

Personites, Plenitude, and Intrinsicity

Cian Dorr and John Hawthorne

Forthcoming in *The Importance of Being Conscious*, ed. Geoff Lee and Adam Pautz (Oxford University Press)

Draft of 17th October 2024

Consult the authors if you wish to cite.

In two thought-provoking papers, Mark Johnston (2016, 2017) has argued on ethical grounds against a wide family of “naturalistic” metaphysical world-views. The common feature of the target views is that they imply that

we are “ontological trash”, i.e. that in our close spatial vicinity there are many persisting things all ontologically on a par, very similar in their features and such that they come into being and cease to exist at various times. (Johnston 2016: 216)

In particular, the target views imply the existence of a vast abundance of “personites”: ‘shorter-lived very person-like things that extend across part, but not the whole, of a person’s life’ (Johnston 2016: 199). Johnston argues that, because of the ontological parity between persons and personites, the target views saddle us with the conclusion that all these personites have *moral status* in much the same way that people do. They have claim-rights, deserve moral consideration, and so on.¹ Having established this to his satisfaction, he goes on to argue that the attribution of such moral status to personites has radically revisionary — indeed, absurd — consequences for questions about what one morally ought to do. For example, he suggests that if one accepted the conclusion, one would have to say that it is morally

¹Johnston has his own favoured way of cashing out the expression ‘moral status’ (Johnston 2017: 621). But since we won’t be discussing Johnston’s claims about what would follow from the attribution of moral status to personites, it won’t matter for our purposes exactly how the expression is precisified.

wrong to take on any task that is gruelling in the short term, since that will involve signing up a host of personites to an unpleasant short-term future that won't be compensated for by any long-term benefits to them, since they won't be around to enjoy such benefits.²

In this paper, we will respond to the first phase of Johnston's argument. We will show how to reconcile the well-motivated metaphysical claim that people are *ontologically* on a par with a multitude of objects in their vicinity with the well-motivated ethical claim that people are *morally* radically different from any such objects.

1 Plenitude

Johnston's target views include widely accepted forms of physicalism, on which physics not only provides a supervenience base for everything else, but is moreover a good guide to relative *naturalness* or *fundamentality*, and thus to which things are "ontologically on a par".³ He holds—correctly, in our view—that physicalism so understood implies that people are "ontological trash", in the sense explained above. Consider, for example, a Sorites sequence of possible situations, connecting an innocuous operation which clearly involves the continued existence of a single person to a drastic operation which clearly involves one person ceasing to exist and being replaced by another (Parfit 1986). Such a sequence need not involve any sharp discontinuities at the level of physics. Nevertheless, by classical logic, there is a first situation after which we have replacement rather than survival of a person. To reconcile the discontinuity as regards the facts about *people* with the lack

²Johnston had already given a brief presentation of the personite worry in (Johnston 2010: 64–5); and several of the main moves are already present in Zimmerman 2003 (§3.2). Olson (2010) gives a detailed presentation of an argument quite close to Johnston's, including similar thoughts about how attributing moral status of personites ('subpeople') would be morally catastrophic. However, Olson simply assumes that the view he is targeting entails that personites have mental states, and does not make the argument from intrinsicity which will be our central focus. Olson does however give an interesting "metasemantic" argument for the claim that personites have moral status, which does not have any close analogue in Johnston's discussion. Our discussion in section 9 will serve, in part, as a response to this argument of Olson's (see note 60).

³We understand the claim that two objects are "ontologically on a par" to mean not only that they are of the same ontological *category*, but that they are roughly equally *natural*. For a general discussion of the relevant notion of metaphysical naturalness, see Dorr and Hawthorne 2013 and Dorr 2019. We assume that the degrees of naturalness of *objects* are proportional to the degrees of naturalness of their haecceities (where the haecceity of *x* is the property *being identical to x*), so we could if we wished avoid ever applying 'natural' to objects.

of any ontologically distinguished lines, one needs to posit a lot of other objects, ontologically on a par with people, but differing from them as regards the transition point from survival to replacement. That way, the transition point where the original *person* stops surviving will not look special from a point of view that takes all of the relevant objects into account.

This kind of argument also applies to many non-physicalist views. For example, even if one is persuaded by arguments that the physicists' roster of fundamental properties is incomplete and needs to be supplemented by some *phenomenal* properties, one might still find it implausible that there are bright lines in the Parfitian Sorites series. There is no forced march from the existence of fundamental phenomenal properties to such bright lines: for example, the fundamental phenomenal properties could be borne by *events* rather than by any one of the multitude of ontologically person-like objects.

You might be puzzled by the claim that only one of the multitude of ontologically very similar objects in the vicinity of a given person (say, Mark Johnston) is a person. Wouldn't the truth of this claim require some kind of incredible "semantic miracle" (Unger 1979), whereby our use of 'person' (and of proper names and personal pronouns) achieves exquisite sensitivity to metaphysically unimportant differences? But if there was a good argument here against the claim that only one of the multitude is a person, there would be an equally good argument against the claim that only one of the multitude is *identical to Mark Johnston*—after all, our use of the predicate 'identical to Mark Johnston' is not especially more sensitive to small differences than our use of 'person'. But the claim that only one thing is identical to Mark Johnston must be accepted by anyone who has made their peace with classical logic. And once one has got used to the idea that the property of *being Mark Johnston* manages to distinguish just one member of a cluster of metaphysically similar objects, there is no further obstacle to preserving the highly intuitive, though not logically compulsory, judgment that the property of *being a person* is similarly discriminating. Of course, given the multiplicity of candidates, it is plausible that the predicates 'person' and 'identical to Mark Johnston', as well as the proper name 'Mark Johnston', are *vague*. But on most accounts of vagueness, acknowledging this vagueness is perfectly compatible with acceptance of the relevant uniqueness claims.⁴

Although Johnston's target is thus very broad, he is occupied much of the time with a narrower target he calls "four-dimensionalism", closely based on

⁴Lewis (1993) develops a view of this sort in the idiom of supervaluationism, though he ends up favouring a hybrid view on which 'person' also has legitimate interpretations that apply to many objects in the relevant chair. See Dorr, Hawthorne, and Yli-Vakkuri 2021 (ch. 3) for arguments against 'many chairs'/'many people' approaches.

the views of Lewis (1976). This view has two important features which go beyond the generic view that persons are ontological trash.

- (i) It denies that any pair of distinct material objects perfectly coincide at all times at which either of them concretely exists.
- (ii) It affirms that all, or at least many, of the shorter-lived objects that *temporarily* coincide with a given person are the subjects of mental states, and indeed that they have (at the relevant time) a rich portfolio of psychological features similar to those of the person with whom they coincide.

We agree that Johnston's ethical arguments constitute a very serious problem for views that accept either of these further claims. But we think that both claims are missteps, which should be avoided even on the assumption that people are "ontological trash".

There are well known apparent counterexamples to (i)—such as Gibbard's (1975) example, where a statue permanently coincides with a piece of bronze while differing from it in its modal properties—which by our lights look as compelling as the reasons for allowing that distinct objects sometimes *temporarily* coincide.⁵ Indeed, we favour a principle of 'plenitude' which every material object permanently coincides with *infinitely* many other material objects, which between them instantiate a vast variety of modal profiles, each specifiable by a partial function from worlds to spacetime regions. See Dorr, Hawthorne, and Yli-Vakkuri 2021 (ch. 11) for several precise versions of such a principle, along with argument that plenitude provides the best way of reconciling ordinary modal judgments with an underlying metaphysics that provides no ontologically distinguished bright lines that conveniently track those judgments. Here, we will simply note that if plenitude provides a way of escaping from Johnston's arguments, that will constitute a further consideration in its favour.

As regards (ii), it seems to us that the reasons for thinking that it rarely happens that multiple *subjects of consciousness* are sitting in the same chair are roughly as strong as the reasons for thinking that it rarely happens that

⁵Lewis of course has his own favoured way of accommodating the relevant modal judgments without allowing for distinct objects that permanently coincide, by appealing to systematic context-sensitivity in modal vocabulary (Lewis 1971, 1986: §4.5). See Fine 2003 and Dorr, Hawthorne, and Yli-Vakkuri 2021 (ch. 10) for arguments against this approach. We note in passing that some of Johnston's arguments, such as the "adventitiousness" arguments we discuss in section 2, invoke *de re* modality in ways that Lewis would diagnose as trading on context-shifts.

multiple *people* are sitting in the same chair. We have already explained why the claim that people are ontological trash does not undermine the latter judgment; the same considerations extend smoothly to the former judgment. Unless some special argument is forthcoming, one should stick with the pre-theoretically plausible judgment that the only conscious being that overlaps any ordinary person is that very person.⁶ Indeed, since a view that multiplies consciousness without multiplying moral status is *prima facie* extremely unpromising, it is vital not to just assume without argument that the target of Johnston's argument will accept that any of the objects overlapping a typical person is a subject of consciousness or any other mental state.

Fortunately, Johnston *does* have some arguments that support a multiplication of conscious beings, on the assumption that people are ontological trash. These arguments are structurally exactly parallel to his main arguments for a multiplication of beings with moral status. The objections we will be making below to these arguments about moral status also apply to the corresponding arguments about consciousness.⁷

2 Intrinsicity

Johnston's most fully-developed argument that the target naturalistic view implies that personites have moral status turns on the concept of *intrinsicity*: specifically, on the claim that given the target view, 'each personite is intrinsically just like some possible person or other' (Johnston 2016: 200). We will spend the next six sections on a careful consideration of this argument. In section 8, we will discuss another argumentative strategy that is suggested by certain of Johnston's remarks. Finally, section 9 will turn to some "metasemantic" arguments which have less of a foothold in Johnston's discussion. Ultimately, we don't think any of the arguments work to undermine the view that people are "ontological trash", but as we will see, they

⁶The anti-physicalist versions of the target view do not provide any particular reason to give this up: one can say that each of the events that instantiate the fundamental phenomenal properties has a unique subject while also saying that 'subject of' is vague in a way that is co-ordinated with the vagueness of 'person' and 'conscious being'.

⁷Johnston (2016: 205) also mentions another argument for the attribution of psychological states to personites that depends on Lewis's (1976) analysis of 'person' as 'maximal sum of *R*-interrelated stages', where being '*R*-related' is (at least in part) a matter of 'mental continuity and connectedness'. As Johnston points out, this seems to presuppose that the *R*-related stages themselves have mental properties. But Lewis's analysis is in any case out of keeping with the view that each person permanently coincides with many non-people, which follows from the combination of plenitude and the claim that no two people permanently coincide.

do yield important insights about ethics, metasemantics, and the nature of intrinsicality.

Here is Johnston's most explicit articulation of the crucial claim about intrinsicality:

For all y, x, t and t' where y is a person and x is one of that person's personites beginning at t and ending at t' , x 's beginning and end in time in corresponds to the beginning and end of a possible person in some world w , one who is exactly similar in all respects intrinsic to x 's actual mental and physical life. (Johnston 2016: 204, premise 3)

For the time being, we will assume that 'exactly similar in all respects intrinsic to x 's mental and physical life' can be simplified to 'exactly similar in all intrinsic respects', or 'such as to have exactly the same intrinsic properties'; while the argument may go better if we treat the qualification as imposing a substantive restriction on the range of relevant intrinsic properties, it will be instructive to first consider the unrestricted version.⁸ Note that Johnston doesn't really need the claim that this holds for *all* personites: his strategies for deriving bizarre conclusions about moral permissibility only require that 'These personites, or at least many, many of them, are therefore sufficiently like persons to deserve respect, along with the other distinctive attitudes which we rightly extend to persons' (Johnston 2017: 620, emphasis ours). To avoid distracting issues about the definition of 'personite' that would be raised by the 'all' claim, we will focus on the 'many' claim, which we can state in modal terms as follows:

Compatibility Many personites have an intrinsic profile compatible with *being a person*.

Here, when F is any property of properties (such as intrinsicality), an F profile is any property of the form *being something whose F properties are all and*

⁸The invocation of *lives* introduces issues relating to the intrinsic properties of events and processes which we don't see much point getting into. We believe that Johnston's primary motivation for putting things in terms of lives is to signal that the operative notion of intrinsicality is an expansive one on which properties like *sometimes being spherical* and *being spherical one year after one's creation and cubical two years later*, and not merely "temporally local" properties like *being spherical*, count as intrinsic, in order to rebut objections to Intrinsicality that turn on the idea that an object's moral status is partly a matter of its history. Let's agree that talk of intrinsicality throughout this paper is to be understood in the expansive way.

only the members of C , for some collection of properties C .⁹

To get from this premise to the conclusion that all personites have moral status, Johnston uses two other premises. Reformulated to match our modal formulation of Compatibility, they amount to the following:

Universality Necessarily every person has moral status.¹⁰

Intrinsicity Everything that lacks moral status has an intrinsic profile that necessitates *lacking moral status*.¹¹

The three premises jointly imply that many personites have moral status. By Universality, *lacking moral status* entails *not being a person*, so by Intrinsicity, everything that lacks moral status has an intrinsic profile that entails *not being a person*. But Compatibility says that many personites have intrinsic profiles that do not entail this.

