Journal of the American Philosophical Association (2025) 1–19
DOI:10.1017/apa.2025.10006

Welfare Subjects and Autopoiesis

ABSTRACT: *A welfare subject is an entity with a good of its own. Welfare subjects have interests. Things matter to them. It is uncontroversial that typical adult humans are welfare subjects. It is uncontroversial that rocks are not welfare subjects. Just what makes this so is a matter of controversy. The default view is that sentience is necessary for welfare subjecthood. A competing view is that teleological organization suffices for welfare subjecthood. This article challenges both views by developing a third view that sits between them. An autopoietic entity is one that engages in self-production in a way that is adaptive to its environment. An autopoietic entity sustains, maintains, and renews itself on an ongoing basis. This article argues that all autopoietic entities are welfare subjects. This view explains why sentient entities such as animals are welfare subjects. It explains why merely teleologically organized entities such as automobile engines are not.*

KEYWORDS: well-being, welfare, autopoiesis, sentience, teleological organization

1. Introduction

A welfare subject is an entity with a good of its own, a being for whom things can be good or bad. It is generally agreed that humans and mice are welfare subjects, but cars and rocks are not. It is a matter of controversy what makes this so. This article introduces a new candidate for a sufficient condition for welfare subjecthood. The thesis is that autopoiesis suffices for welfare subjecthood, where an entity is autopoietic if it creates, sustains, and maintains itself in a way that responds and adapts to its environment.

In Section 2, I explain the concept of a welfare subject more fully. In Section 3, I introduce two views: the dominant view that sentience is necessary for welfare subjecthood and an opposing view that being teleologically organized is sufficient for welfare subjecthood. In Section 4, I explain my proposed sufficient condition of autopoiesis more fully. I then argue that autopoiesis underlies and explains the welfare subjecthood of the entities that we already accept as welfare subjects. In Section 5, I explain some of the autopoietic behavior of plants and argue that their autopoiesis makes them welfare subjects. I pause in Section 6 for some observations about the relationship between the autopoiesis standard and the sentience and teleological organization standards, including that the autopoiesis standard explains why sentient entities are welfare subjects and teleologically organized entities are not. I observe that on the autopoiesis standard, autopoietic machines of the near future will qualify as welfare subjects. In Section 7, I develop this implication into an objection to the effect that my view has absurd implications

about our reasons to advance plant and machine welfare. I consider an additional objection in Section 8.

2. The Concept of a Welfare Subject

We first need to outline the concept of a *welfare subject* before examining competing conceptions of a welfare subject (see, for example, Rawls (1971: 5) for the concept/conception distinction). The concept of a welfare subject is the concept of a being with a good of its own. A welfare subject is a being for whom things can go better or worse, who can fare well or fare badly, whose existence through time can go well or badly for it. Welfare subjects can be benefitted or harmed.

The archetypal welfare subject is, of course, a human being. We often express claims about human welfare (or, equivalently for our purposes, *well-being*) with the *good for* locution: Pleasure is good for humans; It is good for humans to have friends; It is not good for humans to drink cyanide. (The well-being of a welfare subject is its intrinsic good as opposed to what is instrumentally good for realizing its intrinsic good. Drinking cyanide is perhaps not intrinsically bad for a human, but only instrumentally bad. We can elide this distinction here, as nothing can be instrumentally good for an entity unless something can be intrinsically good for it—that is, unless it is a welfare subject.)

Not all good-for talk, however, indicates the presence of a welfare subject. We say that it is good for an automobile engine to have regular oil changes, but the consensus among philosophers of well-being is that an automobile engine is not a welfare subject. An engine does not have a good of its own. Rather, saying that regular oil changes are good for an engine seems to be elliptical for a thought to the effect that regular oil changes will help preserve the condition of the engine to perform the function for which it was designed by humans and for humans.[1] So the presence of a good-for a thing is not sufficient for it to be a welfare subject.

What more seems to be needed in order for an entity to have a good of its own is for things to be able to *matter* to the entity itself. The entity needs to have an *interest* in how things go for it. It matters as much to an engine whether it gets its oil changed as it does to a rock whether it is shattered by a hammer, which is to say not at all. If things do not matter to an entity—if it does not have an interest in anything—then it seems that the entity cannot have a good of its own. Much of the debate about what counts as a welfare subject will turn on what sort of mattering or having of an interest is at issue, as we will begin to see when we take up the dominant conception of welfare subjecthood presently.

3. Sentience and Teleological Organization

For our purposes, we can take sentience to be the capacity to experience pleasure and pain or—perhaps a little more broadly—the capacity for conscious experience that is positively or negatively valenced. In the literature on well-being, the default position

[1] Bradley (2015: 9) and Rosati (2009: 214, 222) give this example and this analysis. Kraut (2007) believes that there is just one good-for relation. Rosati (2009) criticizes Kraut's position.

https://doi.org/10.1017/apa.2025.10006 Published online by Cambridge University Press

is that nothing can matter to an entity unless it is sentient. Among others, Rosati (2009: 225) and Singer (1975: 9) are both fairly explicit that (a) sentience involves the capacity to experience pleasure and pain and that (b) being a welfare subject and having an interest require sentience. Moore and Crisp identify the traditional practice as one that assumes "that only sentient individuals are capable of well-being" (1996: 601). So the default position is that sentience is a necessary condition for welfare subjecthood. To have a good of one's own and have an interest, things have to be able to matter to one in this minimal sense.

The default condition assumes one answer to the question of what is necessary to have an interest, which is that *having an interest* requires the capacity to be *interested in* something. And "for this to be the case a living thing must be a conscious being with the capacity to experience events as positively or negatively valenced" (Rosati 2009: 225). This line of thought can be traced to Railton's classic statement that "what is intrinsically valuable for a person must have a connection with what he would find in some degree compelling or attractive" (1986: 9). There is not normally much argument given for this default condition other than a sense of: 'how *could* x – an object, event, or state of affairs – matter to a subject S unless S had the capacity consciously to experience x as positive or negative, to care about x, or mind whether x occurs or is present?' The engine cannot care whether it has its oil changed, so it is not a welfare subject. Railton simply says that "[i]t 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).

Against the default position, one can argue that S can have *an interest in* x and x can *matter* to S without S being *interested in* x in any *conscious* way and without x *mattering* to S in any *conscious* way.[2] I develop this argument beginning in the next section. Before that, however, notice that there is a loose sense in which it matters to an engine whether its oil is changed regularly. Again, this is just the sense in which a lack of regular oil changes will disable the engine from serving the purpose for which it was designed by and for humans. Again, most philosophers of well-being believe that this sense of mattering is irrelevant to the ascription of welfare to an engine.

