may be hard to swallow.<sup>15</sup>

---

<sup>14</sup>On reasons and enabling conditions, see Dancy (2004: 38–52).

<sup>15</sup>See Lee (2022) for a reply to this concern, offered in defence of the view that only conscious individuals are welfare subjects.

## 5.2 Appealing to Reductive Theories of Consciousness

Perhaps there is some necessary connection between autonomy and phenomenal consciousness. Although it may be unobvious that the cognitive capacities I identified as sufficient for autonomy entail phenomenal consciousness, might that simply reflect our ignorance about the essential nature of phenomenal consciousness?

Suppose we endorse *a posteriori* physicalism: we think the mental reduces without remainder to the physical in light of identity statements that are unknowable *a priori* (Loar 1990; Papineau 2002; Levin 2008; Balog 2012). We might think, for example, that consciousness has a hidden essence consisting in a particular kind of functional organization. The cognitive capacities that suffice for autonomy may entail the instantiation of that kind of functional organization. If so, it is metaphysically impossible that any autonomous agent lacks phenomenal consciousness. However, this would be knowable only *a posteriori*, on the basis of a scientifically confirmed theory of consciousness, and so highly unobvious.

To my mind, two classes of theories of consciousness stand out as most likely to vindicate this approach: *global workspace theories* (Baars 1988; Dehaene 2014; Carruthers 2019) and *higher-order theories* (Lycan 1996; Carruthers 2000; Rosenthal 2005; Lau 2022).

Global workspace theories rest on a picture of the mind as composed of a number of specialized subsystems. These subsystems operate in parallel but are able to pass information to a low-capacity central store. Information passed to the central store is integrated to yield a coherent, holistic representation and made globally available, allowing for the flexible and coordinated control of cognition and behaviour. Very roughly, global workspace theories maintain that mental contents enter into phenomenal consciousness by virtue of being made globally available in this way.

A global workspace architecture may be thought necessary for someone to be able to guide their thoughts and actions according to beliefs and desires that are adequately reasons-responsive (Levy 2014a: 109–130). Without a global workspace, a mind would be a loosely interconnected system of specialized processors, operating in parallel and without guidance from a coherent, integrated point of view that can flexibly direct cognition and behaviour. A mind of that kind would seem unable to achieve the kind of rational self-government required for autonomy, being unable to coordinate the activity of its different parts in accordance with a unitary perspective that takes account of all the agent's reasons.

Consider, next, higher-order theories of consciousness. Very roughly, these are views on which mental contents are able to gain access to phenomenal consciousness by virtue of being the contents of first-order states that are the objects of higher-order states. These higher-order states may be understood as percepts belonging to an ‘inner sense’ modality (Lycan 1996; Lau 2022: 154–157) or as higher-order thoughts (Rosenthal 2005; Carruthers 2000).

As the popularity and influence of hierarchical theories of autonomy attests, there is something compelling in the idea that autonomy requires a capacity for higher-order awareness of first-order mental states. Without this ability, a person wouldn’t be able to step back from their beliefs and desires in order to reflectively assess whether what they believe is really true or what they desire really desirable. It’s therefore natural to suppose that higher-order theories of consciousness support the idea that autonomous agents necessarily instantiate the kind of cognitive architecture that suffices for consciousness.

Both of these lines of argument seem plausible enough that they are worth taking seriously. Still, they are hardly immune to doubt. The argument that a global workspace architecture is necessary for autonomy seems less compelling if taken to mean that autonomous agents must have minds that are modular to some non-trivial extent, as opposed to just needing a coherent, integrated point of view with global effects on cognition and behaviour. There might also be a mismatch between the kind of higher-order states needed for autonomous agency and those that give rise to consciousness according to our preferred higher-order theory. For example, autonomous agents might be able to get by with higher-order thoughts, whereas consciousness may derive from internal monitoring by a non-conceptual ‘inner sense’ faculty.

We might also have general reasons for doubting the overall approach. If we accept *a posteriori* physicalism, we will judge that *zombies* are metaphysically impossible, but conceptually possible, where a zombie is understood as a creature that is physically identical to some conscious human being, but lacks conscious experiences altogether (Kirk 1974; Chalmers 1996: 94–99). It is not unreasonable to think that our ethical theory ought to give the right answers in conceptually possible scenarios, even if they’re metaphysically impossible (Cutter 2017; Hayward 2019). For example, a theory of filial duties presumably ought to give acceptable verdicts about conceptually possible but metaphysically impossible worlds where I have a different ancestry than I actually do.<sup>16</sup> A view on

---

<sup>16</sup>For an argument that it is metaphysically impossible for a person to have a different ancestry than they actually do, see

which autonomy is sufficient for moral standing, whereas the relationship between consciousness and autonomy is merely a *posteriori* necessary, may seem to commit us to the wrong verdicts about some conceptually possible scenarios involving zombies.