Johnston also presents two other, closely related arguments for a similar extension of moral status. The first of these arguments uses, in place of Intrinsicity, the premise that moral status is not an “adventitious” matter, in the sense that ‘a being’s moral status cannot turn on, does not partly supervene on, is not partly grounded in, what happens after it has ceased to exist’ (Johnston 2017: 621). The second argument strengthens this premise by appending ‘...or before it begins to exist’, thereby putatively securing moral status for an even greater variety of personites. Although there are some tricky choices to be made in precisely articulating these notions of adventitiousness, and some subtleties about their relationship with intrinsicity, we are confident that the remarks that follow can easily be adapted to these alternative arguments by simply reinterpreting ‘intrinsic’ throughout as expressing the negation of the relevant notion of adventitiousness.¹²

⁹This talk of “collections” can be formalized using the notion of a *rigid* property of properties see Dorr, Hawthorne, and Yli-Vakkuri 2021 (§1.5). Although for many values of F , F profiles are obviously not themselves F , it does seem plausible that intrinsic profiles are themselves intrinsic properties. In section 4 we will consider one argument that depends on this being the case, but our overall discussion does not depend on it.

¹⁰Cf. Johnston’s (2016: 203) premise 1.

¹¹Cf. Johnston’s (2016: 203-4) premise 2. We could derive Intrinsicity from the simpler premise that *lacking moral status* is itself an intrinsic property; later we will discuss some principles about intrinsicity which make Intrinsicity equivalent to this.

¹²We would suggest understanding adventitiousness as the negation of *future-insensitivity* or *temporal locality*, defined as follows:

F is temporally local [future-insensitive] := necessarily, for any [initial] temporal segment of spacetime r with intrinsic profile G , and any object x whose whole career is confined to r , the proposition that r is G and is a [n initial] temporal segment

3 Haecceities

While there is of course room for debate about the other two premises, it seems to us that the main problem with the intrinsicity argument lies with Compatibility. To evaluate this premise, we will need to get a bit clearer on what it means for a property to be intrinsic. Intrinsic properties are supposed to be those we have ‘purely in virtue of the way we are’, as opposed to ‘the way we interact with the world’ (Marshall and Weatherston 2023). The usual paradigms in the literature are shape properties and mass properties.¹³ But these examples plausibly have two further features as well as being intrinsic. First, they seem to be *qualitative*: intuitively, they do not “depend on the identity of” any specific individuals.¹⁴ Second, they seem to be *undiscriminating*: such that they could not be possessed by only one of a pair of coincident objects.¹⁵ However, insofar as we have any grip on

that wholly contains x necessitates either that x is F or that x is not F .

While dropping ‘that wholly contains x ’ arguably gives a more intuitive notion, this would make *having moral status* a future-sensitive (and hence temporally non-local) matter given Harman’s (1999) view that early fetuses that will later become children have moral status whereas early fetuses that will quickly perish do not. Since Johnston (2016: 202) wants to stay neutral as regards Harman’s view, the above precisification seems to better match his intention.

It is worth noting that given plenitude, certain standard paradigms of intrinsicity are not future-insensitive, and hence also not temporally local. For example, suppose a bronze sphere s and a bronze cube c were both destroyed in the eruption of Vesuvius, and consider an object x that, necessarily, coincides with s if the Allies win WWII and s concretely exists; coincides with c if the Allies lose WWII and c concretely exists; and otherwise does not concretely exist. Although x is spherical, the conjunction of complete truth about the intrinsic properties of the pre-eruption part of spacetime with the true proposition that x is wholly located within this region leaves it open whether the Allies win, and hence whether x is spherical. So sphericity is not future-insensitive. This is a surprising result, but given plenitude, we see no prospect of a plausible gloss on the intuitive idea of future-insensitivity that makes all intrinsic properties future-insensitive.

¹³These are not indisputable, however. Skow (2007) discusses reasons for thinking that many shape properties are not intrinsic. And on some views, having a particular determinate mass is extrinsic, consisting in being related in a certain way to some exemplar, or fitting in a certain way into some holistic pattern, or being “located” at certain point in a ‘mass space’ analogous to physical space (Arntzenius and Dorr 2011).

¹⁴These particular qualitiveness claims are not beyond dispute: given any of the views mentioned in the previous note where determinate masses are non-intrinsic, they are also arguably also non-qualitative.

¹⁵Again, this is not beyond dispute. Fine (2003) suggests that a loaf of bread and the bread that composes it may be “materially” coincident while failing to be spatially coincident (and thus presumably differing in shape). Meanwhile, *spatially* coincident objects can arguably differ in mass at possible worlds with interpenetrating matter.

the notion of intrinsicity, it seems clear that many intrinsic properties lack one or both of these features. Most obviously, the *haecceity* of any object—the property of being identical to that object—should surely count as an intrinsic property.¹⁶ But haecceities, in general, seem to be neither qualitative nor indiscriminating.¹⁷

The intrinsicity of haecceities is bad news for the thesis that all personites have intrinsic profiles compatible with personhood. For since any object's intrinsic profile entails its haecceity, that thesis entails that any personite is such that *it* could be a person (and indeed, such that it could be a person while instantiating exactly the same intrinsic properties). But the wide variety of objects generated by plenitude include some which would seem to qualify as personites, but obviously couldn't have been people at all. For example, consider an object Mark⁻ that coincides with some proper initial segment of Mark's life in the actual world, and doesn't concretely exist in any other world. Mark⁻ seems to qualify as a personite. But Mark⁻ isn't a person in the actual world, since Mark never coincides with any other person. And given that personhood entails concrete existence, Mark⁻ isn't a person in any other possible world either.¹⁸

What about the more cautious claim that every personite that could have permanently coincided with a person could have been a person? Since Plenitude generates an abundant supply of personites of this sort, the claim that all of them have intrinsic profiles compatible with personhood would be sufficient for the truth of Compatibility. But once we bear in mind that per-

¹⁶Francescotti (1999: 593) and Eddon (2011) defend the claim that haecceities are intrinsic. Denby (2014) rejects it. Sider (1996) maintains that there are two senses of 'intrinsic', one on which all haecceities are intrinsic and another on which all intrinsic properties are qualitative. While we are open to the view that metaphysicians have run together multiple senses of 'intrinsic', we are inclined to think all of these senses apply to haecceities. But in the next section we will discuss a modified argument that replaces 'intrinsic' with 'intrinsic and qualitative'; if there is an interpretation of 'intrinsic' that entails 'qualitative', our remarks about this modified argument should carry over to this interpretation of the original argument. Note that the claim that haecceities are *non-adventitious* is on even more secure footing than the claim that they are intrinsic—nothing at all, and *a fortiori* nothing that happens after a thing ceases to exist, can prevent a thing from instantiating its haecceity!

¹⁷We may need to allow that some unusual objects have qualitative haecceities. For example, plenitude implies that there is a unique object such that necessarily it occupies all and only the points occupied by some red object. If there is a *unique* such object, one might plausibly claim that its haecceity is identical to the qualitative property *being an object that necessarily occupies all and only the points occupied by some red object*.

¹⁸Fine (2005) denies that personhood entails concrete existence, but would still agree that Mark⁻ is necessarily not a person, since he affirms that non-personhood entails necessary non-personhood.

manent coincidence does not imply identity, it is far from clear that any of those personites could have *been* people, as opposed to merely permanently coinciding with them. For example, the intrinsic properties of a person's brain are arguably compatible with its always being located in a vat containing no other organic matter, and thus with its permanently coinciding with a person. But an envatted brain would still not have *been* a person, since the person who coincided with it would still have been such that they *could* have had it as just one of many organs. Similarly, consider an object Mark* such that necessarily, it concretely exists at exactly those times at which Mark concretely exists but has not yet visited Budapest, and coincides with Mark whenever it exists. At a possible world where everything is exactly as it actually is until the time when Mark actually first visits Budapest and Mark is then instantaneously vaporized, Mark* plausibly has all its actual intrinsic properties while coinciding with a person, namely Mark. But Mark* would still not have been a person at such a world. For *Mark* would still have been a person, and would still not have coincided with any people distinct from him. But he would have coincided with Mark*, and still been distinct from Mark* (since it would still have been possible for Mark to visit Budapest and impossible for Mark* to visit Budapest).

In fact, if we combine the intrinsicity of haecceities with the popular view that metaphysical necessity obeys the modal logic S5, we can plausibly argue for a claim diametrically opposed to Compatibility:

Incompatibility No non-person has an intrinsic profile compatible with personhood.

If this is true, then we are free, if we wish, to maintain that only people have moral status, while accepting both Universality and Intrinsicity. One widely accepted thesis that will immediately secure Incompatibility under these assumptions is the following:

Person Essentialism Necessarily, every person is such that necessarily, they are a person if and only if they concretely exist.

In S5, this entails that every concretely existing non-person is necessarily not a person. Suppose for contradiction that *x* is a concretely existing non-person that is possibly a person. Then by Person Essentialism, it is possibly necessary that if *x* concretely exists, *x* is a person. But S5 includes the 'B' axiom, according to which whatever is possibly necessarily the case is the case. So we can conclude that if *x* concretely exists, *x* is a person. This contradicts our assumption that *x* concretely exists and is not a person. Given

that haecceities are intrinsic and that non-concreteness (i.e., lack of concrete existence) is intrinsic, it follows that every non-person has an intrinsic property incompatible with personhood: either non-concreteness (if it does not concretely exist) or its haecceity (if it does).

Johnston sometimes seems to assume that his opponent will have to deny Person Essentialism, and hold instead that ordinary people could *very easily* have concretely existed without being people. But his reasons for this seem specific to the Lewisian “four-dimensionalist” view that rejects the possibility of permanent coincidence.¹⁹ Plenitudinous views like ours that embrace massive coincidence are generally assumed to be quite friendly to essentialist claims of various sorts, so it would be surprising if there were some general theoretical argument that would block us from accepting Person Essentialism.²⁰

Nevertheless, Person Essentialism is a rather strong essentialist thesis, and controversial in ways that are not relevant to the present application of it. One might think that although no chimpanzees are in fact people, some chimpanzees could have been “upgraded” in such a way that they would have been people. If so, presumably some people—such as upgraded chimpanzees, if there are any—could have been “downgraded” so as to prevent them from being people, without preventing them from concretely existing. However, any such downgrade would seem to require altering the distribution of matter within the boundaries of the relevant object, in such a way that some of its paradigmatically intrinsic properties would be different.²¹ Thus, even if we reject Person Essentialism, the following principle still seems plausible:

Weak Person Essentialism Necessarily, every person has an intrinsic profile such that necessarily, if they instantiate that profile, they are a person.

Given S5, this is still sufficient to imply Incompatibility. Suppose x has an intrinsic profile F compatible with personhood. Since haecceities are intrinsic,

¹⁹Lewis himself would think that Person Essentialism is true in some contexts and false in others, thanks to different counterpart relations being in play.

²⁰Fine (2005) would reject Person Essentialism, since he maintains that people would still be people even if they had not concretely existed. But he would accept the alternative premise that necessarily every person is necessarily a person, which immediately entails in S5 that the haecceities of all non-people are incompatible with personhood.

²¹This is more plausible given the “whole-life” conception of intrinsicity (note 8). On the temporally local conception, one might object that personhood is “sticky” in such a way that if one upgrades a non-person into a person and later restores them to the intrinsic profile they had before the upgrade, they will remain a person.

it is possible that x is a person with intrinsic profile F . Then by Weak Person Essentialism, it is possibly necessary that if x is F , x is a person. So by the B axiom, if x is F , x is a person. Since x is F , we can conclude that x is a person.

Even Weak Person Essentialism might seem doubtful if we start contemplating exotic possibilities and have a pretty minimalistic conception of what intrinsic properties there are other than haecceities.²² But since we are only trying to argue that Incompatibility is *in fact* true, and not that it is necessarily true, we only really need the restriction of Weak Person Essentialism to *actually instantiated* intrinsic profiles. This can be stated as follows:

Super-weak Person Essentialism For every instantiated intrinsic profile F , necessarily, every F person is such that necessarily, if they are F , they are a person.

This follows from Weak Person Essentialism, since if F is an intrinsic profile, it is necessarily true that every F thing is such that F is its intrinsic profile. (It is no problem that intrinsic profiles instantiated by things like rocks are in the range of the initial quantifier: for them, the consequent is vacuously true.) Given S5 and the intrinsicality of haecceities, Super-weak Person Essentialism is still strong enough to imply Incompatibility.²³

One can certainly imagine theories of personhood that reject even Super-weak Person Essentialism. For example, one might think that being a person requires a certain social environment, such that people can become non-people just by changing their physical surroundings without any intrinsic change. On such a view, there may in fact be things whose intrinsic profiles are compatible with personhood, that are not people for want of the right environment. Universality and Intrinsicality will then imply—plausibly enough—that these unfortunate castaways enjoy moral status despite

²²Suppose for example that one thought that the only intrinsic properties were certain properties necessarily shared by any two spatiotemporally coincident things; haecceities; and Boolean combinations thereof. Then one might consider a counterexample along the following lines. Two chimpanzees, Harpo and Zeppo, are created from special matter that can interpenetrate. In fact they permanently coincide. Harpo is not a person, but Zeppo is endowed with a richer inner organization that makes him a person. But it could have been the other way round, so that Harpo was a person and Zeppo a non-person, while they still coincided. Thus, it could have been that Zeppo had the same intrinsic profile without being a person.