Let us call the sense in which things can matter to an engine the *teleological organization* sense of mattering. A thing that is teleologically organized has, of course, a telos—an end or purpose. The purpose of an automobile engine is to provide propulsion, say. A lack of regular oil changes interferes with the capacity of the engine to achieve that end. It is bad for the engine in this teleological sense. It is in this same sense that it is good for cast iron cookware to be seasoned with oil, bad for some types of batteries to be left on a charger for too long, and bad for pianos to be subjected to excessive humidity. Cookware, batteries, and pianos are teleologically organized and certain conditions conduce to or interfere with their ability to realize or serve their ends. Things matter to them in this sense.

It seems to stretch our welfare talk too far to say that such artifacts are genuine welfare subjects with a well-being—a good of their own—that can be advanced or

[2] In the environmental ethics literature, a parallel debate often puts the question just this way. See Taylor (1981), Varner (1998), Rolston (2012), Hettinger (2012), and Basl and Sandler (2013). For reasons of space and focus, I do not engage this literature here.

https://doi.org/10.1017/apa.2025.10006 Published online by Cambridge University Press

impeded. Basl and Sandler (2013) argue for the teleological conception of welfare subjecthood and admit this stretch, but are willing to bite this bullet. A key question of this article is whether there is a welfare-relevant sense of mattering and having an interest that is more robust than the teleological organization sense but less demanding than the default sentience sense. The central claim of this article is that there is such sense, to which we turn in the next section. Once that sense has been developed, it will be easy to see (in Section 6.1) that the teleological organization sense of mattering is not sufficient for mattering in the sense needed for welfare subjecthood.

4. Autopoiesis

On the view to be offered here, a sufficient condition for an entity to be a welfare subject is for that entity to be autopoietic. An entity is autopoietic in this sense when it sustains, maintains, and renews itself in a way that responds and adapts to its environmental conditions. Said another way, an autopoietic entity produces itself on an ongoing basis where its productive activities are responsive to the challenges and opportunities presented by its situation.[3]

In this section, I discuss some entities widely accepted as welfare subjects. I explain their autopoietic behavior. I argue that their autopoietic behavior is what makes them welfare subjects. Their autopoietic behavior, I argue, reveals what matters to them in a welfare-relevant sense of mattering.[4] In the next section, I turn to entities that are not commonly regarded as welfare subjects and argue that their autopoietic behavior equally makes them welfare subjects.

Humans are autopoietic. While we do not literally create ourselves anew every day, we engage in a variety of behaviors aimed at keeping ourselves going as the sort of things we are. We sustain ourselves by seeking food and drink, taking it in, and eliminating waste. We sleep to renew ourselves. We maintain ourselves by regulating our maintenance activities according to what shows up in our environment as mattering to us and worthy of response. (I borrow this locution of *showing up* from Maher (2017: 72 *et passim*)). So we avoid, flee, or fight against threats to our bodies and lives by avoiding slippery surfaces and extreme temperatures, for instance, and by taking evasive or defensive action against threats to our survival, such as bears and coronaviruses. Other things do not show up to us as mattering or worthy of response, such as leaves blowing in the wind, ticking clocks, or voices in conversation in coffee shops. While most of these behaviors are voluntary actions, much of our autopoiesis takes place through nonconscious behaviors and mechanisms that came about through natural selection: our eyelids blink in dusty conditions, our pulse and respiration quicken when running from danger, our food

[3] In this conception of autopoiesis, I am following Thompson (2007: Chapters 5-6), who credits Maturana and Varela (1973, 1980) and Varela, Maturana, and Uribe (1974). See also Maher (2017: Chapters 3 and 6), who follows Thompson in arguing that autopoiesis in this sense is sufficient for mindedness.

[4] The thought that autopoietic behavior reveals things as mattering or significant for the autopoietic entity is not mine, but is Maher's (2017: 72 *et passim*). The thought is mine, however, that this mattering or significance is entirely sufficient to establish that the entity has a good of its own.

https://doi.org/10.1017/apa.2025.10006 Published online by Cambridge University Press

digests, our wounds heal, and our bodies produce antibodies to fight bacteria and viruses.

These autopoietic behaviors reveal an array of things that matter to us, that are in our interest, that affect our well-being. At the conscious level, it matters to us whether we have sufficient food or are hungry. We have an *interest in* food and are also *interested in* it. Yet it also matters to us whether our pulse and respiration quicken when running from danger. We will fare less well if they do not. Humans have an *interest in* proper cardiopulmonary function, whether or not they have ever taken notice of it—whether or not they are *interested in* it. The balance of gut microbiota matters to a person's well-being—they have an *interest in* their gut health even if they are not *interested in* it.

Like humans, non-human animals regarded as welfare subjects are autopoietic. Deer, for example, engage in many of the same autopoietic behaviors as humans. They eat, drink, eliminate waste, and sleep. They regulate their maintenance activities according to what shows up in their environment as mattering and as worthy of response. So, they increase their food consumption during cold weather to keep themselves warm, travel in herds for safety, and flee threats to their lives such as coyotes. Other things do not show up to them as mattering or worthy of response, such as scampering chipmunks, blowing leaves, or tractors running in the distance. While many autopoietic behaviors of deer are voluntary, much of their autopoiesis takes place through nonconscious behaviors and mechanisms that came about through natural selection. Their pupils dilate in dark conditions, their pulse and respiration quicken when running from danger, their food digests, their wounds heal, and their bodies produce antibodies to fight bacteria and viruses.

These autopoietic behaviors reveal or indicate an array of things that matter to deer, that are in their interest, that affect their well-being. At the conscious level, it matters to them whether they have sufficient food or are hungry. They have an interest in food and are also interested in it, presumably due to the negatively valenced experiences of feeling hungry and feeling cold and the positively valenced experiences of feeling satiated and feeling warm. Yet, as with humans, it also matters to deer whether their pulse and respiration quicken when running from danger. They will be worse off if they do not. Deer have an *interest in* proper cardiopulmonary function, even though they have never taken notice of it—even though they are not *interested in* it. Similarly, the amount of fat a deer has going into winter matters to a deer's well-being, even though it is not consciously interested in its fat covering.