How so? Assuming that there is no *a priori* entailment between the cognitive capacities that define autonomy and phenomenal consciousness, a zombie twin of yours could be autonomous without having any capacity for conscious experience. Nonetheless, we might think that zombies would have to have a very different moral status from Vulcans. Being without consciousness, we might think that they would have no moral status at all. Chalmers (2022: 340–341) seems to think that it's at least reasonable to believe that you can permissibly kill arbitrarily many zombies in order to provide arbitrarily small benefits to conscious individuals.<sup>17</sup> If we share Chalmers' intuition, we may be persuaded that something is awry with a moral theory that treats autonomy as sufficient in itself for moral standing, even if autonomy metaphysically entails phenomenal consciousness.

### 5.3 Linking Consciousness and Autonomy Without Reductionism

Is there a way of linking consciousness and autonomy that doesn't presuppose a reductive theory of consciousness?

One possibility is to appeal to *the phenomenal intentionality hypothesis*, on which all intentional content derives from the intentional content of phenomenally conscious states (Bourget and Mendelovici 2019). Its adherents typically assert that intentional states cannot occur without being at least potentially conscious or otherwise appropriately related to conscious states of the same subject (Searle 1990, 1992; Horgan and Tienson 2002; Loar 2003). Since autonomous agents necessarily have intentional attitudes, it follows straightforwardly that all autonomous agents must have the capacity for conscious experience.

Suppose, however, that we're inclined to reject versions of the phenomenal intentionality hypothesis that have this kind of implication, for the familiar reason that they have trouble accommodating the extent to which cognitive science traffics in representational states that are disconnected from consciousness (Kriegel 2011: 190–198; Smithies 2012). Is there anything else we might say

---

the classic discussion in Kripke (1980: 110–114).

<sup>17</sup>Of course, he also thinks that zombies are metaphysically possible.

instead?

As I see it, the most promising way to link consciousness and autonomy without presupposing a reductive theory of consciousness goes via an argument that links consciousness and rationality.<sup>18</sup> I've stipulated that autonomy involves formal and substantive rationality in both epistemic and practical domains. But what does it take to be rational? Here, I'll focus on epistemic rationality, partly for the sake of simplicity, and partly because the conception of autonomy on which I'm relying emphasizes the idea that the agent's cognitive appraisals are in the driving seat. Could an argument be made that epistemic rationality requires phenomenal consciousness?

In fact, multiple arguments for that claim can be made. For example, a traditional form of epistemic internalism requires that agents must always be in a position to know what propositions they have justification to believe. Both Chisholm (1989) and Smithies (2019) argue on that basis that epistemic justification supervenes on mental states individuated by their relationship to phenomenal consciousness because those are the only empirical factors always knowable to us.<sup>19</sup>

A different line of argument, of which I'm more persuaded, adopts a bottom-up approach that directly interrogates the different sources of knowledge and their relationship to phenomenal consciousness. Consider, for example, the question of how perceptual states justify belief. An intuitively plausible answer is that they do so (at least in part) by virtue of having a distinctive presentational phenomenology (Pryor 2000, 2004; Huemer 2001; Chudnoff 2011; Smithies 2019: 74–112). When I open my eyes and see the garden covered in snow, what justifies me in believing it has snowed has something to do with what it's like for me to be having this experience and the way it seems to present the world to me. That intuition can be bolstered by reflecting on cases that mimic ordinary perceptual experience but involve no perceptual phenomenology (Smithies 2019: 76–90). For example, a patient with blindsight does not seem to have justification to believe that an object of which they have no subjective experience is being presented in their blind field, although they can reliably detect its presence or absence when prompted to guess.

---

<sup>18</sup>Given suitable assumptions about the relationship between cognition and rationality, this might also provide a basis on which to deny that individuals who lack any capacity for phenomenal consciousness have any cognitive states: see Smithies (2012). I won't make any such assumption in the present discussion.

<sup>19</sup>Stated in the terminology popularized by Williamson (2000), this is the assumption that mental states individuated by their relationship to phenomenal consciousness are *luminous*. Williamson famously argues that no mental states are luminous. For Smithies' reply, see Smithies (2019: 345–379)

I'm inclined to believe that what's true of perception is true in general: the justification for any basic belief that *p* involves a conscious state that presents *p* as true in roughly the way perceptual experiences present their contents as true, in line with the doctrine of *phenomenal conservatism* (Huemer 2001, 2006, 2007, 2013). This is obviously a good deal more controversial than the claim that presentational phenomenology explains the ability of perception to justify belief. However, the goal here is merely to exhibit one route to the conclusion that phenomenal consciousness is necessary for the justification of belief. As I've already noted, others are possible.