²³Suppose x has an intrinsic profile F compatible with personhood. Since haecceities are intrinsic, it is possible that x is an F person. By Super-weak Person Essentialism, necessarily every F person is necessarily (if F , then a person); so it is possibly necessary that if x is F , x is a person. By the B axiom, it is true that if x has F , x is a person. Since x has F , we can conclude that x is a person.

not being people. But on the natural ways of developing such views, they do little to undermine the restriction of Super-weak Person Essentialism to intrinsic profiles instantiated by objects *that largely overlap ordinary people*. They thus do little to support the claim that any *personites* have moral status.

4 Qualitativeness

In response to our discussion in the previous section, it would be natural for Johnston to modify the argument by replacing intrinsicity with some more demanding status, not enjoyed by haecceities.²⁴ One idea, suggested by some of the literature on this family of puzzles, is to replace ‘intrinsic’ throughout with something like ‘categorical and intrinsic’ (e.g. Cutter forthcoming: §3) or ‘non-modal and intrinsic’ (e.g. Builes and Hare 2023). But we are skeptical that there is any good way of understanding ‘categorical’ (or ‘non-modal’) on which it is plausible both that (a) Intrinsicity remains true when ‘intrinsic’ is replaced by ‘intrinsic and categorical’, and that (b) the target views (those where people are “ontological trash”) imply that in the vicinity of ordinary people there are many objects whose *categorical and intrinsic* profile is compatible with personhood. The most straightforward way of cashing out ‘categorical’ equates being categorical with being *undiscriminating* (such as to necessarily apply to both or neither of any two coincident objects). This gloss satisfies desideratum (b), since all of the many things that coincide with a person have exactly the same undiscriminating properties as that person. But it does not satisfy desideratum (a), since we see no intuitive reason to think that anything coincident with a person has moral status. For example, *aggregates of molecules* surely do not enjoy moral status, although every person coincides with an aggregate of molecules.²⁵ But if ‘categorical’ is understood in such a way that coincident objects can differ in their categorical properties, it becomes obscure whether the target views satisfy desideratum (b).²⁶

²⁴Less promisingly, one could deny that all haecceities are intrinsic. Our discussion in this section will also be relevant to this response.

²⁵If ‘coincide’ means ‘permanently coincide’, it’s not true that every person coincides with an aggregate of molecules; but there still could even be a person (e.g. a non-biological person made of metal and silicon, or a very short-lived biological person) who permanently coincided with an aggregate of molecules. Also, given plenitude, every biological person coincides with a worldbound object. But it doesn’t seem plausible to think that any worldbound objects have moral status.

²⁶The proposal to restrict the argument to *non-modal* intrinsic properties raises further issues. On an Intensionalist conception of properties, the very idea of an interesting division between modal and non-modal properties (as opposed to modal and non-modal *predicates*)

A more promising modification would replace ‘intrinsic’ throughout with ‘qualitative and intrinsic’. Indeed, it is common in the literature on intrinsic-ity to set aside the non-qualitative (see, e.g. Langton and Lewis 1998: 334-5), and it seems typically more plausible to interpret ‘exact intrinsic similarity’ as sharing of all qualitative intrinsic properties than as sharing of all intrinsic properties. So let us now turn to this more challenging version of Johnston’s argument. Call the revised first and third premises Qualitative Compatibility and Qualitative Intrinsicity.

While we have been treating ‘intrinsic’ as a monadic predicate applicable to properties (so-called ‘global’ intrinsicity), readers may also have come across the notion of ‘local’ intrinsicity, i.e. of an object *having a property intrinsically*. On its standard use, an object can have a property intrinsically even if that property is not intrinsic: for example, assuming sphericity is intrinsic, any spherical object will have the extrinsic property *being either spherical or within six feet of a panda* intrinsically. It is natural to think that these two notions are interdefinable via the following ‘global-to-local’ and ‘local-to-global’ principles (from Marshall and Weatherson 2023, crediting Humberstone 1996):

GTL x is intrinsically F iff x has an intrinsic property that necessitates being F .

LTG F is intrinsic iff necessarily, every F thing is intrinsically F .

The combination of these two principles yields a non-trivial claim just about the global notion:

GTG If necessarily every F thing has an intrinsic property that necessitates being F , then F is intrinsic.

We will discuss some objections to GTG (and hence to the conjunction of GTL and LTG) in section 6 below, but for now let’s accept it and see what follows.²⁷

seems suspect, since every property that can be expressed at all can be expressed by a predicate that contains modal operators. Such a distinction might be in a better standing for a hyperintensional theorist: for example, they might insist that *being a statue* is a non-modal property that necessitates (and perhaps explains) various modal properties. But a hyperintensionalist who makes this distinction and endorses Plenitude has every reason to think that all the modal differences between coincident objects required by Plenitude will be associated with (and perhaps explained by) non-modal differences, such as differences in *essence*, understood according to Fine’s non-modal account of essence.

²⁷For future-insensitivity and temporal locality as defined in note 12, GTG is automatic

GTG provide a plausible route from Person Essentialism, or even Weak Person Essentialism, to the conclusion that *personhood* is intrinsic. For Person Essentialism is equivalent to the claim that necessarily, for any person x , the conjunction of x 's haecceity with concreteness entails personhood and is instantiated by x . Given that, necessarily, haecceities are intrinsic, concreteness is intrinsic, and the conjunction of any two intrinsic properties is intrinsic, this implies that, necessarily, every person has an intrinsic property that entails personhood. Given GTG, it follows that personhood is intrinsic. To argue for the intrinsicity of personhood from Weak Person Essentialism, we need an additional, though highly plausible, assumption, namely that intrinsic profiles — properties of the form *being such that one's intrinsic properties are exactly the members of C* — are themselves intrinsic. Given that necessarily all haecceities are intrinsic, it is necessary that every object's intrinsic profile entails its haecceity, so Weak Person Essentialism is equivalent to the claim that necessarily, every person has an intrinsic profile that entails personhood. If intrinsic profiles are themselves intrinsic, this gives us what we need to conclude, by GTG, that personhood is intrinsic.

By combining the claim that personhood is intrinsic with three additional plausible premises, we can derive the analogue of Incompatibility for qualitative intrinsic properties:

Qualitative Incompatibility No non-person has a qualitative intrinsic profile compatible with personhood.

If this is true, then we are free to maintain that only people have moral status even while accepting both Universality and Qualitative Intrinsicity. The three premises are as follows:

- (1) Personhood is qualitative.
- (2) The negation of any intrinsic property is intrinsic.
- (3) The negation of any qualitative property is qualitative.

if we accept S5. Suppose that necessarily, every F thing has a temporally local property that necessitates being F , and consider a possible situation where some temporal interval r has an intrinsic profile G and contains an object x . If x is F , then x has a temporally local property H that entails F , so the proposition that r is G and is a temporal interval containing x entails that x is H , and hence that x is F . Meanwhile, if x is not F , then that proposition does not entail that x is F , so by S5, it entails that it does not entail that x is F , so it entails that x has no temporally local property that necessitates being F , and hence that x is not F . The case of future-insensitivity is parallel.

Together with the intrinsicality of personhood, these immediately imply that *being a non-person* is qualitative and intrinsic. Since all non-people obviously have this property, we can infer that none of them has a qualitative intrinsic profile compatible with personhood.²⁸

(Note that, unlike the argument in the previous section, this argument did not depend on S5, or indeed on any modal logic at all beyond the very minimal logic K. So the combination of the claim that personhood is intrinsic with (2) also provides a route to Incompatibility that could be compelling to those who reject an S5 logic for metaphysical modality, perhaps because they think that there are some counterexamples to the necessity of distinctness.²⁹)

Resisting the argument by denying (3) seems extremely unpromising: (3) is among the clearest fixed points anchoring our grip on the notion of qualitateness.³⁰ (2) might be accorded a similarly axiomatic status, and certainly falls out immediately from some ways of cashing out the informal notion of intrinsicality. (For example, if one thinks of an intrinsic property as one such that whether something has it is necessarily always settled by what it is like “on the inside”, it’s immediate that intrinsicality is preserved by negation.) Moreover, even if one rejected (2), one could still run the argument in S5 using the principle that intrinsic properties are necessarily intrinsic, which implies that the negation of any intrinsic property is entailed by the intrinsic profile of anything that lacks that property.³¹ And finally, if

²⁸Note that this argument only really needs the claim that the negation of any qualitative intrinsic property is qualitative and intrinsic. If, inspired by Lewis, you think that qualitative intrinsic properties are those that cannot divide duplicates, you will get this automatically.

²⁹If the necessity of distinctness fails, there is reason to doubt that every concretely existing non-person is necessarily not a person. For example, suppose that there are two electrons a and b that could have been identical (and hence necessarily identical), and consider a non-person Mark^\dagger that, necessarily, coincides with Mark whenever it concretely exists, and concretely exists at a time t iff Mark concretely exists at t and either $a = b$ or Mark has not written any papers about Macbeth’s dagger before t . If a had been identical to b , Mark^\dagger would have been necessarily coincident with Mark. On the natural view that there cannot be distinct but necessarily coincident material objects (*pace* Fine 2000), we can infer that if a had been identical to b , Mark^\dagger would have been identical to Mark. Since Mark would presumably still have been a person in that case, this means Mark^\dagger could have been a person. Note however that given (1)–(3), that would involve Mark^\dagger lacking one of its actual qualitative intrinsic properties, namely *not being a person*.

³⁰(3) follows from first two axioms of the ‘Basic Theory of Qualitateness’ in Dorr, Hawthorne, and Yli-Vakkuri 2021 (§14.1).

³¹When F is intrinsic and lacked by an object whose intrinsic profile is *being such that one’s intrinsic properties are exactly the members of C*, that intrinsic profile will entail *not being F unless F is not intrinsic or in C*. But given S5, it is impossible for F to be in C , and given the necessity of intrinsicality, it is impossible for F not to be intrinsic, so this property is

one rejected even that weakening of (2), one would then need to revisit the plausibility of Qualitative Intrinsicity. The claim that *having moral status* is an intrinsic and qualitative property seems to do as much justice as one could wish to the motivating impulse of ethical inclusiveness. If that claim is compatible with there being things that lack moral status whose qualitative intrinsic profiles do not necessitate *lacking moral status*, it is unclear what further considerations could be offered in support of there not being any such things.³²

Finally, there is the possibility of resisting the argument by denying (1). This seems *prima facie* unpromising: in the absence of an uncontroversial definition of ‘qualitative’, the concept is standardly introduced by providing lists of paradigmatically qualitative and paradigmatically non-qualitative properties, and personhood looks like a good candidate to be on the first list. However, as it happens, in the final chapter of Dorr, Hawthorne, and Yli-Vakkuri 2021, we were led to take seriously a view on which a wide range of properties that would normally be regarded as paradigms of qualitative-ness, including *being a person*, *being a table*, and *being conscious*, are in fact non-qualitative. This view provided a solution — albeit not the only workable solution — to a certain family of tricky modal puzzles, concerning the modal behaviour of instances of these properties at worlds that agree with actuality on all qualitative questions. There is no need to rehearse those puzzles here, since a parallel point about Qualitative Intrinsicity applies: if one were willing in response to the puzzles to deny the qualitiveness of *personhood* and *consciousness*, it would then seem completely unprincipled to nevertheless insist that *having moral status* is qualitative. After all, properties like *being conscious* seem like the best candidates to be the difference-makers when it comes to moral status. Thus, although denying (1) is a serious option, it is not a promising way of shoring up the argument for extending moral status to certain non-people.³³

equivalent to *not being F*.

³²(2) becomes automatic if we reinterpret ‘intrinsic’ as ‘future-insensitive’ or ‘temporally local’, as defined in note 12.

³³Even if we adopted the unprincipled-looking package where personhood is non-qualitative while *having moral status* is qualitative, we would still not have to be nearly as expansive in our attributions of moral status to non-people as Johnston’s argument would suggest. The only non-people to whom we would need to assign moral status are certain objects that perfectly spatiotemporally coincide with people at the actual world, and have modal profiles that from a qualitative point of view look just like those of people, but which are friendly in different ways to specific clusters of particle-haecceities. While some of these objects will differ from their coincident person as regards the range of possible worlds in which they *ever* coincide, we might reasonably maintain that each of them is necessarily

The combination of GTG with the intrinsicity of haecceities has other striking consequences for the theory of intrinsicity. One important such consequence also requires the *Iteration* principle according to which anything that is necessarily the case is necessarily necessarily the case. This seems plausible, and we have defended it at length elsewhere.³⁴ Given Iteration and GTG, we can derive that all “necessity properties”—i.e., properties of the form *being necessarily F*—are intrinsic. By Iteration, it is necessary that for all x , if x is necessarily F , x is necessarily necessarily F —in other words, x ’s haecceity necessitates *being necessarily F*. If haecceities are intrinsic, we can conclude that for all x , if x is necessarily F , x has an intrinsic property—its haecceity—that necessitates *being necessarily F*. By GTG, this is sufficient for *being necessarily F* to be intrinsic.

