

Affective Consciousness and Moral Status

Declan Smithies

Forthcoming in *Oxford Studies in Philosophy of Mind*, Volume 5.

Abstract. Which beings have moral status? This paper argues that moral status requires some capacity for affective consciousness. David Chalmers rejects this view on the grounds that it denies moral status to Vulcans – namely, conscious creatures with no capacity for affective consciousness. On his more inclusive view, all conscious beings have moral status. Although we agree that consciousness is required for moral status, we disagree about how to explain this. I argue that we cannot explain why unconscious zombies lack moral status without excluding conscious Vulcans too.

1. Varieties of Sentientism

Which beings have moral status? Although this was not always true in the past, it is now widely acknowledged that all humans have moral status. Moreover, there is increasing recognition that many non-human animals have moral status too. Even so, there is little temptation to extend moral status to all living organisms, including plants, fungi, and bacteria, since these simple life forms don't have the right kinds of minds. This raises the question: what kind of mind do you need for moral status?

According to *sentientism*, moral status requires the capacity for consciousness.¹ Consciousness is understood here in the phenomenal sense, according to which consciousness is experience. As defined, however, sentientism leaves open what kind of consciousness you need for moral status.

David Chalmers (2022) argues for an *inclusive* version of sentientism, according to which the capacity for consciousness is not only necessary but also sufficient for moral status:

Inclusive Sentientism: A being has moral status if and only if it has the capacity for consciousness.

On this view, all and only conscious beings have moral status. Any conscious being has moral status, no matter how simple it is, although its degree of moral status may vary depending on its psychological complexity. Meanwhile, no unconscious being has any degree of moral status no matter how complex its life might be in other ways. Not even *zombies* – that is, unconscious beings that function just like conscious beings – can have moral status.

¹ Proponents of sentientism about either moral status or welfare include Singer (1975), Rosati (2009), Shepherd (2018), van der Deijl (2020), Lin (2021), Siewert (2021), and Chalmers (2022).

While I agree with Chalmers that consciousness is necessary for moral status, I deny that it is sufficient. My own view is that moral status depends on the capacity for consciousness that is *affectively valenced* in the sense that it involves feeling positively or negatively about things:

Affective Sentientism: A being has moral status if and only if it has the capacity for affectively valenced consciousness.

This resembles, and may even be coextensive with, Peter Singer's (1975) view that moral status depends on the capacity for consciousness that is *hedonically valenced* in the sense that it involves feeling pleasure or displeasure:

Hedonic Sentientism: A being has moral status if and only if it has the capacity for hedonically valenced consciousness.

Whether these views are distinct depends on an issue that I cannot adjudicate here: namely, whether hedonic consciousness is the only species of affective consciousness. I'll remain agnostic about whether the affective valence of feelings of desire and emotion is reducible to their hedonic valence.² If so, my view collapses into Singer's. If not, it is more inclusive. Either way, it is more restrictive than Chalmers' view, since not all consciousness is affectively valenced. As such, it is vulnerable to the same objection that Chalmers raises against Singer's view: namely, that it wrongly excludes conscious beings who are incapable of affectively valenced consciousness of any kind.

My preliminary goal in this paper is to answer this objection. In the course of answering it, however, I hope to achieve a more ambitious goal – namely, explaining why sentientism is true in the first place. With this explanation in place, we will see that it forces us to abandon Chalmers' inclusive sentientism and to endorse affective sentientism instead.

2. The Moral Status of Vulcans

A *Vulcan* is defined by Chalmers as “a conscious creature who experiences no happiness, suffering, pleasure, pain, or any other positive or negative affective states” (2022: 343). There is no incoherence

² Smithies (MS) defends a hedonic theory of desire and applies it to debates about welfare and moral status, but this commitment plays no role in the arguments of this paper.

in this definition given the plausible assumption that not all experience is affectively valenced. Vulcans may lack any capacity to experience pleasure, desire, or emotion, but they can nevertheless experience many sensations, perceptions, and conscious thoughts. Our question is whether Vulcans have any moral status: inclusive sentientism implies that they do, whereas hedonic or affective sentientism implies otherwise. Hence, Vulcans provide a useful test case for deciding between competing versions of sentientism.

When we imagine Vulcans, we must be careful to avoid conflating important distinctions. One such distinction concerns *sensory* versus *attitudinal* pleasure. Sensory pleasure may be regarded as one species of attitudinal pleasure. According to higher-order theories, you feel sensory pleasure when you experience a sensation that you feel pleased about for its own sake. On a first-order theory, in contrast, you feel sensory pleasure when you experience a sensation that represents something else that you feel pleased about for its own sake. Either way, sensory pleasure may be reduced to attitudinal pleasure directed towards sensations themselves or whatever those sensations represent. We can give a similar account of sensory *displeasure*, including the unpleasant feeling of sensory pain.³

Crucially, however, sensory pleasure is not the *only* species of attitudinal pleasure. You can feel pleased about something without thereby feeling sensory pleasure. Fred Feldman (1988) gives the example of an injured motorcyclist who feels pleased that he survived his accident, although he has taken a powerful anesthetic that prevents him from feeling any sensations at all. When we imagine Vulcans, it is not enough to imagine creatures who are incapable of sensory pleasure and pain. We must imagine creatures who are incapable of feeling pleased or displeased about anything at all. Otherwise, we artificially inflate the plausibility of the claim that they have moral status.

Another distinction concerns *pleasure* versus *affect*. Pleasure is one species of affective experience, but it is perhaps not the only one. Feelings of desire and emotion represent their objects positively or negatively just like feelings of pleasure and displeasure. And yet it is by no means obvious that feelings of desire and emotion can represent their objects positively or negatively only by feeling pleasant or unpleasant. Your desires typically feel pleasant when you believe they are satisfied but not when you believe they are frustrated. Feelings of apparently frustrated desire, such as hunger or unrequited love, can feel unpleasant despite representing their objects in a positively valenced way.

³ Feldman (1988) and Heathwood (2007) defend higher-order theories of pleasure, whereas Jacobson (2019) defends a first-order theory of why pain feels unpleasant: namely, it represents a bodily condition towards which we feel aversive desire. See Lin (2020) for the claim that feelings of pleasure and displeasure have a distinctive attitudinal phenomenology.

When we imagine Vulcans, it is not enough to imagine creatures who are incapable of feeling pleasure or displeasure. We must imagine creatures who are incapable of affective experience of any kind, including feelings of desire or emotion.

A third distinction concerns having a *diminished* capacity for affective experience versus having no such capacity at all. Chalmers is careful to distinguish his philosophical Vulcans from the Star Trek characters that inspired them: Spock may be less emotional than his human counterparts, but he has some minimal capacity for affective experience, including pleasure and pain. In contrast, Vulcans are defined as having no capacity for affective experience. In this way, they are unlike people suffering from real-life affective disorders, such as depression or anhedonia, which diminish your capacity for affective experience without eliminating the capacity altogether.