It's also worth noting that the conclusion that justified perceptual belief requires phenomenal consciousness in and of itself gets us a big part of the way to the conclusion that all autonomous agents must exhibit phenomenal consciousness. It rules out the possibility that my zombie twin is epistemically rational. More generally, it's challenging to imagine an autonomous agent that doesn't have any need to rely on perception to inform their beliefs about the world. If you agree that justified perceptual belief requires phenomenal consciousness, you should agree that phenomenal consciousness is necessary for autonomous agency *modulo* some relatively recherché possibilities.

The arguments to which I've appealed so far to support the claim that consciousness is necessary for epistemic rationality are associated with epistemic internalism, understood as the claim that epistemic justification is fully determined by 'internal' factors (BonJour 1980; Feldman and Conee 2001). Nonetheless, the claim that phenomenal consciousness is necessary for epistemic rationality does not commit us to internalism and is perfectly consistent with thinking that epistemic justification also depends on 'external' factors, such as reliability (compare Alston 1988). Moreover, insofar as the claim that phenomenal consciousness is necessary for epistemic rationality turns out to conflict with particular externalist theories of epistemic justification, such as process reliabilism (Goldman 1979, 1986), there are ways of softening it so as to remove this inconsistency while keeping the argument intact.

Philosophers have long suspected that internalists and externalists are talking past one another (BonJour 1980; Goldman 1987; Alston 2005). Bach (1985), Engel (1992), and Littlejohn (2009) argue that internalists are concerned with what it takes for a *person* to be justified in holding the belief that *p* (*personal justification*), whereas externalists are concerned with what it takes for the *belief* that *p* to be justified (*doxastic justification*). It's thus open to us to claim that there is some sense of 'epistemic rationality' for which phenomenal consciousness is necessary, and that 'epistemic



rationality,' so understood, defines a necessary condition for autonomy, but may not be the target of analysis for epistemological theories like process reliabilism. Indeed, this claim seems highly plausible if we suppose that process reliabilism is a theory of doxastic justification, whereas there is a distinct concept of personal justification in light of which an agent can be appraised as epistemically rational or irrational. Since autonomy is a property of persons, it stands to reason that it ought to be tied specifically to epistemic rationality understood in terms of personal justification.

## 6 Conclusion and Practical Significance

I've argued that both welfare subjectivity and autonomy are able to confer moral status, and that autonomy doesn't entail welfare subjectivity. Although my conclusions on the relationship between autonomy and phenomenal consciousness are more tentative, I've shown that there are plausible views on which autonomy entails phenomenal consciousness.

As I noted in my introduction, debates about the relationship between phenomenal consciousness and moral standing have focused largely on the relationship between consciousness and welfare subjectivity.<sup>20</sup> If what I've argued is correct, this tells only half the story, at most. There is a separate set of questions about the relationship between autonomy and phenomenal consciousness with which we need to engage.

Setting aside recommendations for the conduct of philosophical inquiry, what is the practical significance of the conclusions drawn in this paper? Suppose everything I've said is true. So what? What difference does it make?

In some sense, the ethical importance of determining whether some kind of being has moral standing seems clear enough. Nonetheless, I can think of two reasons you might be sceptical of the practical significance of the conclusions drawn in this paper. The first is that autonomous agents who aren't welfare subjects seem merely hypothetical. The second is that whether we think of such agents as welfare subjects might not seem to make a difference to how we ought to treat them.

Let's start with the first worry, and let's grant for the sake of argument that there aren't cur-

---

<sup>20</sup>For some exceptions, see McLaughlin (2018: 363), Kagan (2019: 34–35), and Sinnott-Armstrong and Conitzer (2021: 279–285).

rently any autonomous agents who aren't welfare subjects. Nonetheless, there might be soon enough. A recent expert report concludes that if consciousness supervenes on functional organization, then "conscious AI systems could realistically be built in the near term." (Butlin et al. 2023: 6) Although conscious, these systems might lack any capacity for happiness, suffering, pleasure, or pain, while still being rational agents with goals and plans. They might have neither physical nor virtual bodies, whereas a range of theories in philosophy, psychology, and neuroscience analyse affective states as involving interoceptive registrations of bodily states and physiological changes (James 1884; Damasio 1994; Prinz 2004; Barrett 2017). In this vein, the neuroscientist Michael Graziano (2019) argues that because "emotion is anchored to the physical sensations in the body" (131), we should expect that "[e]arly attempts at artificial consciousness will probably lack convincing emotion. ... The Hollywood cliché of the emotionless android might turn out to be spot on, at least for a while." (132)

Let's turn now to the second worry. To reach the conclusion that some autonomous persons have moral standing but aren't welfare subjects, I've had to take a stand on what kind of welfare goods there are and aren't. Some philosophers will disagree with my position and think that the kind of autonomous agents I'm envisioning are in fact welfare subjects. Set aside the question of who is right. It might not be obvious that anything is at stake. Suppose, for example, that we thought of these agents as benefited by achieving their goals. Would it really make a difference to what we thought of ourselves as owing to them?