The claim that all necessity properties are intrinsic would be entirely implausible if we were talking about some run-of-the-mill restricted modality. For example, on the interpretation of ‘couldn’t have learnt Finnish’ on which it applies to people who have no access to Finnish lessons but not to people who have such access, it surely expresses an extrinsic property. But for *metaphysical* necessity—especially given our favoured view that metaphysical necessity is the broadest form of necessity—the idea that necessity properties are intrinsic strikes us as quite attractive, even setting aside the above derivation from GTG. Given plenitude, it is natural to think that objects are “individuated” by their modal profiles—in specifying an object’s modal profile, one “says what it is”—and so the intuitions that support the claim that haecceities are intrinsic seem to carry over to modal profiles.

Counterpart theorists like Lewis (1968, 1986) identify necessity properties with properties specified in terms of quantification restricted to an object’s counterparts, e.g., equating *being necessarily not a swan* with *not having any counterparts that are swans*. Proponents of such views will probably find it unintuitive to classify necessity properties as intrinsic. On their view, an object’s having a necessity property *is*, intuitively speaking, a matter of how it is related to the rest of reality rather than a matter of how it is in itself.

such that if it *ever* coincides with the person, it *always* coincides with the person. Many of Johnston’s strategies for extracting bizarre conclusions about permissibility from attributions of moral status to personites will not apply if moral status is as limited as this. The only possible exceptions are certain arguments turning on the *cardinality* of things with moral status. But we don’t find these arguments very strong, for reasons we will explain in section 8 below.

³⁴See Dorr, Hawthorne, and Yli-Vakkuri 2021 (ch. 8). We argue there that metaphysical necessity should be identified with the *broadest* necessity, and that this in turn vindicates Iteration for metaphysical necessity.

But Lewis's views about the metaphysics of modality are notoriously idiosyncratic and *prima facie* implausible: it is not a good idea to treat him as representative of all views that regard people as "ontological trash".³⁵

The claim that all necessity properties are intrinsic provides an alternative argument for Qualitative Incompatibility. This alternative route requires even weaker essentialist commitments than those needed by the above argument via the intrinsicity of personhood. Like that argument, it requires (1)–(3), and does not require S5. The required essentialist premise is as follows:

Qualitative Super-weak Person Essentialism Every instantiated qualitative intrinsic profile F is such that necessarily, every F person is such that necessarily, if they are F , they are a person.

This looks roughly as plausible as the earlier version without 'qualitative'. *Prima facie*, it would seem impossible to turn a person into a non-person without affecting its qualitative intrinsic properties. And even if there are exotic possibilities where this fails, it seems unlikely to fail for any of the *actually instantiated* qualitative intrinsic profiles.

To derive Qualitative Incompatibility, consider any non-person x with qualitative intrinsic profile F . Let G be the property *necessarily being a person if one is F* . G is intrinsic since it is a necessity property, and qualitative since personhood and F are both qualitative. So by (2) and (3), the negation of G —which is instantiated by x , since x is F and not a person—is also qualitative and intrinsic, and therefore entailed by x 's qualitative intrinsic profile F . But by Qualitative Super-weak Person Essentialism, *being an F person* entails G . *Being an F person* is thus impossible (since it entails both G and its negation): in other words, F is incompatible with personhood. Note that this argument did not rely on any modal logic beyond KT, although we did appeal to Iteration (the characteristic axiom of S4) in the earlier argument for the intrinsicity of necessity properties.

³⁵Setting aside arguments like the one we gave, we don't have strong *intuitions* one way or the other as regards whether necessity properties are intrinsic. Eddon (2011), by contrast, finds it intuitive that *being such that there possibly exists something greater in mass* is extrinsic, giving this judgment equal billing with the judgment that haecceities are intrinsic as objections to Lewis's duplication-theoretic account. We suspect that this intuition reflects some residual influence of Lewis's theory of modality.

5 Maximality

The claim that personhood is intrinsic is surprising. Ted Sider (2001, 2003) has influentially argued that a certain wide range of properties including *being a person*, *being a house*, and *being a rock* are extrinsic, on the grounds that these properties are “maximal”. He roughly glosses ‘*F* is maximal’ as ‘large parts of *F* things are not themselves *F*’.³⁶ Maximality, he says, is a special case of ‘border-sensitivity’, where ‘*A* property is border-sensitive iff whether it is instantiated by an object depends on what is going on outside that object at its borders’ (Sider 2001: 2). And according to Sider, ‘Any analysis of intrinsicity ought to count maximal and border-sensitive properties as extrinsic’ (4).

The heuristic Sider is relying on in deriving this extrinsicity claim is something like the following: if *being F* entails *being such that one’s environment is G*, and *G* is intrinsic, and *F* and *G* are both qualitative, then *being F* is not intrinsic.³⁷ It is quite a challenging task to formulate this thought in a way that is rigorous, consistent, and strong enough to sustain arguments like Sider’s. Lewis’s ‘principle of recombination’ (Lewis 1986: §1.8) is the best-known attempt in this direction, though it is not entirely rigorous (thanks to a weaselly ‘size and shape permitting’ proviso), and is not actually strong enough to sustain Sider’s argument.³⁸ But we need not further explore this, since the combination of plenitude with the intrinsicity of necessity properties rules out even very weak versions of recombination, along with obvious precisifications of Sider’s claim that maximal and border-sensitive properties are extrinsic. Consider for example the following modest recombination principle: if both *F* and *G* are qualitative, intrinsic, possibly instantiated, and entail *being spatiotemporally compact*, then it is possible that there is both an *F* thing and a *G* thing. Let *G* be some ordinary qualitative, intrinsic, un-

³⁶He also suggests a more refined gloss that adds the proviso ‘...because they are large parts of *F* things’. The interpretation of this ‘because’ raises tricky questions: see note 41 below.

³⁷Here, we might identify an object’s environment with the complement of its exact spatiotemporal location. This way of talking goes most smoothly if we don’t mind positing an empty region—then an object exactly located at the whole spacetime manifold has the empty region as its environment.

³⁸The problem is that Lewis’s official principle, despite some of his informal glosses on it, is silent about the possible spatiotemporal arrangements of instances of intrinsic properties. It says that if *being a bronze statue with shape S* and *being a large bronze cube with an S-shaped cavity* were both intrinsic, then it would be possible for an instance of the former to coexist with an instance of the latter, but it does not imply that it would be possible for there to be an instance of the former fitting exactly into the cavity inside an instance of the latter, which is what would be needed for a *reductio* of the intrinsicity claim.

discriminating, possibly instantiated, possibly uninstantiated, compactness-entailing property—say, *being cube-shaped and having a mass of 1 kg*. Let F be the conjunction of the three properties *concreteness* (concrete existence), *necessarily being compact if concrete*, and *being, necessarily, such that nothing is G if one is concrete*. Given that necessity properties are intrinsic, F is qualitative and intrinsic (since all three conjuncts are). Moreover, it is obviously necessary that any F thing is compact, and impossible for there to be both an F thing and a G thing. So, the only way for our weak recombination principle to be true would be for F not to be possibly instantiated. But plenitude says that F is possibly instantiated, since it says (for example) that for any object x , there is an object y that necessarily coincides with x if it concretely exists, and concretely exists iff x concretely exists, x is compact, and nothing is G .³⁹

Indeed, even without any form of plenitude, the intrinsicity of necessity properties fits very poorly with this sort of recombination principle.⁴⁰ Consider *being concrete and such that, necessarily, if one is concrete, there is a sperm and an egg outside one's spatiotemporal boundaries*. If we are right this is intrinsic, and if Kripke is right it is possibly (since actually) instantiated; but obviously there are substantial modal constraints on the environments in which this property can be instantiated. For example, it could not be instantiated by a *lonely* object (one with nothing spatiotemporally outside its boundaries). So, the picture of intrinsicity we have been led to from plenitude and the intrinsicity of haecceities is deeply at odds with the kinds of modal tests for intrinsicity that are tacitly in play in Sider's argument and that inspire the intrinsicity-theoretic recombination principles. The maximality, or more generally border-sensitivity, of properties like *being a person*, *being a house*, and *being a rock* is not a good reason to classify them as extrinsic.⁴¹

³⁹Similarly, for any qualitative, undiscriminating, concreteness-entailing property F —even a paradigmatically extrinsic one like *being six miles north of a bronze sphere*—if there is something that could be F , there is a possibly-instantiated intrinsic property that entails being F , namely *being concrete and, necessarily, F if concrete*. For example, where x is some arbitrary object that could be F , this property will be possibly instantiated by the object that necessarily concretely exists iff x is F and necessarily coincides with x if it concretely exists.

⁴⁰Of course there are other kinds of recombination principles that are compatible with our picture, such as “combinatorialist” principles that try to spell out the idea that any logically consistent pattern of fundamental properties and relations is possible (Hawthorne and Russell 2018; Bacon 2020). One could also endorse intrinsicity-theoretic recombination principles that restrict the kinds of objects that matter—for example, the principle that when F and G are qualitative, intrinsic, possibly instantiated properties that entail *being a compact spatiotemporal region*, it is possible for there to be both an F region and a G region.

⁴¹It is actually not so clear that we should accept that these properties count as ‘maximal’ according to Sider's most careful explanation of that term. In a footnote, Sider (2001: n. 1)

6 Hyperintensionality

Our arguments in section 4 for the claims that personhood is intrinsic and that necessity properties are intrinsic were based on the following principle:

GTG If necessarily every F thing has an intrinsic property that necessitates being F , then F is intrinsic.

GTG immediately implies that intrinsicity is *non-hyperintensional*: any property necessarily equivalent to an intrinsic property is itself intrinsic. For if F is intrinsic and necessarily equivalent to G , then necessarily any G thing has an intrinsic property (namely F) that necessitates G , which by GTG suffices for G to be intrinsic.

The claim that intrinsicity is non-hyperintensional is very controversial. Of course, it follows immediately from Intensionalism, the view that necessarily coextensive properties, relations, and propositions are always identical. But Intensionalism is itself deeply controversial. Various supposed counterexamples to the non-hyperintensionality of intrinsicity can be found in the literature (see Francescotti 1999; Eddon 2011; Bader 2013; Marshall and Weatherson 2023; and for a contrary perspective, Hoffmann-Kolss 2014). Although none of them strike us as compelling enough to amount on their own to a strong case against Intensionalism, some do look suggestive against the background of some worked-out non-Intensionalist view. This casts some doubt on GTG, and on the conjunction of the principles GTL and LTG from which it follows.

However, GTG is really more than we need. It is easy to check that our arguments for [Qualitative] Incompatibility would still go through if one replaced the claim that personhood or necessity properties are [qualitative

clarifies that for F to be maximal, it is not sufficient that no F thing could be a large proper part of another F thing; additionally, 'large parts of F s should be disqualified as being F s because they are large parts of F s'. The 'because' here opens a huge can of worms. For example, if one thought that no necessary truth is true because of a contingent one, and that no large proper parts of people are even possibly people, one would deny that any large proper parts of people are non-people *because* they are large proper parts of people. Sider's definition of 'border-sensitive' using 'depend' will raise similar issues on some ways of cashing out 'depend'. While we guess that Sider intended a purely modal interpretation of 'depend', even giving a modal gloss on 'border-sensitive' will raise some delicate issues, similar to those raised by our definition of 'future-insensitive' in section 2. Consider an object which, necessarily, coincides with a certain bronze sphere s if s is a bronze sphere surrounded by water, coincides with an outer layer of s if s is a bronze sphere not surrounded by water, and otherwise does not concretely exist. On the most straightforward modal ways of spelling out 'border-sensitive', examples like this will mean that *pace* Sider, even *being a solid sphere* counts as border-sensitive.

and] intrinsic with the claim that these properties are necessarily equivalent to some [qualitative and] intrinsic property. These claims could in turn be motivated by the following hyperintensionality-friendly weakenings of GTG:

Weak GTG If necessarily every F thing has an intrinsic property that necessitates being F , then F is necessarily equivalent to an intrinsic property.

Qualitative Weak GTG If F is qualitative, and necessarily every F thing has an intrinsic property that necessitates being F , then F is necessarily equivalent to a qualitative intrinsic property.

Both weakenings seem fairly attractive even on the assumption that intrinsicity is hyperintensional.⁴²

Granted, neither of these weakenings seems anything close to being undeniable. For example, one can imagine a principled view that rejects Qualitative Weak GTG on the grounds that the only qualitative intrinsic properties are *undiscriminating* properties like *being cubical*. If so, modal properties that distinguish coincident objects aren't even coextensive with, and a fortiori are not necessarily equivalent to, any qualitative intrinsic properties. But the qualitative version of Intrinsicity looks deeply tendentious on this picture: aggregates of molecules lack moral status, even when they coincide with people.