With these clarifications in mind, let's now consider whether Vulcans have moral status. Chalmers' strategy is to pump intuitions about moral status by using variations on the classic trolley problem. First, he argues that *zombies* lack moral status:

You're at the wheel of a runaway trolley. If you do nothing, it will kill a single conscious human, who is on the tracks in front of you. If you switch tracks, it will kill five nonconscious zombies. What should you do? (2022: 340)

Chalmers' intuition, and mine too, is that you should kill the five zombies to save the one conscious human. He is even tempted to say that you should destroy a whole planet of humanoid zombies in order to save one conscious chicken, although he admits to wavering on this point. When we substitute Vulcans for zombies, however, his intuitions become much clearer:

Would it be morally acceptable to kill a planet of philosophical Vulcans to save one human with ordinary affective consciousness? I think the answer is obviously no. More simply, suppose you're faced with a situation in which you can kill a Vulcan in order to save an hour on the way to work. It would obviously be morally wrong to kill the Vulcan. In fact, it would be monstrous. (2022: 344)

What makes all the difference between these cases, according to Chalmers, is that Vulcans have moral status, whereas zombies have none.

While I feel the force of these moral intuitions, I doubt they provide any reliable evidence for the moral status of Vulcans. To understand why, it is important to understand how Chalmers connects the concepts of *moral status* and *welfare*: “A being has moral status when . . . it’s a being whose welfare we need to take into account in our moral deliberations” (2022: 339). The plausible suggestion here is that being a welfare subject is what gives you moral status in the first place. Following Chalmers, then, let’s assume that all and only welfare subjects have moral status.⁴

My complaint is that Chalmers’ intuitions about the trolley problem depend on conflating another important distinction between *moral status* and *moral significance*. Not everything with moral significance has moral status: only welfare subjects can have moral status, but other things can have moral significance when they impact the welfare of beings with moral status. Crucially, we needn’t suppose that Vulcans have moral status in order to explain why it would be monstrous to destroy them merely to save time on your morning commute. After all, it would be monstrous to destroy a beautiful rainforest or valuable artwork just for your own convenience. And yet rainforests and artworks don’t have moral status, since they are not welfare subjects: they don’t satisfy the requirement for consciousness imposed by even the most inclusive version of sentientism.

Indeed, this distinction can explain why Chalmers finds his own intuitions wavering when he considers extreme versions of the zombie trolley problem. According to inclusive sentientism, conscious chickens have welfare status, whereas unconscious zombies have none. And yet Chalmers is reluctant to embrace the consequence that we should destroy a whole planet of humanoid zombies to save just one conscious chicken. We can resolve any apparent tension here by recognizing that zombies – like rainforests and artworks – may have some moral significance without any moral status. If we conscious beings value these things highly enough, then it may even be morally permissible to countenance some loss of conscious life in order to preserve them. And if this can be true for zombies, then all the more so for Vulcans. Vulcans may have some moral significance if they are valued by welfare subjects, but it doesn’t follow that they are welfare subjects themselves and thereby endowed with moral status. Indeed, I’ll argue that Vulcans have no moral status at all.

3. Against Inclusion

Rather than simply trading intuitions about the moral status of Vulcans, my goal is to move the debate forward by confronting inclusive sentientism with an explanatory question. What could possibly

⁴ Anyone who rejects this assumption can simply recast the whole debate in terms of welfare rather than moral status.

explain why Vulcans, unlike zombies, are welfare subjects that have moral status? There are two things we need to explain here:

- (1) **Inclusion:** All conscious beings are welfare subjects.
- (2) **Sentientism:** No unconscious beings are welfare subjects.

What I'll argue is that Chalmers cannot explain both claims without equivocating between distinct concepts of value. Once we have this distinction in clear focus, we'll see that we cannot plausibly explain sentientism without abandoning inclusion.

The materials for Chalmers's explanation of inclusive sentientism are contained within the following passage:

Consciousness is the ground of all value. Whenever anything is good or bad for someone, it's because of their consciousness. Consciousness has value, what a conscious creature values has value, and relations between conscious creatures have value. If a creature has no capacity for consciousness, nothing can be good or bad for it from its own perspective. And it's natural to conclude that if nothing can be good or bad for a creature, then the creature has no moral status. (2022: 341-2)

This passage is compressed enough that it needs some fleshing out, but we can reconstruct the proposed explanation of inclusive sentientism by extracting two central claims:

- (1) **Intrinsic Value:** All consciousness has intrinsic value.
- (2) **Grounding:** All other value is grounded in the intrinsic value of consciousness.

The first claim is designed to explain inclusion. The idea is that all consciousness has value – not just instrumentally but intrinsically – for any conscious being. On this view, consciousness is an intrinsic welfare good that always makes your life better for you to some extent, although this may of course be outweighed by other factors that negatively impact your welfare. Consequently, all conscious beings are welfare subjects, since they can be harmed by depriving them of an intrinsic welfare good – namely, consciousness. Meanwhile, the second claim is designed to explain sentientism. Since all value –

including welfare value – is grounded in consciousness, no unconscious beings are welfare subjects. Zombies may have value for conscious beings like us, but their lives have no value *for them*.

As elegant as it is, this explanation trades on equivocation. We cannot plausibly maintain both claims without shifting between distinct concepts of value: namely, *intrinsic value* and *welfare value*. Admittedly, some welfare values are intrinsic values, but not all intrinsic values are welfare values. To illustrate the distinction: like many epistemologists, I accept that knowledge has intrinsic epistemic value in addition to its instrumental prudential value in promoting the satisfaction of your desires, but I'm inclined to deny that knowledge of random trivia always makes a positive contribution to welfare. Is this a coherent combination? There need be no incoherence so long as we recognize that knowledge can have intrinsic *epistemic* value without having any intrinsic *welfare* value.

With this distinction in mind, let's consider Grounding: is it true that all value is grounded in consciousness? This seems dubious without any further restriction. Following Moore, we can contrast two worlds devoid of consciousness, one full of beauty and the other replete with filth:

Let us imagine one world exceedingly beautiful. Imagine it as beautiful as you can; put into it whatever in this world you most admire – mountains, rivers, the sea; and sunsets, stars and moon . . . And then imagine the ugliest world you can possibly conceive. Imagine it simply one heap of filth . . . The only thing we are not entitled to imagine is that any human being ever has or ever, by any possibility, *can*, live in either, can ever see and enjoy the beauty of the one or hate the foulness of the other. (1903: 84)

As Moore observed, the first world seems better than the second insofar as it is more beautiful. After all, beauty is a kind of value – namely, *aesthetic* value. Moreover, the beauty of the first world must be intrinsically valuable, rather than merely extrinsically valuable for conscious beings, since it contains no conscious beings to admire its beauty. At the very least, it remains deeply obscure how the value of consciousness could ground the value of a beautiful world that is devoid of consciousness by necessity. We cannot simply assume without further argument that consciousness is the only intrinsic value from which all other values are derived.