Yes. Although Kant (1785 [1998]) construes the virtue of beneficence as reducible to a disposition to share in others' morally permissible ends, our reasons to provide aid to others seem to depend crucially on the extent to which it's specifically their welfare that's at stake, as opposed to some other end of theirs (Arpaly 2018). Thus, if someone is asking for money they need to buy food or medicine, they have some kind of claim on us, but not if they are asking for money to give to their church, even if they prioritize their faith over their basic needs (compare Scanlon 1975; Darwall 2002: 43–45). Given that welfare considerations have a distinctive role to play in grounding positive obligations to aid, it makes an important difference whether some individual is a welfare subject or merely an autonomous agent whose ability to pursue their own projects ought to be respected.

## References

- Alston, W. P. 1988. An internalist externalism. *Synthese* 74(3): 265–283.
- Alston, W. P. 2005. *Beyond "Justification": Dimensions of Epistemic Evaluation*. Ithaca, NY: Cornell University Press.
- Aristotle. 350 BCE [1980]. *The Nicomachean Ethics*, ed. L. Brown, transl. D. Ross. Oxford: Oxford.
- Arneson, R. J. 1999. Human flourishing versus desire satisfaction. *Social Philosophy and Policy* 16(1): 113–142.
- Arpaly, N. 2003. *Unprincipled Virtue: An Inquiry Into Moral Agency*. Oxford: Oxford University Press.
- Arpaly, N. 2018. On benevolence. *Aristotelian Society Supplementary Volume* 92(1): 207–223.
- Ayer, A. J. 1936. *Language, Truth, and Logic*. London: Dover Publications.
- Baars, B. J. 1988. *A Cognitive Theory of Consciousness*. Cambridge: Cambridge University Press.
- Bach, K. 1985. A rationale for reliabilism. *The Monist* 68(2): 246–263.
- Balog, K. 2012. In defense of the phenomenal concept strategy. *Philosophy and Phenomenological Research* 84(1): 1–23.
- Barrett, L. F. 2017. *How Emotions Are Made: The Secret Life of the Brain*. London: Macmillan.
- Basl, J. 2019. *The Death of the Ethic of Life*. Oxford: Oxford University Press.
- Bayne, T., J. Hohwy, and A. M. Owen. 2016. Are there levels of consciousness? *Trends in Cognitive Sciences* 20(6): 405–413.
- Bermúdez, J. P. forthcoming. What is the feeling of effort about? *Australasian Journal of Philosophy*.
- Birch, J. 2020. Global workspace theory and animal consciousness. *Philosophical Topics* 48(1): 21–37.
- Blackburn, S. 1993. *Essays in Quasi-Realism*. Oxford: Oxford University Press.
- Blackburn, S. 1998. *Ruling Passions: A Theory of Practical Reasoning*. New York: Oxford University Press UK.
- Block, N. 1978. Troubles with functionalism. *Minnesota Studies in the Philosophy of Science* 9: 261–325.
- BonJour, L. 1980. Externalist theories of empirical knowledge. *Midwest Studies in Philosophy* 5(1): 53–73.
- Bourget, D. and A. Mendelovici. 2019. Phenomenal intentionality. In *The Stanford Encyclopedia of Philosophy* (Fall 2019 ed.), ed. E. N. Zalta. Metaphysics Research Lab, Stanford University.
- Bradford, G. 2015. *Achievement*. Oxford: Oxford University Press.
- Bradford, G. 2017. Problems for perfectionism. *Utilitas* 29(3): 344–364.
- Bradford, G. 2022. Consciousness and welfare subjectivity. *Noûs* 4(75): 905–921.
- Bradley, B. 2009. *Well-Being and Death*. Oxford: Oxford University Press.
- Bramble, B. 2016. A new defense of hedonism about well-being. *Ergo: An Open Access Journal of Philosophy* 3.
- Brandt, R. B. 1979. *A Theory of the Good and the Right*. Oxford: Oxford University Press.
- Butlin, P., R. Long, E. Elmoznino, Y. Bengio, J. Birch, A. Constant, G. Deane, S. M. Fleming, C. Frith, X. Ji, R. Kanai, C. Klein, G. Lindsay, M. Michel, L. Mudrik, M. A. K. Peters, E. Schwitzgebel, J. Simon, and R. VanRullen. 2023. Consciousness in artificial intelligence: Insights from the science of consciousness. <https://arxiv.org/abs/2308.08708>.
- Carruthers, P. 2000. *Phenomenal Consciousness: A Naturalistic Theory*. Cambridge: Cambridge University Press.
- Carruthers, P. 2019. *Human and Animal Minds: The Consciousness Questions Laid to Rest*. Oxford: Oxford University Press.
- Chalmers, D. J. 1996. *The Conscious Mind: In Search of a Fundamental Theory*. Oxford: Oxford University Press.
- Chalmers, D. J. 2022. *Reality +: Virtual Worlds and the Problems of Philosophy*. London: Allen Lane.