On the other hand, if the hyperintensionalist is willing to allow that *some* coinciding objects differ in their qualitative intrinsic properties, it seems overwhelmingly natural to think that *all* coinciding objects that differ in their

⁴²Indeed, given S5 together with the higher-order Barcan Formula (which can be derived in our favoured logic from S5), there is a plausible argument for Weak GTG from the very plausible claim that disjunctions (including infinite disjunctions) of intrinsic properties are themselves intrinsic. (Suppose that F entails *having an intrinsic property that entails F* , and consider the disjunction of all intrinsic properties that entail F . This also entails F , and is intrinsic. Moreover, since every intrinsic property that entails F entails the disjunction, we can appeal to S5, the higher-order Barcan formula, and the necessity of intrinsicity to conclude that *necessarily* every intrinsic property that entails F entails the disjunction, and hence that F entails the disjunction. So the disjunction witnesses the truth of Weak GTG.) Note however that if we accept S5, we can already argue for Incompatibility using the arguments from section 3, so Weak GTG is redundant: the interesting claim in this setting is Qualitative Weak GTG. If we didn't want to assume S5, we could also derive Weak GTG directly from the claim that properties of the form *having an intrinsic property that entails F* are all intrinsic. However, we might expect those who have doubts about the intrinsicity of necessity principles to also have doubts about this claim.

qualitative modal profiles also have different qualitative intrinsic properties. For example, it would be natural for hyperintensionalists who follow Fine (1994) in employing a notion of essence thought of as drawing hyperintensional distinctions to claim that *essentiality* properties (like *being essentially a person if concrete*) are intrinsic, even if certain of the necessity properties they entail are not. Given this view, it is also natural to adopt some essence-theoretic strengthening of plenitude on which the multiplicity of instantiated modal profiles also goes along with a corresponding multiplicity of *essentiality* profiles.⁴³ The general idea here seems very plausible: the differences in qualitative necessity properties among the objects coincident with a given object should go hand in hand with qualitative differences in their *essential nature*; and *having such-and-such essential nature* looks like it should be intrinsic for the same reason that haecceities are. But Qualitative Incompatibility still looks *prima facie* quite plausible (and Compatibility entirely unmotivated) if coinciding objects are in this way distinguished by intrinsic essentiality properties.

In sum, even in a hyperintensional setting, it is very hard to imagine a principled theory of intrinsicity that is demanding enough to make Intrinsicity plausible, but not so demanding as to make Compatibility implausible.⁴⁴

⁴³Similarly, a hyperintensionalist might embrace the doctrine of 'Identity-strength Location Plenitude' (Dorr, Hawthorne, and Yli-Vakkuri 2021: 277-8), and maintain that properties of the form *being an object such that for it to be located at a point is for the point to be F* are intrinsic even if necessity properties are not.

⁴⁴One possible intermediate view would claim that the modal properties *being possibly F* and *being necessarily F* are necessarily equivalent to qualitative intrinsic properties when *F* itself is qualitative and intrinsic, but not otherwise. So, for example, *being possibly cubical* and *being necessarily cubical if concrete* are intrinsic, but *being possibly adjacent to something cubical* and *being necessarily adjacent to something cubical if concrete* are not. This should meet the desideratum of providing enough qualitative intrinsic properties to distinguish a person from a coinciding aggregate of molecules, but not enough to distinguish it from everything coincident with it. But this picture makes Intrinsicity look dubious. Consider an object Mark[†] that, necessarily, coincides with Mark if it concretely exists, and concretely exists iff someone with Mark's qualitative profile lived exactly one billion years earlier in the same place as Mark. Clearly, every qualitative intrinsic indiscriminating property that could be instantiated by Mark[†] could be instantiated by Mark, and it seems plausible that the converse is true too: for any world where Mark has a certain qualitative intrinsic property, one should be able to find a world where Mark[†] also has that property by insert a duplicate of Mark a billion years earlier. So, when Mark[†] concretely exists, the view we are considering should say that it has the same qualitative intrinsic profile as Mark. But Mark[†] seems far too modally fragile to be a plausible subject of moral status: holding fixed what things were like a billion years ago, Mark[†] couldn't have concretely existed while behaving even slightly differently.

7 Consciousness

Even if one feels one has an intuitive grasp of the notion of intrinsicity, it is hard to form a clear pre-theoretic judgment one way or the other as regards whether personhood is intrinsic. And even the question whether *moral status* is intrinsic seems highly theory-laden. By contrast, many philosophers have a very clear judgment that *consciousness*, as well as various determinates of consciousness such as *being in pain*, are intrinsic (e.g. Merricks 2001; Pautz 2013; Chalmers 2004). Although others have developed theories of consciousness that seem to imply its extrinsicity (Dretske 1996; Lycan 2001), this has widely been seen as a major problem for such theories.⁴⁵ The claim that consciousness is intrinsic entails that everything that lacks consciousness has an intrinsic profile that necessitates *lacking consciousness*, and thus gives rise to an argument for the claim that many personites are conscious, structurally parallel to Johnston's argument for the claim that many personites have moral status. If the consciousness argument worked, then the moral status argument would to some extent be redundant, since it is independently plausible that all conscious beings have moral status on at least some precisifications of 'moral status'. Meanwhile, if one could stomach the claim that consciousness is extrinsic as a way of blocking the latter argument, it would seem rather unprincipled to remain committed to the intrinsicity of moral status. So, in a way, the consciousness argument is where the real action is. But now that we have seen how to sustain the view that no non-people that largely overlap people have moral status without giving up on the intrinsicity of moral status, we can equally well see how to sustain the parallel view for consciousness.

It is worth noting that we can also give a rather plausible argument for the intrinsicity of consciousness, parallel to our earlier argument for the intrinsicity of personhood. This argument uses GTG, the intrinsicity of haecceities, and the consciousness-theoretic analogue of Weak Person Essen-

⁴⁵Our discussion of intrinsicity does however suggest that the route from so-called "externalist" accounts of various mental properties to the claim that these properties are *extrinsic* is in need of some rethinking. In principle, one might even maintain that *thinking about water* is an intrinsic property. Although *being a region that contains someone thinking about water* is extrinsic, only might maintain that people have extremely demanding essences, so that a person thinking about water could not concretely exist without being related to water in a certain way—such that nothing further beyond appropriate goings-on within their spatiotemporal boundaries is needed for them to think about water. Although this doesn't seem at all plausible, similar essentialist ideas may be more plausible when it comes to the considerations that might be thought to support the extrinsicity of consciousness.

tialism:

Weak Consciousness Essentialism Necessarily, any conscious object has an intrinsic profile such that necessarily, if it has that profile, it is conscious.

Informally: making something stop being conscious always requires somehow changing what's going on within its boundaries. This seems quite plausible, even from the point of view of a committed physicalist who hasn't yet given any thought to the question whether consciousness is intrinsic. After all, all the familiar ways of losing and gaining consciousness manifestly involve uncontroversially intrinsic physical changes.

The foregoing reflections help us evaluate certain variants of Johnston's argument that might be thought to circumvent our objections to the versions considered so far. In particular, Johnston's favoured formulation in terms of 'mental and physical lives' might suggest a restriction where the relevant intrinsic properties are required to be not only qualitative, but also to be either physical or mental. The first thing to note about this version of the argument is that it will not be persuasive to people who think that all mental properties are extrinsic, e.g. because they think that all mental properties are *maximal* and they accept Sider's argument from the previous section. But even on our view where some mental properties (including consciousness) are intrinsic, the argument will still go nowhere, so long you as you also agree with us that ordinary people do not overlap anything else with mental properties. For in that case, the intrinsic mental properties that people have and personites lack should make it unmysterious that they differ in moral status.

The question exactly what goes wrong with this version of the argument depends on whether one thinks that there could be a person that never has any mental properties.⁴⁶ If you think this isn't possible, then you should presumably also think that being a person requires being conscious at some time (or at least, having some intrinsic mental property that works analogously to consciousness). In that case you should reject the Compatibility premise: indeed, you should think that there is no non-person whose lifetime profile of physical and mental intrinsic properties is compatible with personhood. If on the other hand you think that there could be people in a lifelong vegetative state that involves lacking all mental properties, your

⁴⁶If one counts negations of mental properties, such as *not being conscious*, as themselves mental, then the relevant question is whether there could a person that never has any "positive" mental properties, in the sense that their mental properties are those of a stone.

assessment of the argument will turn on how broadly you interpret the category of 'physical' properties. If you interpret it broadly to include super-discriminating properties like *being an organism*, then these properties can take over the role of mental properties in explaining the failure of Compatibility. If you interpret 'physical' narrowly, so that all qualitative intrinsic physical properties are undiscriminating, then vegetative people (assuming that there actually are some) will share all of their physical and mental intrinsic properties with everything that coincides with them. On this interpretation, the problem lies with one of the other premises (Universality or Intrinsicity). Either the vegetative people do not have moral status, or their moral status isn't explained by their physical and mental intrinsic properties, but by other properties (such as *possibly being conscious*).

Although our picture of consciousness respects both the judgment that it is intrinsic and the judgment that there are no other conscious things in the immediate vicinity of an ordinary person, it is in other ways quite disorienting. For us, consciousness is just one member of a huge family of properties, all of which are metaphysically on a par. All of these properties are like consciousness in being intrinsic and having at most one instance overlapping any ordinary person. Most of them are not instantiated by any people, but by objects very similar to people in all underlying metaphysical respects but differing slightly either in their actual or counterfactual spatiotemporal locations. All the properties in the family are, or could easily have been, expressed by the word 'conscious'.⁴⁷ This conflicts with the tempting thought that consciousness is an especially "attention grabbing" property, such that any sufficiently sophisticated language-using creature who instantiated it would be likely to have a word that referred determinately to it.⁴⁸

⁴⁷There is an anti-physicalist version of the view that people are ontological trash that accepts that certain *events* have phenomenal properties that are attention-grabbing in this way. But this view doesn't provide a path from determinate reference to these properties of events to determinate reference to properties of people, owing to the vagueness of 'subject of'.

⁴⁸Our view of consciousness and other mental states also conflicts with some standard presentations of functionalism, according to which all it takes for an object to have some mental state such as consciousness is for it to have a family of internal states which have some appropriate network of dispositions to cause and be caused by one another and certain physically characterized "inputs" and "outputs". Typical sketches of such theories seem to imply that any ordinary conscious object will overlap many other conscious objects. Indeed, as Zimmerman (2003: §3.2) observes, such theories would seem to predict that even the *mass of matter* coincident with a person—assuming there is such an object—will have some mental properties. (Given its greater disposition to become fragmented, these mental properties may not be the same as those of the coinciding person—Zimmerman suggests that the mass of matter it may instead be mentally 'more like a monkey than a man'). The

You might surmise that even if the intrinsicity argument doesn't work, these surprising commitments will somehow provide an alternative route to a disastrous proliferation of moral status. The next two sections will consider some possible strategies for developing such an argument.

8 Similarity

We have been focusing so far on arguments concerning *exact* similarity in certain respects (e.g. intrinsic respects) between actual non-people and possible people. In response to our objections, Johnston might fall back on an appeal to a different kind of argument, crucially invoking a notion of similarity that comes in degrees. The thought is that given physicalism (or one of the other target views), some personites are *similar* to people in all underlying respects (or at least in all underlying *intrinsic* respects), and that objects that in this way “similar from a God's-eye point of view” must also have similar degrees of moral status.⁴⁹

There are suggestions in Johnston's discussion of this sort of fallback argument. He argues for his version of Intrinsicity by deriving it from a logically stronger principle according to which ‘that it cannot be that two beings are significantly alike intrinsically while the one has a moral status and the other does not’ (Johnston 2016: 211). He suggests that this principle is needed for the rational reconstruction of our recognition of moral status for those not of our race or species.

Given what we have learnt by thinking about the intrinsicity argument, we should not expect the similarity-based argument to secure moral status for all of the wide variety of objects that Johnston would count as ‘personites’. For example, any object that (in the actual world) coincides with Mark

conflict between the relevant kind of functionalism and our view is not unique to the philosophy of mind, but arises for a wide range of candidate definitions of ordinary properties that might seem roughly correct if one were not focusing on the need to discriminate coincident and near-coincident objects. For example, the OED defines a ‘lake’ as a ‘large body of water entirely surrounded by land’, but one might worry that—unless ‘body of water’ is being used in a specially discriminating way—this definition implies that Lake Ontario coincides with many other lakes, including the aggregate of water molecules currently coincident with Lake Ontario. The lesson we draw is that the usual reasons for caution about the project of *defining* ordinary vague words apply with special force for the many words that express highly discriminating properties, whose instances do not ordinarily overlap. However, it seems plausible that the general idea of functionalism can be understood in such a way that this is not an objection to it.