This objection might seem uncharitable, since Chalmers is primarily concerned with welfare value, and its connection with moral value, rather than aesthetic value.⁵ A beautiful world devoid of consciousness might be good in itself, but it is not good *for* anyone in the sense that is relevant to welfare and morality. So, perhaps what Chalmers intends is a restricted version of Grounding, according to which all moral and welfare value is grounded in consciousness. If so, then we are in full agreement. In a world without consciousness, nothing has value in this restricted sense: nothing is good or bad *for* anyone, so nothing is morally good or bad.

Once we impose this restriction, however, we undercut the plausibility of Intrinsic Value. I have no dispute with the claim that consciousness has intrinsic value in something like the way that unconscious life does. Perhaps it has aesthetic value or perhaps intrinsic value of some other kind. Other things being equal, a conscious world seems better than a zombie world just as a world teeming with unconscious life seems better than a world comprising only rocks and dust. And yet we can accept that life has intrinsic value without supposing it is a welfare good for any living organism. Similarly, we can accept that consciousness has intrinsic value without supposing it is a welfare good for any conscious being.

To further illustrate the point, we can adopt an example from Andrew Lee:

Consider again two worlds that are empty save for a single creature inhabiting the universe. In the first world, the creature has a painful experience, and the character of its experience is fully exhausted by this pain phenomenology . . . In the second world, the creature is not conscious at all. (2018: 664)

Is the painful world better or worse than the zombie world with no consciousness at all? My own reaction is one of ambivalence. On the one hand, how awesome it is that consciousness emerged in this world! The painful world contains a source of intrinsic value that is entirely missing from a zombie world. On the other hand, how unfortunate for the creature in question! It would be better for this creature not to be conscious or not to exist at all. There need be no conflict between these evaluations, since they use distinct and perhaps incommensurable evaluative concepts. Consciousness can have

⁵ Chalmers (2022: 315) says he is mainly concerned with personal or prudential value, which I'm calling "welfare value". Although he connects this with moral value, he mentions aesthetic value only to set it aside, and doesn't refer back to it when he introduces his proposal that "all value arises, one way or another, from consciousness" (2022: 329).

intrinsic value without being a welfare value for all conscious beings. We cannot draw conclusions about welfare or moral status solely from premises about the intrinsic value of consciousness.⁶

Chalmers might insist that consciousness always has positive welfare value, although it can sometimes be outweighed by the negative welfare value of specific ways of being conscious, such as being in pain. However, we can make trouble for this view by varying the intensity of displeasure in the painful world (Lee 2018: 664-67). If consciousness has non-trivial welfare value, then this can in principle outweigh the disvalue of pain when its unpleasantness is mild enough. And yet it seems implausible that being conscious is good overall for a creature who experiences nothing but mildly unpleasant pain. Arguably, this is because the determinable property of being conscious – as distinct from any more determinate way of being conscious – has no positive welfare value at all.

My main objection, however, is more straightforward. We can imagine conscious beings who are not welfare subjects at all. Consider *the Blob*, an alien life form superficially resembling a blobfish that drifts around on the bottom of the ocean. Blobs are not agents who are capable of self-movement. Instead, they are passive observers who experience an impoverished stream of conscious sensations, perceptions, and thoughts about themselves and their environment. Crucially, however, they are incapable of affective experience of any kind. When Blobs see something red, they might think, “That is red,” or even, “That looks red to me,” but they feel wholly indifferent to the presence or absence of red things. Indeed, we may even suppose that they feel indifferent to their own continued existence. We needn’t imagine that Blobs evaluate these things and find them to be worthless. Rather, Blobs are constitutionally incapable of feeling positively or negatively about anything at all.

I find it highly implausible to suppose that Blobs are welfare subjects. After all, welfare subjects are the kinds of beings who can possess welfare values. But what could make the lives of Blobs go better or worse for them? We cannot benefit them by moving red objects into their vicinity, or harm them by taking them away, since they are indifferent to the presence or absence of red things. And we cannot harm them by destroying them, since they are indifferent to their own continued existence. We might talk about what’s good for Blobs in the same way that we talk about what’s good for plants

⁶ Lee himself doesn’t use this example to distinguish between intrinsic value and welfare value. Indeed, he seems to conflate them in much the same way as Chalmers. While Chalmers assumes that consciousness has welfare value because it has intrinsic value, Lee seems to argue in the opposite direction: that consciousness has no intrinsic value because it has no welfare value. In my view, both inferences are fallacious.

and machines, but not all relational goods are welfare goods (Rosati 2009). Like plants and machines, Blobs can be damaged or destroyed, but they cannot be harmed or wronged.

The obvious rejoinder is that we can harm Blobs by depriving them of consciousness. To my mind, however, this reply loses much of its plausibility once we recognize that not all intrinsic values are welfare values. We can acknowledge that knowledge has intrinsic epistemic value, for example, without supposing that we harm Blobs by depriving them of the capacity to know about the location of red things. Similarly, we can recognize that consciousness has intrinsic value without supposing that we harm Blobs by depriving them of the capacity to experience red things.

Some might protest: isn't better for you to be a conscious Blob than an unconscious zombie? As Charles Siewert (2021) observes, many of us would decline a "radical pheno-ectomy" that surgically removes all our conscious experience while leaving our behavioral capacities intact. We prefer to remain conscious, rather than zombified, because we value conscious experience for its own sake and not just because of its behavioral consequences. Even so, it is doubtful that we intrinsically value the determinable property of consciousness itself, rather than some of its more specific determinates. Being zombified would deprive me of ways of being conscious that I value for their own sake, including the pleasures of eating steak or thinking about philosophy. Without conscious perception and conscious thought, these things would be impossible. Speaking for myself, however, I can see no welfare value in conscious thought or perception *per se*.

Suppose an accident will render me permanently comatose: I have no preference for a future in which I experience only a monotone or an earworm over one in which I remain unconscious. Indeed, I cannot see how this impoverished stream of consciousness would make my life any better for me than the alternative. Perhaps others have different preferences: a conscious future may be better for them, since it would satisfy their current desires. It doesn't follow, however, that consciousness has any welfare value for creatures who are completely indifferent towards it. Intuitively, for instance, there is nothing good or bad for the Blob about being conscious.

This is already enough to cast significant doubt on inclusion. Even so, I don't wish to rely on solely on intuitive verdicts about cases. My main argument against inclusion is an explanatory one: we must reject inclusion in order to explain sentientism.

4. Explaining Sentientism

Although Chalmers and I disagree about inclusion, we agree about sentientism. Nevertheless, we disagree about how to explain this. Indeed, I'll argue that Chalmers presupposes what needs to be explained. The challenge is to explain why all welfare subjects must be conscious.⁷

Suppose we assume what I've been questioning so far – namely, that consciousness is an intrinsic welfare good. Even so, it is surely not the only one. At a minimum, these welfare goods include specific ways of being conscious, such as feeling pleasure. We can define *experientialism* as the view that all such welfare goods are more or less specific ways of being conscious. One prominent version is *hedonism*, according to which the only welfare good is feeling pleasure, but there are also more liberal versions of experientialism that count some non-hedonic experiences as welfare goods, such as novelty, self-understanding, or compassion (van der Deijl 2019). To explain inclusion, of course, we need a maximally liberal version of experientialism that includes any way of being conscious as a welfare good. Otherwise, we cannot explain why all conscious beings are welfare subjects. Nevertheless, the challenge remains to explain sentientism.