- Chisholm, R. M. 1989. *Theory of Knowledge*, 3rd ed. Englewood Cliffs, NJ: Prentice-Hall.
- Choi, S. and M. Fara. 2021. Dispositions. In *The Stanford Encyclopedia of Philosophy* (Spring 2021 ed.), ed. E. N. Zalta. Metaphysics Research Lab, Stanford University.
- Christman, J. 1991. Autonomy and personal history. *Canadian Journal of Philosophy* 21(1): 1–24.
- Chudnoff, E. 2011. The nature of intuitive justification. *Philosophical Studies* 153(2): 313–333.
- Crisp, R. 2006. *Reasons and the Good*. Oxford: Oxford University Press.
- Cutter, B. 2017. The metaphysical implications of the moral significance of consciousness. *Philosophical Perspectives* 31(1): 103–130.
- Damasio, A. R. 1994. *Descartes' Error: Emotion, Reason, and the Human Brain*. New York, NY: Avon Books.
- Dancy, J. 2004. *Ethics Without Principles*. Oxford: Oxford University Press.
- Darwall, S. 2002. *Welfare and Rational Care*. Princeton, NJ: Princeton University Press.
- Davidson, D. 1963. Actions, reasons, and causes. *Journal of Philosophy* 60(23): 685.
- DeGrazia, D. 2021. An interest-based model of moral status. In *Rethinking Moral Status*, eds. S. Clarke, H. Zohny, and J. Savulescu, 40–56. Oxford: Oxford University Press.
- DeGrazia, D. and J. Millum. 2021. *A Theory of Bioethics*. Cambridge: Cambridge University Press.
- Dehaene, S. 2014. *Consciousness and the Brain: Deciphering How the Brain Codes Our Thoughts*. New York, NY: Viking Press.
- Dorsey, D. 2021. *A Theory of Prudence*. Oxford: Oxford University Press.
- Dworkin, G. 1976. Autonomy and behavior control. *Hastings Center Report* 6(1): 23–28.
- Ekstrom, L. W. 1993. A coherence theory of autonomy. *Philosophy and Phenomenological Research* 53(3): 599–616.
- Engel, M. 1992. Personal and doxastic justification in epistemology. *Philosophical Studies* 67(2): 133–150.
- Enoch, D. 2016. What's wrong with paternalism: Autonomy, belief, and action. *Proceedings of the Aristotelian Society* 116: 21–48.
- Fanciullo, J. 2023. Alienation, engagement, and welfare. *The Philosophical Quarterly*: pqad115.
- Feinberg, J. 1989. *The Moral Limits of the Criminal Law Volume 3: Harm to Self*. Oxford: Oxford University Press.
- Feldman, F. 2004. *Pleasure and the Good Life: Concerning the Nature Varieties and Plausibility of Hedonism*. Oxford: Oxford University Press.
- Feldman, F. 2010. *What is This Thing Called Happiness?* Oxford: Oxford University Press.
- Feldman, R. and E. Conee. 2001. Internalism defended. *American Philosophical Quarterly* 38(1): 1–18.
- Fischer, J. M. and M. Ravizza. 1998. *Responsibility and Control: A Theory of Moral Responsibility*. Cambridge: Cambridge University Press.
- Fletcher, G. 2013. A fresh start for the objective-list theory of well-being. *Utilitas* 25(2): 206–220.
- Frankfurt, H. G. 1971. Freedom of the will and the concept of a person. *Journal of Philosophy* 68(1): 5–20.
- Frankfurt, H. G. 1998. *Necessity, Volition, and Love*. Cambridge: Cambridge University Press.
- Friedman, M. 2003. *Autonomy, Gender, Politics*. Oxford: Oxford University Press.
- Gibbard, A. 1990. *Wise Choices, Apt Feelings: A Theory of Normative Judgment*. Cambridge, MA: Harvard University Press.
- Gibbard, A. 2003. *Thinking How to Live*. Cambridge, MA: Harvard University Press.