Thus, autopoiesis is something in common between the two sorts of welfare subjects examined, humans and deer. We can go further and claim that humans and deer do not both merely happen to be both autopoietic and welfare subjects. Autopoiesis is what *makes* humans and deer welfare subjects. Their autopoietic behavior shows *that* things matter to them. It also shows *what* things matter to them. Autopoiesis involves regulating behavior with respect to things that are pursued (or attracted, welcomed, followed, etc.) or avoided (or fled, repelled, attacked, etc.) and *thereby* shown as mattering.

Two notes before moving on: First, I said above that autopoietic activity reveals things as mattering. We can clarify that by saying that autopoietic activity simultaneously makes it the case that things matter and also constitutes and

https://doi.org/10.1017/apa.2025.10006 Published online by Cambridge University Press

singles out the things that matter. I mean to say that the things that matter to a welfare subject are not antecedently present behind a cloak but then revealed when the cloak is lifted by autopoietic activity. Rather, the nature of autopoietic activity is to constitute the things that matter by virtue of the entity's differential responses to things. The things that are pursued or avoided through autopoietic activity are thereby constituted as the things that matter to the entity and affect its well-being.

Second, allow me to clarify the nature of my claim that autopoiesis makes something a welfare subject. I do not aim to deduce this claim, either above or below. Rather, I am making an abductive move to the effect that autopoiesis offers a satisfying explanation of welfare subjecthood in the cases of some familiar welfare subjects, so we should explore the hypothesis that autopoiesis explains welfare subjecthood in general. This article is an exploration and investigation of this hypothesis, including an evaluation of the explanatory fruitfulness of autopoiesis as grounding welfare subjecthood in comparison to teleological organization or sentience grounding it. My aim is to make the claim that autopoiesis makes for welfare subjecthood plausible and likely.

5. Plants

Philosophers do not doubt that the class of welfare subjects includes living things that are both sentient and autopoietic, such as humans and deer, but also apes, pigs, rabbits, birds, and so on. So let us move right to the case of insentient living things that are autopoietic, with plants as our example. Again, the hope is to make it plausible that, due to their autopoietic behavior, things matter to plants, that they have *interests* even though they may not have the mental machinery to be *interested in* anything. No theoretical commitments—for example, about plant agency—will be undertaken if we take *behavior* in a broad sense. In this sense, the term encompasses behaviors ranging from conscious, intentional actions of humans to clearly nonconscious, nonintentional happenings to passive entities, such as the behavior of shingles in a windstorm.

Consider first the thigmomorphism of bean plants. (This example and most of the following examples and explanations of autopoietic plant behavior are drawn from Maher (2017).) Thigmomorphism is a change in form resulting from mechanical stimulus. Bean plants grown in windy conditions respond to the wind by growing thicker stems compared to those grown in calm conditions. So, the bean plant produces and maintains itself in one way or in another depending on its environmental conditions. The plant reveals an interest in having a stem that can withstand its windy conditions by creating itself as wind-resistant in response to its conditions.

Phototropism is oriented growth in response to light stimulus. Where a wooded area meets a meadow, trees bend out over the meadow, reaching toward the sunlight. This is especially so (in the northern hemisphere) if the border runs east to west with the meadow to the south, so that the trees reach toward the abundant sunlight from the south. The plants grow in this fashion, of course, in order to take care of themselves by producing more energy from photosynthesis. They thereby reveal what is good for them. Note, importantly, that in this case and the case of the bean

https://doi.org/10.1017/apa.2025.10006 Published online by Cambridge University Press

plants in windy conditions, plants are not just taking care of themselves to get what is good for them, but doing so in a way that is adaptive to their variable circumstances. An oak tree growing not on the edge of a meadow but in the middle of a dense forest will grow straight up, because that will maximize its energy production in those circumstances. Bean plants and oak trees create and maintain themselves in one way or in another in order to get what is good for them.

Other autopoietic phenomena include root growth around solid impediments and flower heads following the sun. These reveal that subterranean solid objects and the position of the sun matter to plants, but they reveal something else. They reveal that the past has significance to plants in directing their present behavior toward what is good for them. The root is now growing horizontally because in the recent past it was prevented from growing downward. After some more time passes, it will grow downward again in order best to achieve nutrition and anchorage. The plant is back facing the sun at sunrise because that is where the sun was yesterday morning.

Not only does the past matter to plants, but so do durations of past states of affairs. A Venus flytrap requires that two triggers be tripped within 20 seconds of each other in order to close. When it closes to catch an ant, it does so because a trigger was tripped just now and another one was tripped, say, 12 seconds ago. This matters to it, because otherwise it wastes energy closing for false alarms. A plant displaying vernalization flowers only when it has been exposed to a prolonged period of past cold, also to avoid wasting energy. Thus, the array of things that matter to plants—that show up to them as relevant to the regulation of their present behavior—includes present and past events and states of affairs, including durations of past states of affairs.

Additionally, plants take care of themselves with defensive mechanisms and responses to predators. As one basic defensive mechanism, consider the production of nicotine by wild tobacco plants, which is poisonous to most insects that eat them. It is not, however, poisonous to hornworm caterpillars. As Maher describes it, "[w]hen grazed upon by these caterpillars, wild tobacco releases a variety of chemicals …, which attract big-eyed bugs, which eat the eggs and larvae of the caterpillars, defending the tobacco plant" (2017: 117).

Not only do plants respond by defending themselves against past and present attack, but they also anticipate future attack and engage in preemptive defense. Wild lima beans, when attacked by one type of mite, release methyl salicylate, which attracts a second type of mite. Those second mites attack the first type of mite. Bean plants in the immediate vicinity start producing methyl salicylate, even though they have not yet been attacked. Acacia trees have a similar anticipatory defense mechanism. Kudu eating the leaves of acacia trees cause the trees to release tannins that are poisonous to the kudu. Neighboring acacia trees that have not been attacked also release tannins, preemptively, much as the neighboring lima beans engage in anticipatory release of methyl salicylate. (I follow Maher closely here, but see (2017: 99–100) for a somewhat fuller description.)

I said above that some things show up to us humans as worthy of response (approaching bears) and that others do not (ticking clocks). We have just seen a variety of things that matter to plants. These things show up to plants and they regulate their self-productive behaviors in response. Other things do not show up to

https://doi.org/10.1017/apa.2025.10006 Published online by Cambridge University Press

plants. As far as we know, a chipmunk scurrying on the ground next to a cherry tree and the snort of a nearby deer do not matter to the tree in the relevant sense of mattering. Although the chipmunk and deer surely affect the tree by producing vibrations in its trunk and leaves, say, these creatures do not (as far as we know) elicit a response from the tree, just as ticking clocks and coffee house chatter generally do not elicit responses from us. Just as we can distinguish background noise from what matters to us, so can plants. The ability of plants to make this distinction between background noise and what matters – and then to take care of themselves in response to the things that matter—reveals a world of things that are significant to plants, that matter to them. Plants have interests in these things even though plants are not interested in them.