Willem van der Deijl (2020) argues that experientialism is the only adequate explanation of sentientism. On this view, all welfare subjects must be conscious because only conscious experience can make any intrinsic contribution to welfare. The main problem is that experientialism has the implausible consequence that welfare supervenes on experience alone. As Robert Nozick's (1974) experience machine illustrates, experiential duplicates can diverge in welfare depending on whether their experiences correspond to objective reality. Like many philosophers, I'm inclined to regard this as a compelling – perhaps even decisive – objection to experientialism. If experientialism is the only way to explain sentientism, then so much the worse for sentientism.

Proponents of sentientism who reject experientialism confront a more difficult explanatory challenge. If not all welfare goods are ways of being conscious, then why should consciousness be a prerequisite for possessing any welfare goods at all? Suppose there are some non-experiential welfare goods, such as friendship or achievement, in addition to experiential welfare goods, such as pleasure. By definition, zombies cannot feel pleasure, but why can't they have friends or achievements or any

⁷ This challenge is discussed by van der Deijl (2020), Lin (2021), Bradford (2023), and Kriegel (forthcoming).

other non-experiential welfare goods? What we need is some unifying explanation that makes it intelligible why none of these welfare values can be possessed by any unconscious being.⁸

To explain sentientism, Chalmers invokes his restricted version of Grounding: all welfare values are grounded in consciousness in such a way that only conscious beings can possess them. And yet this explanation only pushes the crucial question back one stage further. What is the principled reason to suppose that any welfare value must be grounded in consciousness in such a way that only a conscious being can possess it? To maintain this without any further argument is tantamount to endorsing sentientism without explaining it.

Perhaps some proponents of sentientism will reject any demand for explanation. Quietism is a risky strategy, however, since sentientism may be regarded with suspicion in the absence of any principled explanation. For instance, Gwen Bradford (2023) rejects sentientism on the grounds that it has no principled and plausible explanation, while Uriah Kriegel (forthcoming) comes dangerously close to the same conclusion. This paper aims to develop an explanation that neither Bradford nor Kriegel considers in any significant detail.⁹ As I'll explain, however, this explanation of sentientism undercuts the case for inclusion.

My explanatory strategy bears comparison with one recently proposed by Luke Roelofs (2023). He argues for sentientism on the grounds that only consciousness can rationally motivate action by providing its subject with justifying reasons for action. On his view, which he calls *motivational sentientism*, a being has welfare and moral status only if it has the right kind of consciousness to rationally motivate action. At the same time, he officially remains agnostic about which kinds of consciousness are required to play this normative role in rationally motivating action. This affords him some degree of flexibility in accommodating the moral status of Vulcans. While my position has much in common with his, I'll take a harder line against the moral status of Vulcans by arguing that only *affective consciousness* can play the requisite normative role in providing justifying reasons for action.

On this view, being a welfare subject with moral status requires the capacity for affective consciousness. Ultimately, we cannot explain why zombies are not welfare subjects without excluding Vulcans too, since they both violate this condition in slightly different ways. Zombies are not welfare

⁸ To my mind, it is not enough to argue on a case-by-case basis, as Lin (2021: 880) does, that consciousness is required for each basic welfare good considered by itself. This strategy is too disjunctive to provide any unifying explanation.

⁹ Bradford invokes the functionalist theory of desire that I argue against, whereas Kriegel denies that desiring (or “caring”) is an affectively valenced attitude because he treats hedonic valence as the only kind of affective valence while also rejecting hedonic theories of desire.

subjects because they have no capacity for consciousness of any kind. Of course, Vulcans are conscious creatures. Even so, they are not welfare subjects because they have no capacity for consciousness of the right kind – namely, affective consciousness.

5. The Resonance Constraint

My preliminary objection to inclusion was that Blobs are not welfare subjects. But what exactly is wrong with supposing that the Blob’s welfare consists in staying alive, or being conscious, or seeing red things? We can explain what’s wrong with this idea by endorsing a plausible constraint on welfare values, according to which they must “resonate” with your intrinsic desires:

The Resonance Constraint on Welfare Goods: Something is good for you only if it resonates with your intrinsic desires.¹⁰

We violate the Resonance Constraint if we suppose that it’s good for the Blob to be alive, or conscious, or acquainted with red things, regardless of whether it desires these things or desires anything at all. This conception of welfare is implausible precisely because it implies that what is good for the Blob fails to resonate with the Blob in any way. As Peter Railton remarks, it is implausible that anyone can be alienated in this way from their own good: “It would be an intolerably alienated conception of someone’s good to imagine that it might fail in any such way to engage him” (1986: 9).

The Resonance Constraint undermines the idea that conscious experience is always good for you regardless of whether you want it or, indeed, whether you want anything at all. A more plausible view is that conscious experience is good for you only when it resonates with your desires. According to desire-based theories of pleasure, for example, pleasant experiences are good for you because you desire those experiences or because you desire whatever they represent. Indeed, hedonically neutral experiences can be good for you too so long as they resonate with your desires.

The Resonance Constraint is flexible enough to combine with multiple theories of welfare.¹¹ It is consistent with *hedonism*, the view that pleasure is the only welfare good, so long as pleasure is

¹⁰ The Resonance Constraint is a recurring theme in the literature on welfare: it is endorsed by Railton (1986), Rosati (1996), and Velleman (1998), among many others. Roelofs (2023) applies a version of the Resonance Constraint to Vulcans, but without committing himself to the further claim that affective consciousness is required to satisfy it.

¹¹ Here I adopt Parfit’s (1984: 493-502) familiar taxonomy of theories of welfare.

appropriately connected to desire: on a higher-order theory, it is the kind of experience that you desire, whereas on a first-order theory, it is the kind of experience whose objects you desire. Hedonism is not mandated, however, since pleasure is not the only thing we desire. The Resonance Constraint is equally compatible with *desire-satisfaction theories*, according to which the satisfaction of desire is good for you even when you desire other things besides pleasure. It is compatible with *hybrid theories*, which say that desire satisfaction is good for you only when the object of your desire is objectively worth desiring. Indeed, it may even be compatible with *objective list theories* so long as all objective welfare goods are ones that resonate with your desires. You cannot plausibly possess objective welfare goods such as friendship or achievement, for example, while remaining wholly indifferent towards them.¹² The Resonance Constraint excludes only the more extreme view that there are welfare goods, including consciousness, that are objectively good for you regardless of what you desire and, indeed, whether you desire anything at all. And yet this view is implausible insofar as it countenances welfare subjects, such as Blobs, who are completely indifferent to their own welfare.