- Goldman, A. 1987. Strong and weak justification. *Philosophical Perspectives* 2: 51–69.
- Goldman, A. I. 1979. What is justified belief? In *Justification and Knowledge: New Studies in Epistemology*, ed. G. Pappas, 1–25. Boston, MA: D. Reidel.
- Goldman, A. I. 1986. *Epistemology and Cognition*. Cambridge, MA: Harvard University Press.
- Goodpaster, K. E. 1978. On being morally considerable. *Journal of Philosophy* 75(6): 308–325.
- Graziano, M. S. A. 2019. *Rethinking Consciousness: A Scientific Theory of Subjective Experience*. New York, NY: W.W. Norton & Company.
- Gregory, A. 2021. *Desire as Belief: A Study of Desire, Motivation, and Rationality*. Oxford: Oxford University Press.
- Griffin, J. 1986. *Well-Being: Its Meaning, Measurement, and Moral Importance*. Oxford: Oxford University Press.
- Griffiths, P. E. 1997. *What Emotions Really Are: The Problem of Psychological Categories*. Chicago, IL: University of Chicago Press.
- Haji, I. 1998. *Moral Appraisability: Puzzles, Proposals, and Perplexities*. Oxford: Oxford University Press.
- Haybron, D. M. 2008. *The Pursuit of Unhappiness: The Elusive Psychology of Well-Being*. Oxford: Oxford University Press.
- Hayward, M. K. 2019. Immoral realism. *Philosophical Studies* 176(4): 897–914.
- Heathwood, C. 2005. The problem of defective desires. *Australasian Journal of Philosophy* 83(4): 487–504.
- Heathwood, C. 2019. Which desires are relevant to well-being? *Noûs* 53(3): 664–688.
- Hooker, B. 2015. The elements of well-being. *Journal of Practical Ethics* 3(1): 15–35.
- Horgan, T. and J. Tienson. 2002. The intentionality of phenomenology and the phenomenology of intentionality. In *Philosophy of Mind: Classical and Contemporary Readings*, ed. D. J. Chalmers, 520–533. Oxford: Oxford University Press.
- Huemer, M. 2001. *Skepticism and the Veil of Perception*. Lanham, MD: Rowman & Littlefield.
- Huemer, M. 2005. *Ethical Intuitionism*. Basingstoke: Palgrave Macmillan.
- Huemer, M. 2006. Phenomenal conservatism and the internalist intuition. *American Philosophical Quarterly* 43(2): 147–158.
- Huemer, M. 2007. Compassionate phenomenal conservatism. *Philosophy and Phenomenological Research* 74(1): 30–55.
- Huemer, M. 2013. Phenomenal conservatism über alles. In *Seemings and Justification: New Essays on Dogmatism and Phenomenal Conservatism*, ed. C. Tucker, 328. Oxford: Oxford University Press.
- Hume, D. 1739–40 [1978]. *A Treatise of Human Nature*, 2nd edn., eds. Selby-Bigge and Nidditch. Oxford: Oxford University Press.
- Hurka, T. 1993. *Perfectionism*. Oxford: Oxford University Press.
- Hurka, T. 2011. *The Best Things in Life: A Guide to What Really Matters*. Oxford: Oxford University Press.
- James, W. 1884. What is an emotion? *Mind* 9(34): 188–205.
- Kagan, S. 2009. Well-being as enjoying the good. *Philosophical Perspectives* 23(1): 253–272.
- Kagan, S. 2014. An introduction to ill-being. *Oxford Studies in Normative Ethics* 4: 261–88.
- Kagan, S. 2019. *How to Count Animals, More or Less*. Oxford: Oxford University Press.
- Kant, I. 1785 [1998]. *Groundwork of the Metaphysics of Morals*, transl. M. Gregor. Cambridge: Cambridge University Press.
- Kauppinen, A. 2022. Moral Sentimentalism. In *The Stanford Encyclopedia of Philosophy* (Spring 2022 ed.), ed. E. N. Zalta. Metaphysics Research Lab, Stanford University.
- Killmister, S. 2013. Autonomy and false beliefs. *Philosophical Studies* 164(2): 513–531.