⁴⁹We have in mind here some notion of similarity that goes hand in hand with considerations of naturalness—see Dorr and Hawthorne 2013 for some discussion of how this connection might be cashed out.

until his arrival in Budapest and then passes out of concrete existence seems, in certain respects that seem quite legible even from the God's-eye point of view, *very* dissimilar from any possible person—for no person could cease to exist without undergoing physical changes dramatically unlike any of the changes undergone by Mark during his arrival in Budapest.⁵⁰ And we should not think that the power of the similarity argument to multiply bearers of moral status will be much increased if we restrict the relevant underlying respects of similarity to *intrinsic* respects of similarity: for as we have argued, differences in modal profile are, or are accompanied by, intrinsic differences. Indeed, it will be quite challenging to use similarity-based arguments to justify ascribing a significant level of moral status to any non-people other than things whose modal spatiotemporal profiles are *very* close to those of people. Given this, many of Johnston's techniques for generating disastrous conclusions about moral permissibility from the ascription of moral status to personites aren't going to work. The only events that will be survived by some but not all of the objects within the cluster of overlapping objects similar to a given person will be unusual ones near the boundaries of what the person can survive. In these hard cases, we do not have clear pre-theoretic judgments about whether, for example, costs prior to the event can be justified by benefits afterwards.

Johnston does however have one argumentative technique that remains applicable, namely arguments that turn on the idea that the *number* of things with person-like moral status is much greater than we normally take it to be, with morally destabilizing results. Since modal space is infinite, there will plausibly be *infinitely* many objects very similar to, and permanently coincident with, any given person. So, in situations where we would normally think 'If I do A, I will save a hundred people from an early death, and if I do B, I will only save three people, so I should do A', we should be thinking 'If I do A, I will save infinitely many moral-status-bearers, and if I do B, I will also save the same infinite number of moral-status-bearers'. It is natural to worry that this thought will undermine the intuitive judgment that doing A is morally obligatory in the relevant situations.

⁵⁰There is a temptation to say that the person Mark is very similar to an object that coincides exactly with Mark at all non-actual world-histories and actually coincides with Mark up to his arrival in Budapest. But we don't think that's the right way to conceive 'similarity in all objective respects': it's like saying that the property *being a dog* is very similar to the property *being a dog that is not such that: things are just as they actually are and it is the closest dog to the Eiffel Tower*. Modal 'shape' matters to similarity; and a change at one world far from the boundary can make a big difference to shape, just as removing one point can turn a topological sphere into a topological torus.

However, the only obvious way to turn this worry into an *argument* will require appealing to some premise along the lines of ‘When two actions will each save the same number of moral-status-bearers from an equally bad fate, and the two groups do not overlap, and there are no other morally relevant differences such as personal relationships or promises, then both are morally permissible if either is’. And in cases involving infinite collections of people, this premise is deeply problematic even on the assumption that no non-people have moral status. For in an infinite world, it is all too easy for things to be set up in such a way that any two of our options will differentially affect infinite groups of people in such a way that if the premise were true, all sorts of morally appalling options would turn out to be permissible. And this disaster is not an inevitable consequence of any systematic moral theory applicable to infinite populations, since there are competing ideas that avoid the disaster and also have considerable intuitive support. For example, it seems quite intuitive to think that outcomes where the spatial or temporal *density* of people living good lives is higher than the density of people living bad lives are *ceteris paribus* better and to be promoted, even when the cardinalities are equal.⁵¹ The cardinality-theoretic considerations are thus very far from providing a knock-down route from the infinite multiplication of moral-status-havers to any disastrous, or even controversial, claims about the moral permissibility of particular actions. They do nevertheless present those who accept the multiplication with an important *challenge*, namely to find a systematic theory that explains why ordinary person-counting is such a good guide to morally permissible action despite ignoring almost all the relevant bearers of moral status. The view that spatial density matters does not directly help address this challenge, though it does provide a precedent for a systematic theory that does not succumb to ‘infinitary paralysis’.⁵²

⁵¹For example, in a world where every hundredth house along an infinite road is red and the rest are green, one should choose an outcome where the people in green houses have happy lives and the ones in red houses have sad lives over an outcome where the people in red houses have happy lives and the ones in green houses have sad lives. See Kagan and Vallentyne 1997 and Arntzenius 2014. Such views pair naturally with views on which spatial density matters to *rational self-locating uncertainty* within such worlds (Arntzenius and Dorr 2017), especially given Harsanyi-style thoughts about the relation between ethical preference and selfish preference under conditions of ignorance (Harsanyi 1955).

⁵²If one is open to the idea that spatial organization matters, one might also be open to the idea that the organization of the moral-status-bearers in some “similarity space” matters too. Perhaps, for instance, we could define some reasonable notion of volume in similarity space in such a way the total volume of the infinite set of moral-status-bearers who largely overlap some finite set of the people is roughly proportional to the cardinality of the latter

Of course this challenge will only need to be faced if the similarity argument works, and it is by no means straightforward to flesh this argument out in a compelling way. Generally speaking, even when *perfect* similarity in underlying respects suffices for perfect similarity in certain supervening respects, *imperfect* similarity in underlying respects can be compatible with radical dissimilarity in supervening respects. The reason is that for some scales, even being on the scale at all requires being a thing of some specific kind, where (assuming plenitude) each thing of that kind is extremely similar to things not of that kind. For example, consider banknotes which have a certain value. Each banknote is coincident with innumerable other things which are not banknotes. But it seems wrongheaded to think 'This bill has a monetary value of \$20; there are innumerable things that largely overlap the bill and are also very similar to it in all underlying respects; therefore there are innumerable things that largely overlap the bill and have a monetary value of at least \$19'. It seems more plausible to think of 'has monetary value' as super-discriminating in the same way as 'is a banknote' and 'is a person'. Similarly, it seems silly to think of similarities in *social class* in the analogous way, so that many of the non-people who coincide with some upper class person must be at least upper middle class. On our view, all mental scales are super-discriminating. Since the non-people who are similar in underlying respects to a very intelligent person experiencing severe pain don't have any (positive) mental properties, they are not even moderately intelligent, and do not experience even moderate pain. Against this background, it would be no surprise for the scale of moral status also to be super-discriminating. If there is a compelling objection to this picture, it will have to turn on some specific considerations having to do with mental properties or moral status, rather than some general metaphysical doctrine on which sufficient similarity in underlying respects always secures similarity in high-level respects.

Johnston's main strategy for arguing that there is something specifically *morally* bad about perspectives like ours involves an extended analogy between such perspectives and forms of 'racism, sexism, and speciesism' that draw the 'circle of respect' in some repugnantly narrow way. According to Johnston's analogy, our view is like that of speciesists who think that creatures without the highly specific biological features characteristic of *Homo sapiens* are altogether lacking in moral status, and justify this by positing that such creatures are also not conscious (and lack the capacity for consciousness). There certainly is an analogy here, though of course there is also an ana-

set.

logy between us and sensible non-panpsychists who deny that rocks are conscious. And there is also an analogy between the three debates in that it is hard to find theoretical arguments that will be persuasive to those fully committed to the denial of consciousness and moral status to the relevant groups. Intrinsicity-based arguments are completely ineffective against the speciesists we have in mind, since (unlike the ones Johnston considers) they attribute the differences to certain intrinsic biological characteristics that are in fact highly correlated with membership in *Homo sapiens*, rather than to differences in evolutionary history. And it is also challenging to get anywhere with some general doctrine about degrees of similarity. Even if one granted that *sufficient* underlying similarities would be enough to secure mental and moral similarities, it is hard to think of a dialectically effective argument that the degree of underlying similarity is too great to permit the particular high-level differences that the speciesist alleges.

Still, theoretical arguments from neutral premises aren't everything; and 'Your position is analogous to that of a racist, sexist, or speciesist!' may on occasion be a helpful way of shaking someone's confidence in a misguided view by appealing to their better nature. But notice that if we are analogous to speciesists who say 'These fine-grained biological features are a key part of the supervenience base of consciousness and moral status', Johnston is analogous to equally pernicious speciesists who say 'Since obviously these non-humans aren't conscious and lack moral status despite being *physically* so similar to us, there must be some additional, fundamental non-physical property — maybe consciousness itself, or some mysterious "quintessence" (cf. Johnston 2016: 202) — that distinguishes us from them'. (Or, even more closely, to those who say 'Those so-called non-human animals are mere aggregates or agglomerations of parts, radically different even in ontological category from genuinely unified substances, like us human beings'.) The only ones who come out of this analogy looking good are the philosophers who stick with physicalism (or one of the other views targeted by Johnston's argument) and embrace a vast multiplication of both consciousness and moral status.

How should we make our peace with our preferred way of drawing the circle of respect in the presence of these analogies? We will conclude this section by pointing to two disanalogies between our own view and that of the speciesists which seem quite telling to us, and have helped us avoid being troubled by feelings of guilt by association.

First, it is natural to imagine that the speciesists' misguided views about non-humans are reflected in their *preferences* among possible situations. For example, so long as the human beings in one world are even slightly better

off than the human beings in another world, they prefer the former world, no matter what is going on with the non-humans. One can imagine ways in which the contrast between people and very similar non-people that overlap them might, similarly, show up in someone's system of preferences. For example, they might conform to a 'Pareto' principle according to which one world is preferred to another when it is at least as good for every person and better for some people, while not conforming to parallel principles featuring properties very similar to personhood. But we argue elsewhere (Dorr and Hawthorne 2024) that such preferences are morally inappropriate: a virtuous person's preferences among possible worlds do not in any natural way single out the property of personhood from other nearby properties. As we see it, a preference-system that did give personhood a special status would be "fetishistic" in a way analogous to the preferences of someone who prefers to be thin, but is altogether indifferent to all differences in their bodily shape other than the particular differences that mark the boundary of thinness (a property that is in no way metaphysically special relative to other properties that draw boundaries in slightly different places). Although we grant that the *objectual* attitudes of a virtuous person will treat personhood very differently from other nearby properties—for example, they will love, or at least respect, the people they know while not loving or respecting any of the non-people that overlap them—the contrasts will be invisible if we just focus on their *propositional* attitudes, including preference (among possible situations).⁵³ We also are attracted to the view that degrees of goodness of possible worlds—the subject-matter of axiology—correspond to virtuous preferences, meaning that personhood will also not be distinguished from other similar properties by any special role in axiology. These contrasts give our view a very different cast from that of the speciesists.

Here is second disanalogy that also seems revealing. Imagine an isolated community of 'Pluraliens' that coincidentally speaks a language that sounds like English, but uses words like 'conscious' and 'person'—and for that matter 'table'—quite differently from us. They unhesitatingly accept sentences like 'Wherever there is one {table/person/conscious thing/rights-bearer}, there are innumerable many other {tables/people/conscious things/rights-

⁵³We are thus opposed to views on which objectual attitudes can be read off in some simple way from propositional attitudes (including preferences). The combination of plenitude with Intensionalism makes the obstacles to such reading-off particularly vivid. There is an object that necessarily coincides with the fusion of all electrons if John is happy and doesn't concretely exist otherwise; but one could *care whether* John is happy without *caring about* this weird object. But by Intensionalism, the proposition that John is happy *is* the proposition that this weird object concretely exists.

bearers} with slightly different modal spatiotemporal profiles'. The Pluraliens' language doesn't even have the first- or second-person singular: when we would say we would say 'I am hungry', 'I love you', or 'Eight people admired one table', they say 'We are hungry', 'Each of us loves each of you', or 'There are eight groups of largely-overlapping people and one group of largely-overlapping tables such that each person in any the former groups admired each table in the latter group'.⁵⁴ While the Pluraliens' practice may have some practical disadvantages, we doubt that they would be making any *mistakes*: the properties expressed by 'person' and 'conscious' in their mouths really do have a vast multiplicity of instances.⁵⁵ And we think that they should return the compliment by recognizing that we are also not making any systematic mistakes. They should understand that the workings of vagueness mean that no "semantic miracle" is needed for a community to use 'person' in such a way that 'No ordinary person coincides with any other person' is true on all precisifications. Consider, by contrast, a possible isolated community of humans, the Humanists, who speak an English-like language but are acculturated to treat 'Non-human animals are not conscious and lack any moral status' as if it expressed an obvious truth, analogous to the way we treat 'Stones are not conscious and lack any moral status'. Suppose they also act accordingly, e.g. not caring (except instrumentally) about whether dogs are trapped in burning buildings. We are inclined to think that the Humanists *would* be making a grievous mistake: 'No non-human animals are conscious' and 'No non-human animals have any moral status' are false in their mouths. And if they are sticking to their guns, they should be equally uncharitable to us, thinking that we are mistaken when we say things like 'Cats and dogs feel pain'.⁵⁶

⁵⁴They might find it useful to have a specialized form 'wehere' of the first-person singular as in Johnston 2017 (633–4), to do some of the communicative work that we do by having both the first-person singular and the first-person plural to choose from.

⁵⁵Assume that the Pluraliens are quite similar to us as regards the actions to which they are willing to apply 'morally permissible'. Given this, it's natural to imagine that systematic ethical theory (at least of a consequentialist sort) is less popular among them than among us, and moral particularism more popular. However, Pluralien ethicists with a theoretical bent might be interested in the research program about 'volume in similarity space' we mentioned in note 52.