The precise formulation of the Resonance Constraint raises questions that are tangential to the central concerns of this paper. What exactly is the relation that your welfare goods must stand in towards your desires in order to *resonate* with them in the right way? Do you actually need to desire what is good for you or is it enough that you would desire it under certain ideal conditions? For current purposes, we can finesse such questions by formulating the Resonance Constraint as a constraint on *welfare subjects*, rather than *welfare goods*:

The Resonance Constraint on Welfare Subjects: A being is a welfare subject only if it has intrinsic desires.

This constraint on welfare subjects may ultimately derive from a corresponding constraint on welfare goods. Even so, it is independently plausible on its own terms. Indeed, the best argument for any corresponding constraint on welfare goods is an abductive one: namely, that we need it order to give a principled explanation of this independently plausible constraint on welfare subjects.

¹² See Fletcher (2013: 215-17). Lin (2021: 880) makes similar points in defending sentientism: perhaps only conscious creatures can have desires of the right kind to possess these welfare goods. Unless he endorses the Resonance Constraint, however, he has no unifying explanation of why all welfare goods require consciousness.

We find similar reasoning in Peter Railton's seminal discussion of the Resonance Constraint. He articulates the driving idea as follows:

It seems to me that notions like good and bad have a place in the scheme of things only in virtue of facts about what matters, or could matter, to beings for whom it is possible that something matter. Good and bad would have no place within a universe consisting only of stones, for nothing matters to stones. (1986: 9)

Inanimate objects like rocks and stones are not welfare subjects because they have no intrinsic desires. And what goes for these inanimate objects goes equally for Blobs. Although they are conscious, they have no intrinsic desires that make things good or bad for them. Nothing matters *for a Blob* because nothing matters *to a Blob*.

6. Do Vulcans Have Desires?

Our next question is whether Vulcans can satisfy the Resonance Constraint. Can Vulcans have desires? Chalmers assumes that they can:

Vulcans' lives may be literally joyless, without the pursuit of pleasure or happiness to motivate them . . . But they nevertheless have serious intellectual and moral goals. They may want to advance science, for example, and to help those around them. They might even want to build a family or make money. They experience no pleasure when anticipating or achieving these goals, but they value and pursue the goals all the same. (2022: 343-4)

If so, then perhaps Vulcans are welfare subjects in virtue of having desires, rather than conscious experiences. This is less inclusive than Chalmers' official view, since it excludes Blobs from welfare status, although it has the compensating advantage of complying with the Resonance Constraint. Moreover, it threatens to generate counterexamples to both hedonic and affective sentientism by including Vulcans as welfare subjects with moral status when they have desires.

Why assume that Vulcans can have desires? Chalmers' main consideration is that we can imagine Vulcans (unlike Blobs) who are motivated to pursue various outcomes. Indeed, we can even imagine humanoid Vulcans who are motivated to act in all the same ways that we are. Perhaps these

Vulcans are *affective zombies* in whom unconscious motivations combine with conscious thoughts and perceptions to play the causal roles normally occupied by our affective experiences.¹³

According to the *motivational theory* of desire, the essence of desire is to combine with belief in the motivation of action. Here is Robert Stalnaker's classic formulation:

The Motivational Theory of Desire: To desire that p is to be disposed to act in ways that would tend to bring it about that p in a world in which one's beliefs, whatever they are, were true. (1984: 15)

Stalnaker combines the motivational theory of desire with a corresponding theory of belief:

The Motivational Theory of Belief: To believe that p is to be disposed to act in ways that would tend to satisfy one's desires, whatever they are, in a world in which p (together with one's other beliefs) were true. (1984: 15)

The idea is that belief and desire are inter-defined by their causal role in combining to motivate action. This is perhaps the default option for any functionalist theory of mind, according to which mental states can be functionally defined by their causal role.

The motivational theory implies that Vulcans desire to achieve their goals just so long as they are motivated to pursue them: they have these desires not because of their conscious experience, but rather because of their motivational dispositions. The problem is that the motivational theory implies that *zombies* can have desires for much the same reason. After all, they can be motivated to act in just the same way that Vulcans can. Hence, the motivational theory includes Vulcans within the class of welfare subjects only by including zombies too. This undermines sentientism, the view that only conscious creatures are welfare subjects.

Instead of abandoning sentientism, we should abandon the motivational theory instead. Here is my argument in a nutshell. Beliefs and desires not only motivate action but also provide *justifying*

¹³ Smithies and Weiss (2019: 30) introduce this definition of an affective zombie. Compare Strawson's (1994: 281) Aldebaranians: "They are not capable of any affect states at all, but they are capable of entering into . . . motivating states that are functionally very similar to states that we normally think of as desire states . . . in respect of the way in which they interact with a being's informational states to cause it to move in apparently goal-directed ways."

reasons for action. After all, our actions are justified, reasonable or rational when they are based on the justifying reasons provided by our beliefs and desires. And yet merely being motivated to do something is not enough for having any justifying to do it. Hence, motivation is not sufficient for belief or desire. This blocks the argument that zombies have beliefs and desires. As I'll explain, however, it also undermines Chalmers' contention that Vulcans have desires.

We can illustrate the point by considering examples of "Vulcanized" desire in which we excise any affectively valenced feelings associated with desire, while leaving their motivational dispositions intact. Here, for instance, is Warren Quinn's example of radioman:

Suppose I am in a strange functional state that disposes me to turn on radios that I see to be turned off. Given the perception that a radio in my vicinity is turned off, I try, all other things being equal, to get it turned on. Does this state rationalize my choices? Told nothing more than this, one may certainly doubt that it does. But in the case I am imagining, this is all there is to the state. I do not turn on the radios in order to hear music or get news. It is not that I have an inordinate appetite for entertainment or information. Indeed, I do not turn them on in order to *hear* anything. (1993: 236)

Radioman is motivated to act in ways that he believes will bring it about that radios are turned on. So, for example, he will press a button when he believes that this will turn on a radio. Nevertheless, I'll argue, he has no desire to turn on radios.¹⁴

Here are my assumptions about Quinn's example. First, radioman's actions are *intentional*, rather than mere reflexes, since his behavior is mediated by his beliefs and desires. He won't press the button when he believes it is wired up to explosives, for example, since he doesn't want to cause an explosion. Second, he has no *instrumental* desire, reason, or motivation to turn on radios. In particular, he is not disposed to derive pleasure from turning them on or displeasure from seeing them turned off. And, finally, although he is disposed to turn on radios, he is not disposed to *feel any desire* to turn them on. As Quinn says, the motivational disposition is "all there is to the state".