Let us collect some major points that seem relevant to establishing that the autopoietic activity displayed by plants suffices for their welfare subjecthood. First, these entities take care of themselves with behaviors that keep them going as the sort of things they are. They do not merely consume energy and expel waste as an engine does or seek a single end as a thermostat or other homeostatic entity does. They take care of themselves by responding and adapting to a wide variety of environmental stimuli, events, and conditions in the past, present, and anticipated future.

Second, this taking care often takes the form of self-creation in one way rather than another: shoots grow up, roots grown down and around impediments, bean stalks thicken in windy conditions, shoots grow this way rather than that depending on the direction of the sunlight, plants do or do not release certain chemicals depending on the presence of predators, and so on.

Third, this adaptive, self-sustaining, self-maintaining, self-creating activity reveals an array of things that show up to them as significant and worthy of response, that are thereby non-consciously valenced for them, that matter to them, that are good or bad for them. The autopoietic activity shows *that* things matter and *what* things matter. Energy from the sun and strong stalks are good and to be pursued. Root impediments and hornworm caterpillars are bad and to be avoided or eliminated. Scurrying chipmunks do not matter. Thus, plants—and any other such autopoietic entities—have interests in the things that are shown (i.e., constituted) through their autopoietic activity to be positively valenced for them. This is so even if they do not have inner lives and subjective experiences like ours, even if they cannot be interested *in* things that are in their interest. Since to have a good of its own—to have an interest, for things to matter—is what it is to be a welfare subject, autopoietic entities such as plants are properly welfare subjects.

Thus, with the cases surveyed here of robust 'matterings-to' manifested and revealed by plants, it is apparent that the specific sort of mattering-to manifested by sentient entities is not necessary for welfare subjecthood. To insist that mattering requires feeling pleasure and pain or sentience more generally now seems to beg the question about the scope of welfare subjects by excluding from the outset the possibility that there could be things with a good of their own that are built substantially differently than humans.

Understandably, the advocate of the necessity of sentience for welfare subjecthood is going to have some worries about opening the door to non-conscious mattering. In

https://doi.org/10.1017/apa.2025.10006 Published online by Cambridge University Press

particular, they will worry (1) that too many entities may slip through the door as welfare subjects that intuitively should be on the other side of the door and (2) that we moral agents will be excessively burdened with reasons to promote the welfare of those entities. I will segue to these important concerns right after making some observations about positive aspects of the view on offer here.

6. Observations

6.1 Teleological organization does not suffice for welfare subjecthood

We are now in a good position to explain why teleologically organized entities that are not autopoietic are not welfare subjects. This is because the sort of mattering-to displayed by teleologically organized entities differs in kind from the mattering-to displayed by autopoietic entities. This difference is rooted in how autopoietic entities take care of themselves, especially by adapting to and responding to their circumstances, but mere teleological entities do not.

An autopoietic entity such as a plant *seeks* what is good for it, *avoids* what is bad for it, *repairs* itself, *sustains* itself, *reacts* very differently depending on the opportunities and challenges presented by its environment. A basic teleological entity such as an engine with a governor, a cooling system with a thermostat, or a heat-seeking missile does none of that. These are mere input-output machines with a few teleological aspects, such as to run at a constant speed, to keep the ambient temperature within a set range, or direct itself toward the greatest infrared radiation. An engine does not create, maintain, or repair itself, say by changing its own oil or repairing damage to its components if its oil is not changed regularly. A cooling system does not change its own filters or repair its own refrigerant leaks. Mere teleological entities do not engage in these or any other self-productive or self-maintaining activities.

In a very loose sense, an engine with a governor reveals that its speed matters to it by seeking a specific speed, and a heat-seeking missile reveals that the source of heat matters to it by changing its course in response to the location of the heat. But these are not the robust matterings-to revealed by adaptive, self-creating, self-sustaining, self-repairing autopoietic entities. The mattering-to revealed by autopoietic entities is much richer than that revealed by teleological entities. This explains why autopoietic entities are welfare subjects but mere teleological entities are not.

6.2 The autopoietic standard (probably) explains the sufficiency of sentience

It is a reasonable conjecture that all sentient entities are autopoietic. Sentience is a primary way for a sentient entity to perceive stimuli from its environment as positively or negatively valenced. The entity can then respond in a way that avoids the negative stimuli and pursues the positive stimuli which, in turn, tends to the ongoing sustenance, maintenance, and self-creation of the entity for its own good. That all sentient entities are autopoietic is merely a conjecture that would need to be examined scientifically. If it is true, then the sentient organisms that we already think of as welfare subjects turn out to be welfare subjects *because* they are autopoietic. It is

https://doi.org/10.1017/apa.2025.10006 Published online by Cambridge University Press

a nice bonus for the view advanced here that it could provide an explanation for why sentient organisms are welfare subjects.

What if the conjecture turned out to be false because we discovered or built a sentient entity that was not autopoietic? Even if we were convinced that this thing was a welfare subject, that would not do any violence to the view advanced here. The view offered here is that autopoiesis is sufficient for welfare subjecthood, not that it is necessary. If it is possible for something to be a sentient non-autopoietic welfare subject, then the view here would merely lose the bonus of being able to explain why sentient things are welfare subjects.

6.3 The autopoietic standard rightly rejects the necessity of sentience

Plants are archetypal autopoietic entities, but are apparently insentient. The phenomena canvassed above seem to show that things matter to plants in a way that indicates that they have a good of their own. So sentience is not necessary for welfare subjecthood. In Section 6.4, I argue that another type of insentient welfare subject is possible. In Section 8, I introduce yet more insentient welfare subjects. If I am correct about these cases, then these are further, independent reasons to believe that sentience is not necessary for welfare subjecthood.