⁵⁶On a simple 'use plus eligibility' view of the role of naturalness in metasemantics (Lewis 1999b), one might think that the impropriety of charity to the Humanists would have to be underwritten by ascribing a high degree of naturalness to the property they express with 'conscious', since only a high degree of naturalness could overwhelm the "use" advantages of a charitable interpretation. We think this goes wrong in a variety of ways that already emerge in connection with ordinary natural-kind words like 'rabbit'. Each of the precisifications of 'rabbit' is highly unnatural, and would be very hard to refer to de-

9 Plasticity

To our minds, the most promising avenue for trying to derive surprising ethical implications from metaphysical views that make persons “ontological trash” is one that invokes *metasemantic* considerations in a way that hasn’t come up so far in this paper. The key idea is that properties and objects that are “ontologically on a par” are also *metasemantically* on a par, in the sense that there are similar as regards how easy it is to refer to them. To use the memorable terminology of Hodes (1984), none of the many metaphysically similar objects in the vicinity of a given person (including the person themselves) is especially more of a *reference magnet* than the others. Likewise, none of the many metaphysically similar *properties* in the vicinity of personhood and consciousness is especially magnetic. A consequence of this claim of metasemantic parity (see Dorr and Hawthorne 2014) is that uses of proper names, personal pronouns, and predicates like ‘person’ and ‘conscious’ are highly *semantically plastic*: for any given such use, there is a large multitude of similar objects or properties each of which either is, or could very easily have been, referred to by that use. On a ‘monistic’ way of developing this idea (Williamson 1994), any given use only refers to a single object or property, but there are tiny shifts in the underlying facts which would induce a shift in reference. On a ‘pluralistic’ version (Dorr and Hawthorne 2014; Dorr, unpublished), a single use at the actual world already refers to a vast multiplicity of objects or properties.⁵⁷

It doesn’t go without saying that metaphysical parity goes hand in hand with metasemantic parity and hence semantic plasticity—there are interesting views that posit a major contrast between a notion of naturalness made for metasemantics and one more closely tied to fundamental metaphysics (see Dorr and Hawthorne 2013). But for the purposes of this paper, let’s assume that persons and personhood are metasemantically undistinguished, so that our devices for referring to them are all highly plastic, and see whether this leads to a problematic multiplication of beings with moral status.⁵⁸

terminately (or uniquely); nevertheless, it is easy for a community exposed to rabbits to evolve a vague word that has rabbithood as one of its precisifications, even if they are prone to systematic errors in the application of the word. These reflections do not preclude more sophisticated accounts of the metasemantic role of naturalness, but this is not the right venue to pursue the matter (see Dorr and Hawthorne 2013; Dorr 2019).

⁵⁷Pluralism provides an especially straightforward interpretation of the technical terminology characteristic of “supervaluationism”; but the abstract formalism of supervaluationism can also be deployed in ways compatible with monism.

⁵⁸Even if it yields no such multiplication, the semantic plasticity of person-theoretic

One not-so-promising argument is based on the premise that ‘moral status’ — or some more specific predicate like ‘bearer of rights’ — is not semantically plastic. It is natural to think that ‘All people have moral status’ is “supertrue”, in the sense that all of the propositions that it expresses or could easily have expressed are true. After all, if it could easily express a falsehood, that would raise serious doubts about our ability to *know* that all people have moral status.⁵⁹ But if ‘person’ is plastic and ‘moral status’ is not, the only way ‘All people have moral status’ could be supertrue is for all of the objects that have *any* of the properties that are or could easily have been expressed by ‘person’ to have moral status. And on the view we are discussing, there is a huge multiplicity of such objects in the vicinity of any person: each candidate referent of ‘I’ in the mouth of that person instantiates some candidate referent of ‘person’ in the mouth of that person, since ‘I am a person’ is also stably true. So assuming that ‘moral status’ is not semantically plastic, each person must belong to a large population of similar, largely overlapping objects that all enjoy moral status.⁶⁰

The problem with this line of argument is that it is unclear what would motivate denying the plasticity for ‘moral status’ in a setting where one has granted that ‘person’ is plastic. There is strong pressure to say that a wide variety of mental and agential predicates (including ‘conscious’, ‘is in

vocabulary may generate other interesting ethical consequences. For example, in Dorr, Hawthorne, and Yli-Vakkuri 2021 (§13.3) we argue that such plasticity makes trouble for a principle suggestively named ‘Death Is Bad’. In Dorr and Hawthorne 2024, we raise similar concerns about widely accepted ‘Pareto’ principles in population axiology.

⁵⁹One could consider a view on which, although most of the propositions that ‘All people have moral status’ could easily have expressed are in fact false, this sentence could not easily have expressed a falsehood, since at each nearby world, it expresses a proposition that is true at that world but false at many neighbouring worlds. (Cf. Yablo’s (forthcoming) similar technique for combining plasticity for personal pronouns with non-plasticity for ‘person’, discussed in Dorr, Hawthorne, and Yli-Vakkuri 2021 (§11.5).) This would go with a picture on which the tiny underlying differences that would induce shifts in the reference of personal pronouns would also prevent us from enjoying moral status. Although this arguably avoids problems about knowledge, it is extremely revisionary to think that our moral status is as fragile as it turns out to be on this view: for example, it is hard to see how we could be confident that we would still have had moral status if we had eaten different foods for breakfast. For this reason we will ignore views where moral status is highly discriminating but highly fragile.

⁶⁰Olson (2010: 267-8) seems to be making something like this argument when he suggests that if the denotation of ‘person’ is determined by ‘the vagaries of linguistic behaviour’, and if the bearers of moral status do not (even temporarily) coincide, that raises a worry about our ability to know that people have moral status. While Olson doesn’t explicitly say that the relevant ethical predicates are non-plastic, it is hard to pinpoint any epistemological worry without that assumption.

pain', 'is an agent', 'is a thinker', etc.) are semantically plastic in a way that is "penumbally connected" with the plasticity of 'person', so as to make sentences like 'Ordinary people do not overlap any conscious beings distinct from themselves' supertrue. And a wide variety of ethical predicates seem in turn to be obviously penumbally connected to these mental and agential predicates: for example, 'Everything that has duties is an agent' is surely supertrue, so 'duty' is going to have to be plastic if 'agent' is. So one would need some very special reason to think that 'moral status' (or 'bearer of rights' or whatever) is not caught up in the same network of penumbral connections.

Some may worry that in conceding the plasticity of 'moral status', we will be forced to give up on a certain popular "moral realist" vision on which interpretative charity is especially problematic in the domain of ethics. For moral realists of this sort, communities that have words that play a similar "action-guiding" role to some of our ethical words, but apply those words very differently—e.g., accepting the sentence analogous to 'Slaves have no moral rights'—should be interpreted as making a mistake, rather than as meaning something different that would make the sentences they are disposed to accept come out true. But, first, this view doesn't have to be combined with a denial of semantic plasticity for any ethical vocabulary: there could be a large cluster of meanings such that using 'moral status' in a certain action-guiding way is enough to give it a referent within that cluster, although the selection of a specific member of the cluster depends on the same delicate factors that make for plasticity in 'person'. If so, the truth of 'Slaves have moral status' could be so robust that even counterfactual communities that are stably disposed to reject it would be making a mistake.⁶¹ And second, the realist still could posit that the meanings of certain ethical expressions are reference magnets, so long as the expressions in question are less tightly tied to 'person', 'agent', and so on. For example, they might think that 'better than', when used as an ethical predicate attaching to possible worlds or states of affairs, expresses a highly magnetic relation and is therefore non-plastic.

For these reasons, the argument based on the non-plasticity of 'moral status' does not look very threatening. But there is a different kind of ar-

⁶¹Moreover, even if there are many properties that could easily have been expressed by 'has moral status', it doesn't immediately follow that there are many propositions that could easily have been expressed by 'Slaves have moral status': the inference from the multiplicity of properties to the multiplicity of propositions seems to depend on a "structured" theory of propositions that is inconsistent given our favoured logic (Dorr 2016: §5). Indeed, on an Intensionalist view of propositional identity, it is fairly plausible that if 'Slaves have moral status' is stably true, it is stably such as to express the one and only necessarily truth.

gument from plasticity in words like ‘person’ to a multiplication of moral status that we find much more intriguing. This argument requires a “monistic” account of the relevant plasticity, where any given use of ‘person’, ‘conscious’, ‘moral status’, proper names, personal pronouns, etc. only refers to one property or object. Assuming this, the practice of referring to people using personal pronouns is extremely fragile—among worlds that differ at most slightly from the actual world in metasemantically relevant respects, most uses of personal pronouns do not refer to people at all. If the semantic relation of *reference* is fragile in this way, it seems plausible that the mental relation of *attention* is fragile in a parallel way. It would be odd to think that people at most nearby worlds never *refer* to their children, but nevertheless often *pay cognitive attention* to them. Once we go this far, it is hard to avoid saying a parallel thing about ethically significant relations like *love* and *respect*. It seems absurd to think that although the chance that we would ever pay any attention to our children was low, the chance that we would love them was still high. But it also seems wild to think that the actual world is very rare in its modal neighbourhood in being such that we love and respect things that are *fitting* objects of love and respect. That would be bizarrely self-congratulatory: the many chance events that, according to the view in question, had to turn out in highly specific ways for us to end up bearing these attitudes towards people were not all narrow escapes from moral tragedy.⁶²

By comparison with such self-congratulation, it seems better to think that all the things we could easily have loved or respected, and would have loved or respected if not for the vicissitudes that made ‘person’ mean exactly what it in fact means, are also fitting objects of love or respect. But this already seems like one important kind of “moral status”. And once we have got to the point of multiplying this particular form of moral status, we will then face further arguments for multiplying other forms, based on plausible connections between these and being a fitting object of love or respect. For example, it might be argued that everything that’s a *fitting object* of respect is also *deserving* of some minimal level of respect, and thus *has a right* to such respect, and hence a right against certain kinds of *treatment* incompatible with such respect.

For our money, this argument from semantic plasticity is the most com-

⁶²One might try to take the sting out of such self-congratulation by going for a picture on which, thanks to semantic plasticity, the sentence ‘Although the things we love at the actual world are typically fitting objects of love, most nearby worlds are not like this’ is robustly such as to express a truth. But while this avoids the need to posit widespread *epistemic* tragedy at nearby worlds, the *moral* tragedy of ill-fitting attitudes seems bad enough.

elling route from the claim that people are “ontological trash” to a multiplication of moral status. Would this kind of multiplication lead to bizarre conclusions about which actions are morally permissible? The answer depends in part on how much modal robustness we want to attribute to the fact that the things we love and respect are generally fitting objects of these attitudes. A minimal desideratum is that at no time in human history was there a high chance that this fact would fail to obtain. Owing to plasticity, this will still generate a multiplication of moral status to an abundance of things differing only slightly from people in their modal profiles, and thus force us to confront whatever challenges are raised by cardinality issues. However, as we discussed in section 8, these challenges are not obviously fatal. The multiplication may also have some worrisome consequences if it turns out that was a substantial chance of usage developing in such a way that ‘person’ would have applied to objects with markedly different spatiotemporal profiles. For example, if someone is reading this paper in a future after the invention of some radical technology like teletransportation, the plasticity-motivated multiplication would motivate investigating whether at the time of the invention there was a substantial chance of the community coalescing either around acceptance of ‘This new technology has enabled a new form of survival’ or of ‘This new technology has enabled a new form of death’. If it turns out that the reaction was chancy in this way, the multiplication might support a refusal to use the technology, even on the assumption that the survival sentence is true.

Stronger forms of robustness will of course have more radical implications for action. For example, it is plausible that there are possible situations in which people use ‘person’ in a way that makes ‘No person could ever survive falling into a coma’ true in their mouths—where we would truly say ‘She woke up from the coma’, they truly say ‘A new person came into being’—although it is implausible that there was ever a substantial chance that we would talk like this. If we want to say that even this unlikely scenario would not have involved unfitting love, then there is pressure to accord moral status to person-like objects for which comas spell doom. This will *prima facie* open the door to Johnston-style arguments for the impermissibility of incurring costs before a coma that will only be compensated for by benefits after the coma. However, it is not so clear how to run these arguments if, as we think, no non-person overlapping a person has mental states. If the respect-worthy object that will end with the coma will not *suffer* (since suffering involves mental states), it is unclear what costs we would be imposing on it by volunteering for painful procedures before the coma. And anyway, it might not be so bad to think that membership in such a weird

and unlikely linguistic community leads to systematically misdirected love and respect.