The motivational theory implies that radioman has a desire to turn on radios, since he is motivated to act in ways that he believes will turn them on. And, since he has no instrumental desire

¹⁴ Smithies and Weiss (2019) provide additional arguments that appeal to higher-level requirements on justifying reasons.

to turn on radios, it must be an intrinsic desire. So we can evaluate the motivational theory by considering whether radioman satisfies the following normative constraint on intrinsic desire:

The Normative Constraint: All intrinsic desires provide justifying reasons for action.¹⁵

The main reason to accept this normative constraint is that it explains why you always have some instrumental reason to do what you have reason to believe will satisfy your intrinsic desires. If you want gin, rather than vodka, then you thereby have some reason to order it from the bar. Of course, these desire-based reasons can be rebutted or undercut – say, by evidence that the gin is mixed with gasoline. Nevertheless, your intrinsic desires give you defeasible reasons that can justify action in the absence of defeating considerations. When you have misleading evidence – say, that your drink contains gin when in fact it’s gasoline – your intrinsic desires can motivate actions that are *bad* for you. Even so, your intrinsic desires give you justifying reasons of a kind that make it instrumentally justified, rational, or reasonable for you to act in ways that may turn out badly for you.¹⁶

Radioman’s motivational disposition violates this normative constraint on intrinsic desire. As Quinn observes, merely being motivated to turn on radios gives him no reason whatsoever to do so:

I cannot see how this bizarre functional state in itself gives me even a *prima facie* reason to turn on radios, even those I can see to be available for cost-free on turning. It may help explain, causally, why I turn on a particular radio, but it does not make the act sensible. (1993: 237)

It is not that radioman has reasons to turn on radios that are defeated by countervailing considerations, since we can stipulate that there is no cost in turning them on. Even so, the act of turning on radios seems utterly pointless when considered from his own perspective. I’m not suggesting that turning on radios seems pointless to him. Since radioman has no affectively valenced feelings about turning on

¹⁵ Note that this doesn’t imply the Humean theory of reasons, according to which all justifying reasons for action are provided by intrinsic desires. It is compatible, for example, with the Kantian thesis that morality provides reasons for action that need not align with your desires.

¹⁶ In this context, I am using the ambiguous or context-sensitive term ‘reason’ in what is often called the *subjective* sense, rather than the *objective* sense (e.g. Schroeder 2008). It counts in favor of an action by making it reasonable, rational, or justified in light of your subjective evidence, rather than good in light of the objective facts.

radios, it seems to him neither pointless nor worthwhile. It may never occur to him to ask whether there's any point in turning on radios. Nonetheless, he is motivated to turn them on anyway.

In summary, here is my argument against the motivational theory of desire. We can stipulate that radioman has no instrumental desire to turn on radios: if he has any such desire, it must be an intrinsic one. But all intrinsic desires provide justifying reasons for action. And yet radioman has no such reason to turn on radios. Therefore, he has no desire to turn on radios. Even so, he is motivated to do whatever he believes will bring it about that radios are turned on. Hence, the motivational theory is false: motivation is not sufficient for desire.

7. How Desires Justify Action

Our next question is what else radioman needs to have an intrinsic desire that gives him justifying reasons to turn on radios. What is the missing ingredient? In this section, I'll argue that radioman has no justifying reason to turn on radios because he has no disposition to feel any desire to do so. More generally, I'll argue for the following constraint on intrinsic desire:

The Affective-Phenomenal Constraint: All intrinsic desires provide justifying reasons for action by disposing you to experience affectively valenced feelings of desire.

When you desire something for its own sake, you are disposed to experience affectively valenced feelings of desire for it.¹⁷ These feelings of desire, and dispositions to have them, give you defeasible reasons for action because they represent what you desire as valuable.

The feeling of desire is something we all know from our own personal experience. We cannot do justice to the experience of being a kid in the candy store, a teenager in love, or a competitor striving for victory without mentioning the feeling of desire. In these examples, our desires are felt with great intensity, but we shouldn't overlook more mundane examples. When I consider whether to listen to the cricket on the radio, for instance, I feel attracted towards doing so. When I think about listening to the shipping forecast, in contrast, I feel mild aversion. These feelings of attraction and aversion, which may vary in their degree of felt intensity, are all feelings of desire.

¹⁷ Others who highlight the affectively valenced feeling of desire include Vadas (1984), Davis (1986), Strawson (1994), Johnston (2001), Chang (2004), Heathwood (2019), Smithies and Weiss (2019), and Siewert (2021).

Any valenced feeling of desire has three essential properties. First, it is *phenomenal*: there is something it's like for you to experience the feeling of desire. Second, it is *intentional*: you cannot feel desire without feeling desire *for* something, which is what your desire is about. And third, it is *affectively valenced*: you cannot feel desire without feeling positively or negatively about whatever your desire is about. Moreover, there is not just an incidental relation between these properties. On the contrary, the affectively valenced intentionality of feeling desire is essential to its phenomenal character. The feeling of desire *just is* a way of feeling positively or negatively about what you desire.¹⁸

This point is crucial for answering an objection. How could a mere *feeling* give you any reason to do anything? This might seem no better than the idea that being disposed to do something gives you a reason to do it. But there are two points to make in reply. First, the feeling of desire is not just a brute sensation with no intentional content. It is a feeling of desire *for something* in particular. And second, feeling the desire for something is not just feeling an urge, compulsion, or motivation to bring it about. It is a way of feeling *positively or negatively* about the intentional object of your desire. As such, it gives you a justifying reason to promote or avoid whatever your desire is about.

What I'm suggesting, in effect, is that feeling the desire to act gives you a justifying reason to act by representing what you desire under "the guise of the good". Feeling attracted to something is a way of representing it as attractive, whereas feeling averse to something is a way of representing it as aversive. One view is feelings of desire represent value in virtue of their *content*: they are neutral representations of the valenced content that their objects are good or bad. An alternative view, which I find more plausible, is that they represent value in virtue of their *attitude-type*: they are valenced reactions to neutral contents represented in other ways, such as perception, imagination, or thought. On either view, feelings of desire provide justifying reasons for action by representing what you desire in a positively or negatively valenced way.¹⁹

Quinn advances the closely related proposal that your desires give you justifying reasons to act only if you believe there is some *value* in doing so. To have such a reason, he says, "I need the *thought* that the direction in which I am psychologically pointed leads to something good" (1993: 242).

¹⁸ This coheres with intentionalism about phenomenal consciousness, according to which all phenomenal properties are identical with intentional properties construed broadly to include attitudes as well as contents (Chalmers 2010: Ch. 11).

¹⁹ The content view for desire is defended by Stampe (1987) and Oddie (2005), while the attitudinal view is defended by Tenenbaum (2007) and Schafer (2013). I don't claim that we always *act* under the guise of the good: perhaps, as Velleman (1992) argues, we can make perverse decisions to do things we don't want to do precisely because we believe there is nothing good about them. And yet such actions are not justified by our desires.

Arguably, however, this is an over-intellectualization. Children and animals can act rationally on their desires without any capacity for evaluative belief, thought, or judgment. Moreover, even mature adults can act for justifying reasons that conflict with their evaluative beliefs. As Nomy Arpaly (2002) argues, Huck Finn acts rationally and morally in helping Jim to escape slavery, despite believing that he ought to turn Jim over to the authorities. If so, then Quinn's condition is not necessary: your desires can give you justifying reasons for action in the absence of evaluative belief, thought, or judgment.