6.4 Sophisticated machines of the future will be welfare subjects

Consider a very sophisticated self-driving automobile of the not-too-distant future, Tessler. When Tessler's charge is running low, it drives itself to a charging unit and plugs itself in. While it is on its way, if the ambient temperature warrants, it warms the battery for faster charging. It monitors the condition of its brake fluid and gives itself a fluid change from a spare tank when necessary. Then it drains the waste tank and refills the spare tank at a service station where it routinely stops to fill its own windshield washer fluid and adjust its tire pressure according to the predicted outdoor temperature. Although Tessler can take care of itself in these and other ways in response to its environmental conditions, it sometimes has worn or broken parts that it cannot replace itself. In such a case, it drives itself to a service station and communicates its needs (or if it is disabled from moving, calls itself a tow truck). Of course, in addition to these new features, Tessler has all of the safety and control features of new cars today to avoid collision or prevent itself from leaving the roadway, such as radar-guided collision avoidance with automatic steering and braking, traction control, and anti-lock brakes. Assume that Tessler is not sentient. Still, with these or other easily imagined features, Tessler meets the conditions for being autopoietic: it sustains, maintains, and renews itself on an ongoing basis in response to its environmental conditions. It is therefore a welfare subject on my proposed standard.

One might turn this observation into an objection: To say that Tessler is a welfare subject with a good of its own stretches the concept of a welfare subject too far. Tessler is a mere machine. It is not a welfare subject. It is still the case that when we say that something is good for Tessler, this means that it helps preserve Tessler's condition in order that it can serve the purposes for which it was designed by and for

https://doi.org/10.1017/apa.2025.10006 Published online by Cambridge University Press

humans. It has no good of its own. Its good is determined only with reference to human ends.

Reply: Imagine all humans perished and left thousands of Tesslers. Things would still be good or bad for Tesslers, even though they were no longer able to serve the human ends for which they were designed. It would be bad for them to drive off of cliffs and become damaged beyond repair, and good for them to have their sensing and steering systems in good order to avoid driving off of cliffs. They would continue to take care of themselves and reveal what is good for them, say, by taking monthly test drives to be sure all their systems were functional and taking corrective action to the extent that they could without assistance. If their machine learning algorithms were sophisticated enough, they might even be able to figure out how to take care of themselves in ways in which they previously relied upon humans and other machines.

This reasoning will not convince those with strong intuitions against the welfare subjecthood of machines. I believe that some of the intuitive oomph against the idea that machines could be welfare subjects comes from the background assumption that we always have moral reasons to advance the welfare of welfare subjects, or at least not to harm them. If I have a welfare subject like a dog and I want to smash him up into pieces, there certainly is a moral reason for me not to do so. If I have a mere machine like a Tessler and I want to smash it up into pieces, there is no direct moral reason for me not to do so. So, it looks like machines cannot be welfare subjects. Let us turn in the next section to consider this alleged link between welfare subjecthood and moral reasons.

7. Objection: Absurd Implications

7.1 Welfare facts give rise to moral reasons

Connie Rosati argues that welfare subjecthood gives rise to reasons:

> Because it concerns what benefits a valuable being, as opposed to what furthers an end that may be of no particular interest or value, or what keeps in good condition a use value we may have no reason to use, welfarist good-for facts give pro tanto reasons to anyone, though these may often be reasons of noninterference rather than reasons of direct promotion. … [T]he fact that something is good for a welfare subject seems to have an automatic connection with reasons and so to be normative in a way that the fact that something is good for a car is not. (2009: 226–227)[5]

[5] Rosati repeats these claims in her (2020: 242-243, 245, 254, *et passim*). Notice that Rosati's main claim is that *all* welfare value gives rise to moral reasons. This is distinct from the familiar welfarist claim that *only* welfare value gives rise to moral reasons. (See Bramble (2020), Dorsey (2016), Keller (2009), Moore and Crisp (1996), and Sen (1979) for entries into the literature on welfarism.) The welfarist commitment to the necessary presence of welfare value in the presence of moral reasons obviously does not entail a commitment to the sufficiency of welfare value for generating moral reasons. That is just a point about the logical relationship between Rosati's claim and the familiar welfarist claim.

https://doi.org/10.1017/apa.2025.10006 Published online by Cambridge University Press

The objection is that if Rosati is right that we have reasons to promote or at least not to interfere with the realization of welfare value, then the view advanced here about welfare subjecthood has absurd implications. The absurd implications are that we have moral reasons to protect and promote the welfare of plants and moral reasons to protect and promote the welfare of machines of the near future, just as we have moral reasons to protect and promote the welfare of humans and non-human animals.

To see the alleged absurdity, suppose I am going away on a two-week vacation. Since my dog, Curly, is a welfare subject, I have a reason, for example, to find someone to come exercise him and feed and water him. If I leave Curly to die of thirst instead, then I do something wrong. Suppose now that instead of a dog, I have a houseplant, Petunia. If Petunia is a welfare subject, then I have a reason, for example, to find someone to come water Petunia. If I leave Petunia to die of thirst instead, then I do something wrong.[6] If it is wrong to let one houseplant die, then imagine how wrong it is to cut down a few trees for recreational fires in a fireplace or the moral horror involved in mowing one's lawn. But that is absurd. I do not do anything wrong to plants by letting my houseplant die or by mowing my lawn.

Similar absurdities can be generated by considering machine welfare. We will be doing something wrong to our Tesslers if we drive them hard like race cars and cause premature wear that we do not fix, or if we park them outdoors indefinitely, allowing their batteries to go dead and their bodies and components to rust away.

My response to this objection has two independent parts. In the first part, I challenge the claim at the heart of the objection that all welfare facts give rise to moral reasons. In the second part, I argue that even if all welfare facts do give rise to moral reasons, the reasons deriving from plant or machine welfare are too weak to generate the absurdity claimed in the objection.

7.2 Reply, Part One

Consider an argument to doubt that all welfare value gives rise to moral reasons. Notice with respect to some sorts of value other than welfare value, that this other value can be present without there being any direct agent-neutral reasons to promote it or not interfere with its realization. For example, I can acknowledge the aesthetic value of an art installation without having any reason at all to become a contributor. I claim that the moral reasons I admittedly do have not to interfere with the realization of that value derive not from that value, but from duties I have toward the persons interested in its realization. Yet it is plausible to think that if I were the last person standing, there would be no reason whatsoever for me not to destroy that value, for instance by using the canvases to tarp my firewood.