- Kirk, R. 1974. Sentience and behaviour. *Mind* 83(329): 43–60.
- Korsgaard, C. M. 2018. *Fellow Creatures: Our Obligations to the Other Animals*. Oxford: Oxford University Press.
- Kraut, R. 2007. *What is Good and Why: The Ethics of Well-Being*. Cambridge, MA: Harvard University Press.
- Kriegel, U. 2011. *The Sources of Intentionality*. Oxford: Oxford University Press.
- Kripke, S. 1980. *Naming and Necessity*. Cambridge, MA: Harvard University Press.
- Lau, H. 2022. *In Consciousness We Trust: The Cognitive Neuroscience of Subjective Experience*. Oxford: Oxford University Press.
- Lee, A. Y. 2018. Is consciousness intrinsically valuable? *Philosophical Studies* 175(1): 1–17.
- Lee, A. Y. 2022. Speciesism and sentientism. *Journal of Consciousness Studies* 29(3-4): 205–228.
- Lee, A. Y. 2023. Degrees of consciousness. *Noûs* 57(3): 553–575.
- Lee, A. Y. forthcoming. Consciousness makes things matter. *Philosophers' Imprint*.
- Levin, J. 2008. Taking type-B materialism seriously. *Mind and Language* 23(4): 402–425.
- Levy, N. 2014a. *Consciousness and Moral Responsibility*. Oxford: Oxford University Press.
- Levy, N. 2014b. The value of consciousness. *Journal of Consciousness Studies* 21(1-2): 127–138.
- Lewis, D. K. 1988. Desire as belief. *Mind* 97(418): 323–32.
- Lewis, D. K. 1996. Desire as belief II. *Mind* 105(418): 303–13.
- Lippert-Rasmussen, K. 2007. Why killing some people is more seriously wrong than killing others. *Ethics* 117(4): 716–738.
- Littlejohn, C. 2009. The externalist's demon. *Canadian Journal of Philosophy* 39(3): 399–434.
- Loar, B. 1990. Phenomenal states. *Philosophical Perspectives* 4: 81–108.
- Loar, B. 2003. Phenomenal intentionality as the basis of mental content. In *Reflections and Replies: Essays on the Philosophy of Tyler Burge*, eds. M. Hahn and B. Ramberg, 229–258. Cambridge, MA: MIT Press.
- Lycan, W. G. 1996. *Consciousness and Experience*. Cambridge, Mass.: MIT Press.
- McDowell, J. 1978. Are moral requirements hypothetical imperatives? *Aristotelian Society Supplementary Volume* 52(1): 13–29.
- McGrath, S. 2004. Moral knowledge by perception. *Philosophical Perspectives* 18(1): 209–228.
- McLaughlin, B. P. 2018. Could an android be sentient? In *Blockheads! Essays on Ned Block's Philosophy of Mind and Consciousness*, eds. A. Pautz and D. Stoljar, 335–373. Cambridge, MA: MIT Press.
- McMahan, J. 2002. *The Ethics of Killing: Problems at the Margins of Life*. Oxford: Oxford University Press.
- McNaughton, D. 1988. *Moral Vision: An Introduction to Ethics*. New York, NY: Blackwell.
- Mele, A. R. 1995. *Autonomous Agents: From Self-Control to Autonomy*. Oxford: Oxford University Press.
- Moss, S. 2016. *Probabilistic Knowledge*. Oxford: Oxford University Press.
- Nagel, T. 1970. *The Possibility of Altruism*. Oxford: Clarendon Press.
- Neely, W. 1974. Freedom and desire. *Philosophical Review* 83(1): 32–54.
- Nussbaum, M. C. 2001. *Upheavals of Thought: The Intelligence of Emotions*. Cambridge: Cambridge University Press.
- Nussbaum, M. C. 2022. *Justice for Animals: Our Collective Responsibility*. New York, NY: Simon & Schuster.