More importantly, evaluative belief is not *sufficient* to give you justifying reasons for action. Suppose radioman believes that there is some value in turning on radios. This belief cannot give him any justifying reason to turn on radios. After all, beliefs about value – like beliefs about anything else – are justified or unjustified depending on whether they are based on justifying reasons. If they are unjustified, then they cannot provide justifying reasons for action. If they are justified, they merely transmit justification from the reasons on which they are based. Either way, beliefs about value cannot provide justifying reasons for action in the same way as desires.²⁰

Despite these problems, we can salvage Quinn's proposal by regarding affectively valenced feelings of desire as a more primitive mode of evaluation than belief, thought, or judgment. Children and animals can feel desire, for example, without having any concept of value. Moreover, our feelings of desire sometimes conflict with our evaluative beliefs. Most importantly, our feelings of desire can supply justifying reasons for action in the absence of any corresponding evaluative beliefs. A perfectly good reason for doing a handstand, for example, is that you just feel like it regardless of whether you believe there is any value in doing so. And this feeling of attraction gives you a justifying reason for action because it represents the prospect of doing a handstand as an attractive one.

In summary, radioman has no justifying reason to turn on radios because he has no disposition to experience any positively valenced feeling of desire to turn them on. If you feel attracted to turning on radios just for its own sake, then you thereby have some reason to do so, since you represent the prospect of turning on radios as an attractive one. Similarly, if you have a standing disposition to feel attracted to turning on radios, then you have a standing reason to do so, even when you are thinking about other things. Of course, such reasons can be defeated by evidence that there is no value in turning on radios. But the mere disposition to turn on radios, without any disposition to feel positively about turning them on, gives you no reason – not even a defeasible reason – to do so.

²⁰ Incidentally, this is why Vulcans cannot satisfy the Resonance Constraint by having evaluative beliefs, rather than desires: beliefs about value cannot play the right kind of normative role in providing justifying reasons for action.

8. An Argument for Affective Sentientism

We now have all the ingredients in place to argue against Chalmers' inclusive version of sentientism in favor of a more restrictive version of affective sentientism, according to which the capacity for affective consciousness is required for being a welfare subject with moral status. Using this argument, we can explain why neither zombies nor Vulcans are welfare subjects with moral status.

We can begin with the Resonance Constraint, which is usually formulated as a constraint on welfare goods:

The Resonance Constraint on Welfare Goods: Something is good for you only if it resonates with your intrinsic desires.

As we saw in §5, we can reformulate this as a constraint on welfare subjects, rather than welfare goods:

The Resonance Constraint on Welfare Subjects: A being is a welfare subject only if it has intrinsic desires.

The Resonance Constraint explains why neither inanimate rocks nor conscious Blobs are welfare subjects: if you have no intrinsic desires, then nothing can be good or bad for you.

Can zombies or Vulcans satisfy the Resonance Constraint? Unlike rocks and Blobs, they may be motivated to bring about various outcomes, such as advancing science, building families, or making money. According to the motivational theory of desire, these motivational dispositions are sufficient for desire. If so, then perhaps zombies and Vulcans can satisfy the Resonance Constraint after all. Indeed, the Resonance Constraint is often interpreted as a thesis about motivation. As Connie Rosati observes, "Many have thought it is a necessary condition on something being good for someone that she be capable of being motivated by it" (2010: 227).

As I argued in §6, however, motivation is not sufficient to satisfy the Resonance Constraint. Intrinsic desires play a normative role in providing justifying reasons for action and not just a causal role in motivating action:

The Normative Constraint on Intrinsic Desire: All intrinsic desires provide justifying reasons for action.

The motivational theory of desire violates this normative constraint, since the mere motivation to do something gives you no justifying reason to do it. In Quinn's example, radioman is motivated to turn on radios, but this gives him no justifying reason to turn them on.

Given this normative constraint on intrinsic desire, the Resonance Constraint articulates a connection between the normative concepts of welfare and justifying reasons for action:

The Normative Resonance Constraint on Welfare Goods: Something is good for you only if it resonates with your intrinsic desires, which provide you with justifying reasons for action.

We can now explain why the motivational interpretation of the Resonance Constraint is inadequate. The mere motivation to pursue an outcome is not sufficient to make it even *pro tanto* good for you. After all, there is nothing good for radioman about turning on radios in spite of his motivation to turn them on. Turning on radios is not good for him because his motivation to turn them on gives him no justifying reason to do so.

What radioman is missing is any disposition to experience affectively valenced feelings of desire to turn on radios. Your intrinsic desires normally dispose you experience such feelings. Moreover, these feelings – and your dispositions to have them – give you justifying reasons for action by representing what you desire as valuable. I argued in §7 that this is how your intrinsic desires provide you with justifying reasons for action:

The Affective-Phenomenal Constraint on Intrinsic Desire: All your intrinsic desires provide you with justifying reasons for action by disposing you to experience affectively valenced feelings of desire.

Radioman's motivation to turn on radios gives him no justifying reason to turn them on because he has no disposition to experience the positively valenced feeling of desire to turn them on.

If we combine this with the Resonance Constraint, we get the following result:

The Affective-Phenomenal Constraint on Welfare Goods: Something is good for you only if it resonates with your dispositions to experience affectively valenced feelings of desire.

We can now explain why there is nothing good for radioman about turning on radios despite his motivation to produce this outcome. Turning on radios is not good for him because it doesn't resonate with any of his dispositions to experience positively valenced feelings of desire.

What goes for radioman goes equally for zombies and Vulcans. They violate the following constraint on welfare subjects, which follows from our earlier premises:

The Affective-Phenomenal Constraint on Welfare Subjects: A being is a welfare subject only if it is disposed to experience affectively valenced feelings of desire.

Vulcans may be motivated to advance science, to build a family, or to make money, but they are not disposed to experience any positively valenced feelings about these outcomes. While zombies are not disposed to feel anything at all, Vulcans are not disposed to feel positively about anything. As a result, their motivations give them no justifying reasons to pursue these outcomes. That is why achieving these outcomes is not good for them. More generally, nothing is good or bad for them because they are not disposed to feel positively or negatively about anything.

As promised in §4, this argument gives us an explanation of sentientism. The Resonance Constraint, together with further constraints on desire, implies that you need some capacity for consciousness in order to be a welfare subject. Zombies are not welfare subjects because they have no capacity for consciousness at all. This explanation of sentientism undercuts inclusion, however, since it implies that Vulcans are not welfare subjects either. Vulcans are conscious beings, unlike zombies, but they have no capacity for affective consciousness. As a result, they violate the constraints on desire needed to satisfy the Resonance Constraint. Since only welfare subjects have moral status, it follows that neither zombies nor Vulcans have moral status.

In summary, we cannot explain why zombies lack moral status without denying that Vulcans have moral status either. So we cannot maintain that consciousness is a necessary condition for moral status unless we also deny that it is a sufficient condition. To have moral status, you need the capacity for consciousness of the right kind – namely, affectively valenced consciousness that represents the world in a positive or negative way. Neither zombies nor Vulcans satisfy this condition.