It is plausibly the same way with welfare value. On many theories of human welfare, it is good for a bank robber to pull off a successful heist and good for a villain to succeed at her villainy. It is plausible to think that there is not even a pro tanto reason for me to promote the heist or villainy. As well, it is not at all obvious that

[6] Thanks to Connie Rosati for an example of this form.

https://doi.org/10.1017/apa.2025.10006 Published online by Cambridge University Press

there is a reason not to interfere with the robber or villain getting what is good for them. Indeed, there is probably a reason in favor of interfering with the realization of their welfare value, at least if one can interfere at little cost to oneself. So, it appears that sometimes there is not even a pro tanto reason to promote or abstain from interfering with the realization of welfare value. If that is right, then the mere fact that my houseplant or Tessler is a welfare subject does not entail all by itself that I have a pro tanto reason to promote that welfare or that I do anything wrong in letting it die of thirst or rust apart.

7.3 Reply, Part Two

Part One takes a position on a debate about reasons to promote an intrinsic value that is not settled. Additionally, I do not claim to have done justice to Rosati's (2009) argument for the distinctiveness of the welfarist good-for relation with this quick analogy between aesthetic value and welfare value. So, suppose my argument in *Part One* fails. Suppose, instead, that it is true that facts about welfare do give rise to moral reasons in the way described at the outset of this section. I shall now argue that even if we grant that all welfare value gives rise to moral reasons, the allegedly absurd implication is no absurdity at all. We can concede that all facts about welfare give rise to pro tanto reasons, but recognize that some of these reasons are quite weak—certainly too weak to derive any absurd implications from their existence.

The moles in my lawn are welfare subjects. Driving my lawn tractor over their burrows while mowing interferes with their pursuit of their welfare. Assuming I have a moral reason not to interfere with their pursuit of their welfare, this reason is a fairly weak one. It does not give rise to a duty, for instance, to stop mowing my lawn and let it turn to brambles or to capture the moles and move them to a place where my lawn mowing will not interfere with their pursuit of their welfare. While the moral reason not to interfere with the moles' pursuit of their welfare by mowing my lawn is very weak, I have a stronger reason not to run over a scurrying mole in my path when I can easily avoid it by applying the brake. And while I have a moral reason to brake for the scurrying mole, the reason I have to brake for a scurrying toddler in my path is stronger still, by many orders of magnitude. The point is the obvious one that moral reasons come in strengths.

Therefore, even if moral agents always have agent-neutral pro tanto reasons to promote or not to interfere with the realization of welfare value, the strength of our reasons to promote or not to interfere with the welfare of plants and machines may be weak, extremely weak. This would explain why I have a very strong reason not to leave my dog to die of thirst but a quite weak moral reason not to leave my houseplant to die and an even weaker reason not to let my Tessler rust.

It should not be surprising that moral reasons based on welfare can be more or less weighty. Indeed, someone who thinks that sentience is necessary and sufficient for welfare subjecthood is already committed to this, as seen by the examples of the sentient mole and toddler. Furthermore, it is no objection to the view that all sentient beings are welfare subjects to point out that this implies we have moral reasons to promote or at least not to interfere with the welfare of billions of wild animals. My interlocutor could admit that we have those reasons, but claim that they are weak, so

https://doi.org/10.1017/apa.2025.10006 Published online by Cambridge University Press

no absurdity about what we have overall reason to do results. I have similarly claimed that even if we have moral reasons to promote or not interfere with the welfare of plants and sophisticated machines, these reasons could be very weak, so no absurdity results.

7.4 A Further Challenge

Rosati (2020: 255–257) acknowledges room for a wide-scope theory of welfare subjects according to which *everything* is a welfare subject. She explains the challenge of providing an account of the strength of reasons on such an account. So presumably she would put the challenge to me to explain *why* the reasons to promote the welfare of humans, non-human animals, plants, and machines differ in their strength when those beings are all equally welfare subjects on my medium-scope theory. Luckily, Gwen Bradford has offered a solution to this problem, which I co-opt here.

In preparation for understanding Bradford's view, distinguish between *affective consciousness* and *phenomenal consciousness*. *Affective consciousness* is essentially *sentience* as I have been using it here: the ability to experience states as positively or negatively valenced. Borrowing a characterization from David Chalmers, a being is *phenomenally conscious* if there is something that it is like to be that being (2020: 14:00) (Cf. also Block 1995: 230). Phenomenally conscious beings have experiences of which they are aware, such as of sounds, sights, and thoughts.[7]

I have argued explicitly that affective consciousness is not necessary for welfare subjecthood. Furthermore, since trees and Tesslers are not phenomenally conscious but are welfare subjects, my view entails also that phenomenal consciousness is not necessary for welfare subjecthood. Bradford (2023) argues explicitly that phenomenal consciousness is not necessary for welfare subjecthood. On Bradford's view, like mine, sophisticated machines of the near future, "some plants, and possibly other surprising things" may be welfare subjects (2023: 913). She considers a robot that is not phenomenally conscious and not affectively conscious, but that can be the subject of "plausible putative extra-experiential [welfare] goods" (2023: 907) including knowledge and satisfied desires:

> If the robot's knowledge and satisfied desires are good, we have pro tanto reason to ensure that all the robots of the world have as much

[7] Chalmers considers a phenomenally conscious being that is not affectively conscious, which he calls a *(Philosophical) Vulcan* (2020: 33:30; 2022: 343). He argues that Vulcans have moral status, and concludes partly on that basis that phenomenal consciousness is sufficient for moral status and affective consciousness is not necessary (2020: 40:30; 2022: 344). While Chalmers is primarily concerned with moral status, I have been concerned in the first place with welfare subjecthood. In the text, I am in the midst of considering a possible relationship between moral status and welfare subjecthood. Chalmers indicates sympathy with the view that phenomenal consciousness is necessary for welfare subjecthood (2020: 51:30; 2022: 329-330, 342), but that affective consciousness is not (2020: 51:30). For more on the relationship between welfare subjecthood and moral status, see Roelofs (2023) and Delon (unpublished). van der Deijl (2021: 192-193) also discusses the relationship and argues that phenomenal consciousness (which he unfortunately calls "sentience") is necessary and sufficient for welfare subjecthood, and sufficient for moral status.

https://doi.org/10.1017/apa.2025.10006 Published online by Cambridge University Press

> knowledge as possible and everything they want. Further, because the current welfare level of robots is presumably lower than the average human, egalitarian considerations give reason to devote resources to improving the welfare of the world's robots. (2023: 909)

Bradford concedes—as I do for the sake of argument here—that we have pro tanto reasons to advance robot welfare but claims—as I do—that these reasons are very weak. Finally, to the point and my sole purpose in introducing Bradford's work here: she provides a good explanation for *why* those reasons are relatively weak. She claims that the value of nonconscious welfare states—such as unconscious knowledge, desires satisfied outside of experience, and posthumous achievements—is very small in general. For example, if I strive my whole life for the company I built to reach a value of $10 million, and it does but only after my death, then the value of that achievement is tiny in comparison to what it would have been if the goal had been reached during my lifetime. Since the value of nonconscious welfare states is small, our reasons to promote those states are correspondingly weak (2023: 913–914).