- Pallies, D. 2022. Attraction, aversion, and asymmetrical desires. *Ethics* 132(3): 598–620.
- Papineau, D. 2002. *Thinking About Consciousness*. Oxford: Oxford University Press.
- Parfit, D. 1984. *Reasons and Persons*. Oxford: Oxford University Press.
- Prinz, J. J. 2004. *Gut Reactions: A Perceptual Theory of the Emotions*. Oxford: Oxford University Press.
- Pryor, J. 2000. The skeptic and the dogmatist. *Noûs* 34(4): 517–549.
- Pryor, J. 2004. What’s wrong with Moore’s argument? *Philosophical Issues* 14(1): 349–378.
- Pugh, J. 2020. *Autonomy, Rationality, and Contemporary Bioethics*. Oxford, UK: Oxford University Press.
- Quinn, W. 1984. Abortion: Identity and loss. *Philosophy and Public Affairs* 13(1): 24–54.
- Quinn, W. 1993. Putting rationality in its place. In *Morality and Action*, 228–255. Cambridge: Cambridge University Press.
- Raibley, J. 2010. Well-being and the priority of values. *Social Theory and Practice* 36(4): 593–620.
- Railton, P. 1986. Facts and values. *Philosophical Topics* 14(2): 5–31.
- Rawls, J. 1971. *A Theory of Justice*. Cambridge, MA: Harvard University Press.
- Roberts, R. C. 2003. *Emotions: An Essay in Aid of Moral Psychology*. Cambridge: Cambridge University Press.
- Rosati, C. S. 1996. Internalism and the good for a person. *Ethics* 106(2): 297–326.
- Rosenthal, D. M. 2005. *Consciousness and Mind*. Oxford: Oxford University Press.
- Russell, J. 1980. 12). A circumplex model of affect. *Journal of Personality and Social Psychology* 39: 1161–1178.
- Russell, J. A. 2003. Core affect and the psychological construction of emotion. *Psychological Review* 110(1): 145–172.
- Savulescu, J. 2007. Rational desires and the limitation of life-sustaining treatment. *Bioethics* 8(3): 191–222.
- Scanlon, T. M. 1975. Preference and urgency. *Journal of Philosophy* 72(19): 655–669.
- Searle, J. R. 1990. Consciousness, explanatory inversion and cognitive science. *Behavioral and Brain Sciences* 13(1): 585–642.
- Searle, J. R. 1992. *The Rediscovery of the Mind*. Cambridge, MA: MIT Press.
- Shepherd, J. 2024. Sentience, vulcans, and zombies: The value of phenomenal consciousness. *AI and Society*: 1–11.
- Sher, G. 1997. *Beyond Neutrality: Perfectionism and Politics*. Cambridge: Cambridge University Press.
- Shiffrin, S. V. 2000. Paternalism, unconscionability doctrine, and accommodation. *Philosophy and Public Affairs* 29: 205–250.
- Shoemaker, D. W. 2003. Caring, identification, and agency. *Ethics* 114(1): 88–118.
- Singer, P. 1977. *Animal Liberation: A New Ethics for Our Treatment of Animals*. New York, NY: Avon Books.
- Singer, P. 1993. *Practical Ethics*, 2nd. ed. Cambridge: Cambridge University Press.
- Sinhababu, N. 2009. The Humean theory of motivation reformulated and defended. *Philosophical Review* 118(4): 465–500.
- Sinhababu, N. 2017. *Humean Nature: How Desire Explains Action, Thought, and Feeling*. Oxford: Oxford University Press.
- Sinnott-Armstrong, W. and V. Conitzer. 2021. How much moral status could artificial intelligence ever achieve? In *Rethinking Moral Status*, eds. S. Clarke, H. Zohn, and J. Savulescu, 269–281. Oxford: Oxford University Press.
- Sliwa, P. 2012. In defense of moral testimony. *Philosophical Studies* 158(2): 175–195.
- Smith, M. 1994. *The Moral Problem*. Oxford: Blackwell.
- Smithies, D. 2012. The mental lives of zombies. *Philosophical Perspectives* 26(1): 343–372.

- Smithies, D. 2019. *The Epistemic Role of Consciousness*. Oxford: Oxford University Press.
- Smithies, D. forthcoming. Affective consciousness and moral status. *Oxford Studies in Philosophy of Mind*.
- Smithies, D. and J. Weiss. 2019. Affective experience, desire, and reasons for action. *Analytic Philosophy* 60(1): 27–54.
- Solomon, R. C. 1973. Emotions and choice. In *Not Passion's Slave*, 3–24. Oxford: Oxford University Press.
- Solomon, R. C. 1976. *The Passions*. Notre Dame, IN: University of Notre Dame Press.
- Solomon, R. C. 1988. On emotions as judgments. *American Philosophical Quarterly* 25(2): 183–191.
- Solomon, R. C. 2003. Emotions, thoughts and feelings: What is a 'cognitive theory' of the emotions and does it neglect affectivity? *Royal Institute of Philosophy Supplement* 52: 1–18.
- Sumner, L. W. 1996. *Welfare, Happiness, and Ethics*. Oxford: Oxford University Press.
- Tiberius, V. 2018. *Well-Being as Value Fulfillment: How We Can Help Each Other to Live Well*. Oxford: Oxford University Press.
- Vadas, M. 1984. Affective and nonaffective desire. *Philosophy and Phenomenological Research* 45(2): 273–80.
- Velleman, J. D. 1998. Is motivation internal to value? In *Preferences*, eds. C. Fehige and U. Wessels, 88–102. New York, NY: de Gruyter.
- Wall, S. 1998. *Liberalism, Perfectionism and Restraint*. Cambridge: Cambridge University Press.
- Watson, G. 1975. Free agency. *Journal of Philosophy* 72(8): 205–20.
- Watson, G. 1987. Free action and free will. *Mind* 96(382): 154–72.
- Williamson, T. 2000. *Knowledge and its Limits*. Oxford: Oxford University Press.
- Wolf, S. R. 1990. *Freedom Within Reason*. Oxford: Oxford University Press.
- Yelle, B. 2014. Alienation, deprivation, and the well-being of persons. *Utilitas* 26(4): 367–384.