9. Conclusions

Many philosophers endorse sentientism: the thesis that only conscious creatures are welfare subjects endowed with moral status. This paper explores a further question that arises for all proponents of

sentientism: what kind of consciousness is required for welfare and moral status? As we've seen, Chalmers argues for an inclusive version of sentientism, according to which consciousness of any kind is sufficient for moral status. On this view, all consciousness has intrinsic welfare value. Hence, all conscious beings are welfare subjects, since they are capable of possessing at least one welfare value.

This paper contains multiple arguments against Chalmers' inclusive version of sentientism. First, I argued that it is not supported by intuitions about the moral impermissibility of killing Vulcans, since these intuitions are insensitive to the distinction between moral status and moral significance. Second, I argued that there are intuitive counterexamples to inclusion: although Blobs are conscious beings, they are not welfare subjects. Third, I argued that the apparent plausibility of inclusion trades on equivocating between distinct concepts of value: while it is plausible that all consciousness has intrinsic value, it is not so plausible that all consciousness has welfare value.

Fourth, I argued that Chalmers cannot explain why sentientism is true in the first place. If all consciousness has welfare value, then it follows that all conscious beings are welfare subjects, since they are capable of possessing at least one welfare value – namely, consciousness. This doesn't imply that only conscious beings are welfare subjects, however, since it doesn't exclude the possibility that unconscious beings can possess other welfare goods besides consciousness. Hence, this inclusive view provides no principled explanation for sentientism.

Finally, I argued that we cannot explain properly sentientism without abandoning inclusion. The explanation I proposed is that you need consciousness to satisfy the Resonance Constraint: only conscious creatures have justifying reasons for action of the kind that you need in order to be a welfare subject with moral status. More specifically, however, you need *affective* consciousness to satisfy the Resonance Constraint, since only affective consciousness provides justifying reasons for action of the requisite kind. This explanation of sentientism undermines inclusion: it explains why zombies lack moral status only by excluding Vulcans too.

References

- Arpaly, Nomy. 2002. Moral Worth. *Journal of Philosophy* 99: 223-245.
- Bradford, Gwen. 2023. Consciousness and Welfare Subjectivity. *Noûs* 57: 905-921.
- Chalmers, David. 2010. *The Character of Consciousness*. Oxford University Press.
- Chalmers, David. 2022. *Reality+: Virtual Worlds and the Problems of Philosophy*. Norton.

- Chang, Ruth. 2004. Can Desires Provide Reasons for Action? In *Reason and Value: Themes from the Moral Philosophy of Joseph Raz*, edited by Jay Wallace, Philip Petit, Samuel Scheffler, and Michael Smith, 56-90. Oxford University Press.
- Davis, Wayne. 1986. Two Senses of Desire. In *The Ways of Desire*, edited by Joel Marks, 63-82. Precedent.
- Feldman, Fred. 1988. Two Questions about Pleasure. In *Philosophical Analysis: A Defense By Example*, edited by David Austin. Kluwer.
- Fletcher, Guy. 2013. A Fresh Start for the Objective-List Theory of Well-Being. *Utilitas* 25: 206-220.
- Heathwood, Chris. 2007. The Reduction of Sensory Pleasure to Desire. *Philosophical Studies* 133: 23-44.
- Heathwood, Chris. 2019. Which Desires are Relevant to Well-Being? *Noûs* 53: 664-688.
- Jacobson, Hilla. 2019. Not Only a Messenger: Towards an Attitudinal-Representational Theory of Pain. *Philosophy and Phenomenological Research* 99: 382-408.
- Johnston, Mark. 2001. The Authority of Affect. *Philosophy and Phenomenological Research* 63: 181-214.
- Kriegel, Uriah. Forthcoming. The Value of Consciousness to the One Who Has It. In *The Importance of Being Conscious*, edited by Geoffrey Lee and Adam Pautz. Oxford University Press.
- Lee, Andrew. 2018. Is Consciousness Intrinsically Valuable? *Philosophical Studies* 176: 655-671.
- Lin, Eden. 2020. Attitudinal and Phenomenological Theories of Pleasure. *Philosophy and Phenomenological Research* 100: 510-524.
- Lin, Eden. 2021. The Experience Requirement on Well-Being. *Philosophical Studies* 178: 867-886.
- Moore, G. E. 1903. *Principia Ethica*. Cambridge University Press.
- Nozick, Robert. 1974. *Anarchy, State, and Utopia*. Basic Books.
- Oddie, Graham. 2005. *Value, Reality, and Desire*. Oxford University Press.
- Parfit, Derek. 1984. *Reasons and Persons*. Oxford University Press.
- Quinn, Warren. 1993. Putting Rationality in its Place. In *Morality and Action*, by Warren Quinn, 228-255. Cambridge University Press.
- Roelofs, Luke. 2023. Sentientism, Motivation, and Philosophical Vulcans. *Pacific Philosophical Quarterly*. 104: 301-323.
- Railton, Peter. 1986. Facts and Values. *Philosophical Topics* 14: 5-31.
- Rosati, Connie. 1996. Internalism and the Good for a Person. *Ethics* 106: 297-326.
- Rosati, Connie. 2009. Relational Good and the Multiplicity Problem. *Philosophical Issues* 19: 205-234.
- Schafer, Karl. 2013. Perception and the Rational Force of Desire. *Journal of Philosophy* 110: 258-281.

- Schroeder, Mark. 2008. Having Reasons. *Philosophical Studies* 139: 57-71.
- Shepherd, Joshua. 2018. *Consciousness and Moral Status*. Routledge.
- Siewert, Charles. 2021. Consciousness: Value, Concern, Respect. *Oxford Studies in Philosophy of Mind* 1: 3-40.
- Singer, Peter. 1975. *Animal Liberation*. Harper Collins.
- Smith, Michael. 1994. *The Moral Problem*. Blackwell.
- Smithies, Declan and Weiss, Jeremy. 2019. Affective Experience, Desire, and Reasons for Action. *Analytic Philosophy* 60: 27-54.
- Smithies, Declan. MS. [Title to be added later pending blind review.]
- Stalnaker, R. 1984. *Inquiry*. MIT Press.
- Stampe, Dennis. 1987. The Authority of Desire. *Philosophical Review* 96: 335-381.
- Strawson, Galen. 1994. *Mental Reality*. Oxford University Press.
- Tenenbaum, Sergio. 2007. *Appearances of the Good: An Essay on the Nature of Practical Reason*. Cambridge University Press.
- Vadas, Melissa. 1984. Affective and Non-Affective Desire. *Philosophy and Phenomenological Research* 45: 273-279.
- Van der Deijl, Willem. 2019. Is Pleasure All That is Good About Experience? *Philosophical Studies* 176: 1769-1787.
- Van der Deijl, Willem. 2020. The Sentience Argument for Experientialism about Welfare. *Philosophical Studies* 178: 187-208.
- Velleman, David. 1992. The Guise of the Good. *Noûs* 26: 3-26.
- Velleman, David. 1998. Is Motivation Internal to Value? In *Preferences*, edited by Christoph Fehige and Ulla Wessels, 88-102. Walter de Gruyter.