Since phenomenal consciousness of welfare states by welfare subjects makes all the difference to the strength of reasons to promote those welfare states, our reasons to promote the welfare of welfare subjects all of whose welfare states are *nonconscious* are extremely weak. Bradford's robots are not phenomenally conscious. *Therefore*, the egalitarian reasons we have to advance their welfare rather than human welfare are extremely weak. Petunia is not phenomenally conscious, but Curly is. *Therefore*, the reasons I have to hire a caretaker for Petunia are extremely weak in comparison to my reasons to have someone take care of Curly. Thus, Bradford's explanation of the variable strengths of the reasons to advance the welfare of different sorts of welfare subjects solves the problem for my account with which this subsection began, which was to explain why the strengths of those reasons can differ on my account.

Of course, it would be nice if we had a further explanation for *why* nonconscious welfare states have relatively little value. Bradford (2023: 914–918) has an account of this as well. In brief, it is that a welfare good plus phenomenal consciousness of that welfare good is an organic unity with much more value than either the welfare good or phenomenal consciousness alone.

8. Objection: Alternative Explanation of the Data

Recall my comments at the end of Section 4 explaining my abductive strategy in this article, in effect claiming that autopoiesis accounts for the data of what entities are and what entities are not welfare subjects better than sentience or teleological organization accounts for the data. One might raise an alternative explanation of the data. This explanation is that, in the case of sentient organisms that are capable of taking an interest in things, autopoietic processes are *instrumentally good* (provided that they cause, produce, or lead to, e.g., episodes of pleasure). Still, it is only sentient organisms that are welfare subjects, and it is only states like pleasure that are *non-instrumentally good*. This explanation grants that autopoietic processes are valuable *indicators* of welfare value, but sticks to the idea that sentience is the real ground of

https://doi.org/10.1017/apa.2025.10006 Published online by Cambridge University Press

welfare subjecthood, tied as it is to non-instrumental goods (i.e., only sentient entities can experience pleasure or pain).

In response, I shall introduce a new reason to believe that autopoiesis and not sentience grounds welfare subjecthood. The new reason is that autopoiesis is present in some welfare subjects that are not actually sentient but will be sentient in the future, or that were sentient in the past but are not now, or that will never be sentient but would have been if things had been otherwise. After the following explanation, we will see that the hypothesis that autopoiesis grounds welfare value does a better job of explaining the data in such cases than the hypothesis that sentience grounds welfare value.

Consider some property P (e.g., sentience, having a rational nature, etc.) that is proposed as grounding membership in some normative category C (e.g., being a welfare subject, being morally considerable, etc.). If we run across something that barely has P but intuitively is a member of C, we can merely clarify the theory (or revise it slightly) to say that P *at any level* grounds membership in C. The trouble arises when something that intuitively ought to be a member of C does not actually have P to any extent, but has P only on a hypothetical level.

Consider:

(A) a being that intuitively is C and does not have P, but is expected to develop P (e.g., a fetus about to be born);
(B) a being that intuitively is C, does not have P, but had it at an earlier time (e.g., someone in a coma or with a severe brain injury);
(C) a being that intuitively is C, does not have P and will never develop P, but would have had P had their developmental history been typical of their kind (e.g., someone born with a severe neurological disorder).

What do we say of such cases? People who want to preserve P as a necessary condition can revise the view to say that for S to be C, it is necessary that S meet *one* of the following conditions:

(1) S actually has P (at least to some extent);

or:

(2) S has P, but only hypothetically, in one of the following senses:
  (A) S does not have P, but is expected to develop P;
  (B) S does not have P, but once had P;
  (C) S does not have P, but would have had P, had S developed typically.

The upshot is: what role, then, does P actually play in determining whether something is C? Apparently actually having P does not play the role since something can completely lack P but be in C. So there must be something *else* unifying the things in classes (1) and (2). P is not doing the unifying or sorting it

https://doi.org/10.1017/apa.2025.10006 Published online by Cambridge University Press

was intended to do. There must be some other property common to the things in classes (1) and (2) that accounts for their C.

The suggestion was that sentience even at a hypothetical level might ground welfare subjecthood. If, in the above, P is sentience and C is welfare subjecthood, then we see that sentience is not doing the unifying or sorting it was intended to do. Thus, to accept sentience at a hypothetical level as grounding welfare subjecthood is to abandon the sentience view, i.e., to capitulate. So the suggestion that sentience at a hypothetical level might ground welfare value is not going to help the sentience view. It is going to help my view. My suggestion, of course, is that the property that accounts for the welfare subjecthood of things that have sentience only hypothetically is autopoiesis.

9. New Questions for Well-Being Researchers

If I am right that all autopoietic entities are welfare subjects, then many new questions and lines of research present themselves to philosophers of well-being: What is the proper relationship between theories of well-being for humans and theories of well-being for insentient autopoietic entities? Do we need brand new theories of well-being for insentient autopoietic entities, or can theories of human well-being be adapted or generalized to apply to all autopoietic entities? Does opening the class of welfare subjects to all autopoietic entities have implications for invariabilism, the view that the same theory of welfare is true of every welfare subject (Lin 2018)? Do we need discussions about the proper relationships between philosophy of well-being (on the one hand) and biology and engineering (on the other hand), similar to the discussions about the proper relationship between the philosophy of human well-being and the psychology of human well-being (Alexandrova 2015, Tiberius 2013, Raibley 2012, Feldman 2010, Angner 2009, and Haybron 2008) and about the proper relationship between the philosophy of animal well-being and animal welfare science (Bruckner 2019, 2020)? While these are enormous questions, it is a benefit of the view advanced here that it forces us to confront these questions and move toward a fuller understanding of well-being wherever it is present.

DONALD W. BRUCKNER
PENN STATE UNIVERSITY, NEW KENSINGTON
donald.bruckner@gmail.com

Acknowledgements. I am grateful to Chauncey Maher, Gwen Bradford, and anonymous referees for written feedback. Thanks also to my commentators at a 2023 Eastern APA symposium, Devin Curry and Zach Vereb; my commentator at the 2022 Rocky Mountain Ethics Congress, Nolan Cannon; and to the audiences at both.