So far, then, we have found no terrible consequences of the plasticity-driven multiplication of moral status as regards which *actions* are morally permissible. Nevertheless, the package does seem in other ways to be quite problematic. Unless we jettison our commitment to the idea that there is only one conscious being in the vicinity of an ordinary person, the package implies will imply there are things that are fitting objects of love and respect despite lacking any capacity for consciousness or any other mental states. This is strange! While of course there is a sense in which we can love even things like mountains and ice cream flavours, 'love' also seems to be capable of expressing a more demanding and morally significant relation, such that it is hard to even understand what it could mean to bear that relation to something that one does not take to be even metaphysically capable of thinking or feeling.⁶³ The package also risks quite pervasive moral tragedies at the level of feeling. If the relevant moral status includes *deserving* or *having a right to* some minimal level of respect (from at least someone or other), but the *respect* relation is too demanding for finite beings like us to bear it to more than a few objects within any cluster of largely overlapping objects, then we seem to be inevitably saddled with a situation in which myriad objects will never receive the respect which they deserve, or to which they have a right. Even if this doesn't lead to many wrongful actions, it is still strange and disturbing to think of ourselves as inevitably falling short in this way.

The plasticity-motivated multiplication of moral status is thus not easy to shrug off. But it bears emphasis that the argument for this multiplication depended crucially on the monistic version of plasticity. As we mentioned, our own preference is for a pluralist form of plasticity, on which uses of the relevant expressions simultaneously refer to many candidate meanings, even in the actual world. Given pluralism, it becomes much easier to sustain the idea that people are "love magnets" despite not being reference magnets. We could say that it is impossible, or at least very difficult, for anyone to love or respect any object that has never had any capacity for mental states.⁶⁴ *Love*

⁶³While we don't talk about respecting ice-cream flavours, we do talk about people respecting things like nature, traditions, and so on. But here again, it seems plausible that 'respect' also has a narrower, more "interpersonal" interpretation.

⁶⁴Thus for example, we are inclined to think that even the Pluraliens from the previous section only love and respect other people, though of course they bear the relations *they* express by 'love' and 'respect' to many non-people. Here, we are of course attempting to focus on the relatively demanding uses of 'love' and 'respect', setting aside the uses in play in 'He loves chocolate' and 'She respects tradition'. One motivation for saying 'very difficult'

and *respect* could in this regard work like relations like *child of* and *employer of*, that arguably necessarily hold only between people and animals, and certainly never hold between a person and a non-person that largely overlaps a person.⁶⁵ Of course, given the plasticity of ‘person’, this will go hand in hand with the idea that ‘love’ and ‘respect’ also express a multitude of candidate meanings, co-ordinated with the meanings of ‘person’ in such a way as to make ‘People are love and respect magnets’ true on all candidate interpretations.

Given this pluralist package, the argument that non-people are *fitting* objects of love and respect breaks down. We can, in a principled way, sustain the thought that on all relevant definitions of ‘has moral status’, this predicate is caught up in the cluster of penumbrally-connected expressions that includes ‘person’, in such a way that what Johnston’s “principle of ethical singularity”—according to which “The only being with a moral status to be found within a person’s spatio-temporal envelope is that person” (Johnston 2017: 628)—could not easily have expressed a falsehood.

Our conclusion, then, is that while the moral considerations that support the principle of ethical singularity do not provide a strong argument against metaphysical views on which people are “ontological trash”, these considerations do constitute a surprising argument that if any such metaphysical view is true, semantic pluralism is true as well.⁶⁶

References

- Arntzenius, Frank (2014). ‘Utilitarianism, Decision Theory and Eternity’, *Philosophical Perspectives*, 28/1: 31–58.
- Arntzenius, Frank, and Dorr, Cian (2011). ‘Calculus as Geometry’, in Frank Arntzenius, *Space, Time, and Stuff* (Oxford: Oxford University Press).
- (2017). ‘Self-Locating Priors and Cosmological Measures’, in Khalil Chamcham, John Barrow, Simon Saunders, and Joe Silk (eds.), *The Philosophy of Cosmology* (Cambridge: Cambridge University Press), 396–428.

rather than ‘impossible’ is the thought that one could love or respect something—maybe a sophisticated puppet—that one mistakenly believes to have the capacity for mental states.

⁶⁵This view of love and respect could be combined either with a view of *attention* as similar to reference (in that typically when one attends to a person one attends to many other things in the vicinity), or as similar to love and respect (in that it is vastly easier to attend to people than to the non-people that largely overlap them). A parallel choice arises for many other object-directed attitudes: see for example the discussion of *seeing* in Dorr, Hawthorne, and Yli-Vakkuri 2021 (§12.4).

⁶⁶Thanks to Brian Hedden, Jake Nebel, Ted Sider, and Collis Tahzib, who provided helpful feedback on earlier versions of this material.

- Bacon, Andrew (2020). 'Logical Combinatorialism', *The Philosophical Review*, 129/4: 537–89.
- Bader, Ralf M. (2013). 'Towards a Hyperintensional Theory of Intrinsicity', *Journal of Philosophy*, 110/10: 525–63.
- Builes, David, and Hare, Caspar (2023). 'Why Aren't I Part of a Whale?', *Analysis*, 83/2: 227–34.
- Chalmers, David J. (2004). 'The Representational Character of Experience', in Brian Leiter (ed.), *The Future for Philosophy* (Oxford: Oxford University Press), 153–81.
- Cutter, Brian (forthcoming). 'The Many-Subjects Argument Against Physicalism', in Geoffrey Lee and Adam Pautz (eds.), *The Importance of Being Conscious* (Oxford University Press).
- Denby, David (2014). 'Essence and Intrinsicity', in Robert M. Francescotti (ed.), *Companion to Intrinsic Properties* (De Gruyter), 87–109.
- Dorr, Cian (2016). 'To Be F Is To Be G', in John Hawthorne and Jason Turner (eds.), *Philosophical Perspectives 30: Metaphysics* (Oxford: Blackwell), 1–97.
- (2019). 'Natural Properties', *Stanford Encyclopedia of Philosophy*, Fall 2019, ed. Edward N. Zalta, <https://plato.stanford.edu/entries/natural-properties/>.
- (unpublished). 'The Multiplicity of Meaning', The Hempel Lectures, delivered 2020.
- Dorr, Cian, and Hawthorne, John (2013). 'Naturalness', in Karen Bennett and Dean Zimmerman (eds.), *Oxford Studies in Metaphysics*, viii (Oxford: Oxford University Press), 3–77.
- (2014). 'Semantic Plasticity and Speech Reports', *Philosophical Review*, 123: 281–338.
- (2024). 'Pareto Principles and the Shallowness of Personal Identity' (unpublished).
- Dorr, Cian, Hawthorne, John, and Yli-Vakkuri, Juhani (2021). *The Bounds of Possibility: Puzzles of Modal Variation* (Oxford: Oxford University Press).
- Dretske, Fred (1996). 'Phenomenal Externalism, or If Meanings Ain't in the Head, Where Are Qualia?', *Philosophical Issues*, 7: 143–58.
- Eddon, Maya (2011). 'Intrinsicity and Hyperintensionality', *Philosophy and Phenomenological Research*, 82: 314–36.
- Fine, Kit (1994). 'Essence and Modality', in James Tomberlin (ed.), *Philosophical Perspectives 8: Logic and Language* (Atascadero: Ridgeview), 1–16.
- (2000). 'A Counterexample to Locke's Thesis', *The Monist*, 83/3: 357–61.
- (2003). 'The Non-Identity of a Material Thing and Its Matter', *Mind*, 112: 195–234.

- Fine, Kit (2005). 'Necessity and Non-existence', in id., *Modality and Tense: Philosophical Papers* (Oxford: Oxford University Press), 321–57.
- Francescotti, Robert (1999). 'How to Define Intrinsic Properties', *Noûs*, 33/4: 590–609.
- Gibbard, Allan (1975). 'Contingent Identity', *Journal of Philosophical Logic*, 4: 187–221.
- Harman, Elizabeth (1999). 'Creation Ethics: The Moral Status of Early Fetuses and the Ethics of Abortion', *Philosophy and Public Affairs*, 28/4: 310–24.
- Harsanyi, John C. (1955). 'Cardinal Welfare, Individualistic Ethics, and Interpersonal Comparisons of Utility', *Journal of Political Economy*, 63/4: 309–21.
- Hawthorne, John, and Russell, Jeffrey Sanford (2018). 'Possible Patterns', in Karen Bennett and Dean W. Zimmerman (eds.), *Oxford Studies in Metaphysics*, xi (Oxford: Oxford University Press), 149–92.
- Hodes, Harold (1984). 'Logicism and the Ontological Commitments of Arithmetic', *Journal of Philosophy*, 81: 123–49.
- Hoffmann-Kolss, Vera (2014). 'Is the Intrinsic/Extrinsic Distinction Hyperintensional?', in Robert M. Francescotti (ed.), *Companion to Intrinsic Properties* (De Gruyter), 157–73.
- Humberstone, Lloyd (1996). 'Intrinsic/Extrinsic', *Synthese*, 108: 205–67.
- Johnston, Mark (2010). *Surviving Death* (Princeton: Princeton University Press).
- (2016). 'Personites, Maximality And Ontological Trash', *Philosophical Perspectives*, 30: 198–228.
- (2017). 'The Personite Problem: Should Practical Reason Be Tabled?', *Noûs*, 51: 617–44.
- Kagan, Shelly, and Vallentyne, Peter (1997). 'Infinite Value and Finitely Additive Value Theory', *Journal of Philosophy*, 94: 5–26.
- Langton, Rae, and Lewis, David (1998). 'Defining "Intrinsic"', *Philosophy and Phenomenological Research*, 58: 333–45.
- Lewis, David K. (1968). 'Counterpart Theory and Quantified Modal Logic', in id., *Philosophical Papers*, i (Oxford: Oxford University Press, 1983), 26–46. From *Journal of Philosophy*, 65 (1968): 113–26.
- (1971). 'Counterparts of Persons and their Bodies', in id., *Philosophical Papers*, i (Oxford: Oxford University Press, 1983), 47–54. From *Journal of Philosophy*, 68 (1971): 203–11.
- (1976). 'Survival and Identity', in id., *Philosophical Papers*, ii (Oxford: Oxford University Press, 1986), 55–77. From Amelie Oksenberg Rorty (ed.), *The Identities of Persons* (University of California Press, 1976), 17–40.
- (1983). *Philosophical Papers* (Oxford: Oxford University Press).
- (1986). *On the Plurality of Worlds* (Oxford: Blackwell).

- Lewis, David K. (1993). 'Many, But Almost One', in id., *Papers in Metaphysics and Epistemology* (Cambridge: Cambridge University Press, 1999), 164–82. From Keith Campbell, John Bacon, and Lloyd Reinhardt (eds.), *Ontology, Causality, and Mind: Essays on the Philosophy of D. M. Armstrong* (Cambridge: Cambridge University Press, 1993).
- (1999a). *Papers in Metaphysics and Epistemology* (Cambridge: Cambridge University Press).
- (1999b). 'Putnam's Paradox', in id., *Papers in Metaphysics and Epistemology* (Cambridge: Cambridge University Press), 56–7. From *Australasian Journal of Philosophy*, 62 (1984): 221–36.
- Lycan, William G. (2001). 'The Case for Phenomenal Externalism', *Philosophical Perspectives*, 15: 17–35, <http://www.jstor.org/stable/2676166>, accessed 10 July 2024.
- Marshall, Dan, and Weatherson, Brian (2023). 'Intrinsic vs. Extrinsic Properties', *The Stanford Encyclopedia of Philosophy*, Fall 2023, ed. Edward N. Zalta and Uri Nodelman, <https://plato.stanford.edu/archives/fall2023/entries/intrinsic-extrinsic/>.
- Merricks, Trenton (2001). *Objects and Persons* (Oxford: Oxford University Press).
- Olson, Eric T. (2010). 'Ethics and the Generous Ontology', *Theoretical Medicine and Bioethics*, 31/4: 259–70.
- Parfit, Derek (1986). *Reasons and Persons* (Oxford: Oxford University Press).
- Pautz, Adam (2013). 'The Real Trouble for Phenomenal Externalists: New Empirical Evidence', in Richard Brown (ed.), *Consciousness Inside and Out: Phenomenology, Neuroscience, and the Nature of Experience* (Dordrecht: Springer Studies in Brain and Mind), 237–98.
- Sider, Theodore (1996). 'Intrinsic Properties', *Philosophical Studies*, 83/1: 1–27.
- (2001). 'Maximality and Intrinsic Properties', *Philosophy and Phenomenological Research*, 63: 357–64.
- (2003). 'Maximality and Microphysical Supervenience', *Philosophy and Phenomenological Research*, 66: 139–49.
- Skow, Bradford (2007). 'Are Shapes Intrinsic?', *Philosophical Studies*, 13: 111–30.
- Unger, Peter (1979). 'I Do Not Exist', in G. F. MacDonald (ed.), *Perception and Identity* (London: MacMillan), 235–51.
- Williamson, Timothy (1994). *Vagueness* (London: Routledge).
- Yablo, Stephen (forthcoming). 'Everything Exists: Plenitude, Kinds, and Paradox', in Javier Cumpa (ed.), *The Question of Ontology: The Contemporary Debate* (Oxford: Oxford University Press).

Zimmerman, Dean (2003). 'Material People', in Michael J. Loux and Dean W. Zimmerman (eds.) (Oxford: Oxford University Press).