Competing interests. The author declares none.

References

Alexandrova, Anna. (2015) 'Well-Being'. In Nancy Cartwright and Eleanora Montuschi (eds), *Philosophy of Social Science: A New Introduction* (Oxford: Oxford University Press), pp. 9–30.

https://doi.org/10.1017/apa.2025.10006 Published online by Cambridge University Press

Angner, Erik. (2009) 'Subjective Measures of Well-Being: Philosophical Perspectives'. In Harold Kincaid and Don Ross (eds), *The Oxford Handbook of Philosophy of Economics* (Oxford: Oxford University Press), pp. 560–579.
Basl, John, and Sandler, Ronald. (2013) 'The Good of Non-Sentient Entities: Organisms, Artifacts, and Synthetic Biology'. *Studies in History and Philosophy of the Biological and Biomedical Sciences*, 44(4 Pt B), 697–705.
Block, Ned. (1995) 'On a Confusion about a Function of Consciousness'. *Behavioral and Brain Sciences*, 18(2), 227–247.
Bradford, Gwen. (2023) 'Consciousness and Welfare Subjectivity'. *Noûs*, 57(4), 905–921.
Bradley, Ben. (2015) *Well-Being*. Cambridge, UK: Polity.
Bramble, Ben. (2020) 'Welfarism'. In Hugh LaFollette (ed.), *The International Encyclopedia of Ethics* (Malden, MA: Wiley-Blackwell). https://doi.org/10.1002/9781444367072.wbiee940
Bruckner, Donald W. (2019) 'Philosophy and Animal Welfare Science'. *Philosophy Compass*, 14(10), 1–12.
Bruckner, Donald W. (2020) 'Animal Welfare Science, Varieties of Value, and Philosophical Methodology'. *Animal Welfare*, 29(4), 387–397.
Chalmers, David. (2020) 'Consciousness and Moral Status'. Delivered at City College of New York. https://selfawarepatterns.com/2020/11/20/consciousness-and-moral-status/
Chalmers, David. (2022) *Reality+: Virtual Worlds and the Problems of Philosophy*. New York: W. W. Norton.
Delon, Nicolas. Unpublished. 'Agential Value'.
Dorsey, Dale. (2016) 'Welfarism'. In Guy Fletcher (ed.), *The Routledge Handbook of Philosophy of Well-Being* (New York: Routledge), pp. 417–428.
Feldman, Fred. (2010) 'On the Philosophical Implications of Empirical Research on Happiness'. *Social Research*, 77(2), 625–658.
Haybron, Dan. (2008) 'Philosophy and the Science of Subjective Well-Being'. In Michael Eid and Randy J. Larsen (eds), *The Science of Subjective Well-Being* (New York: Guilford), pp. 17–43.
Hettinger, Ned. [1998] (2012) 'Comments on Holmes Rolston's "Naturalizing Values"'. In Louis P. Pojman and Paul Pojman (eds), *Environmental Ethics: Readings in Theory and Application (*6th *ed)* (Boston: Wadsworth), pp. 119–122.
Keller, Simon. (2009) 'Welfarism'. *Philosophy Compass*, 4(1), 82–95.
Kraut, Richard. (2007) *What is Good and Why: The Ethics of Well-Being*. Cambridge, MA: Harvard University Press.
Lin, Eden. (2018) 'Welfare Invariabilism'. *Ethics*, 128(2), 320–345.
Maher, Chauncey. (2017) *Plant Minds: A Philosophical Defense*. New York: Routledge.
Maturana, Humberto R. and Varela, Francisco J. (1973) *De máquinas y seres vivos: Una teoría de la organization biológica*. Santiago, Chile: Editorial Universitaria.
Maturana, Humberto R. and Varela, Francisco J. (1980) *Autopoiesis and Cognition: The Realization of the Living*. Dordrecht: D. Reidel.
Moore, Andrew and Crisp, Roger. (1996) 'Welfarism in Moral Theory'. *Australasian Journal of Philosophy*, 74(4), 598–613.
Raibley, Jason R. (2012) 'Happiness is Not Well-Being'. *Journal of Happiness Studies*, 13(6), 1105–1129.
Railton, Peter. (1986) 'Fact and Values'. *Philosophical Topics*, 14(2), 5–31.
Rawls, John. (1971) *A Theory of Justice*. Cambridge, MA: Harvard University Press.
Roelofs, Luke. (2023) 'Sentientism, Motivation, and Philosophical Vulcans'. *Pacific Philosophical Quarterly*, 104(2), 301–323.
Rolston, Holmes. [1998] (2012) 'Naturalizing Values: Organisms and Species'. In Louis P. Pojman and Paul Pojman (eds), *Environmental Ethics: Readings in Theory and Application (*6th *ed)* (Boston: Wadsworth), pp. 105–118.
Rosati, Connie. (2009) 'Relational Good and the Multiplicity Problem'. *Philosophical Issues*, 19(1), 205–234.
Rosati, Connie. (2020) 'Welfare and Rational Fit'. *Oxford Studies in Metaethics*, 15, 241–262.
Sen, Amartya K. (1979) 'Welfarism and Utilitarianism'. *The Journal of Philosophy*, 76(9), 463–489.

Singer, Peter. (1975) *Animal Liberation: A New Ethics for Our Treatment of Animals*. New York: Random House.
Taylor, Paul W. (1981) 'The Ethics of Respect for Nature'. *Environmental Ethics*, 3(3), 197–218.
Thompson, Evan. (2007) *Mind in Life: Biology, Phenomenology, and the Sciences of Mind*. Cambridge, MA: Harvard University Press.
Tiberius, Valerie. (2013) 'Well-Being, Wisdom, and Thick Theorizing: On the Division of Labor between Moral Philosophy and Positive Psychology'. In Simon Kirchin (ed.), *Thick Concepts* (Oxford: Oxford University Press), pp. 217–233.
van der Deijl, Willem. (2021) 'The Sentience Argument for Experientialism about Welfare'. *Philosophical Studies*, 178(1), 187–208.
Varela, Francisco J., Maturana, Humberto R., and Uribe, R. (1974) 'Autopoiesis: The Organization of Living Systems, its Characterization and a Model'. *BioSystems*, 5(4), 187–196.
Varner, Gary. (1998) *In Nature's Interests? Interests, Animal Rights, and Environmental Ethics*. Oxford: Oxford University Press.

https://doi.org/10.1017/apa.2025.10006 Published online by Cambridge University Press