Philosophers' Imprint

volume 16, no. 1
january 2016

LOGICALLY SIMPLE PROPERTIES AND RELATIONS

Jan Plate

Université de Neuchâtel

1. Introduction

Metaphysicians generally agree that not all predicates are created equal. In the *Parmenides*, young Socrates affirms that there are Forms of the beautiful, the just, and the good, but denies that there is a Form of hair or of mud. In classical Indian metaphysics, Udayana's followers distinguished 'real' universals (*jāti*) from those that are merely 'constructed' (*upādhi*).[1] And in recent Western philosophy, Goodman (1955) has distinguished projectible from non-projectible predicates, Armstrong (1978) predicates that correspond to universals from those that don't, Shoemaker (1980) genuine from 'mere Cambridge' properties, and David Lewis (1983; 1986) perfectly natural attributes from those that are less than perfectly natural.[2] It is easy to notice that in each of these distinctions, one of the two respective classes of predicates (or universals, etc.) is in some way privileged. The distinction that this paper is concerned with also fits into this roster. I shall try to give an account of it, and offer the beginnings of an argument for the view that it would make for a viable substitute for Lewis's distinction between the perfectly and the less-than-perfectly natural attributes.

Although Lewis's writings contain various remarks designed to explain what it is for an attribute to be perfectly natural, he has provided no official analysis of this notion, and was apparently willing to treat it as a primitive.[3] Those parts of Lewis's philosophy that make use of

1. Bartley (2001, 543).
2. For the sake of brevity, the term 'attribute' is here (following Carnap [1942]) used interchangeably with 'property or relation', practically regardless of context. So I shall speak of 'Lewis's concept of a perfectly natural attribute', irrespective of the fact that (i) Lewis himself tends not to use the word 'attribute', and (ii) the conception of attributes operative in this paper is markedly different from Lewis's.
3. Cf. Lewis (1983, 347), (1986, 63). For Lewis's attempts at explicating perfect naturalness, see, *e.g.*, his (1983, 346f.), (1986, 60f.), and (2009, 204). As for taking the notion as a primitive: it is of course possible to take instead the concept of *comparative* naturalness as primitive, and then to say that any given property or relation is perfectly natural just in case no attribute is more natural than it. However, the difference between the two approaches will not matter much for our purposes.

the notion of a perfectly natural attribute are thereby rendered vulnerable to the charge of resting on an esoteric and even obscure concept.[4] By contrast, there is no reason whatsoever to take the concept of a *logically simple* attribute as primitive. Since the account that will here be given of the distinction between logically simple and complex attributes amounts to an analysis (at least on one understanding of this term), it will be possible to make non-trivial general claims about logically simple attributes that can be regarded as *analytically true*, rather than to be based on speculation and/or complex 'job descriptions'—*i.e.*, descriptions of the different kinds of theoretical work that the respective concept is supposed to do—that are only questionably satisfied. Inevitably, the account will make use of some primitive notions of its own, such as those of state of affairs, negation, and conjunction. But I think that these will not be particularly hard to grasp. Like the concept of parthood (and despite their technical-sounding names), they seem to have their origin not so much in arcane metaphysics as in fairly commonplace thought and talk.[5]

In the first place, then, the concept of a logically simple attribute is arguably less obscure than Lewis's concept of perfect naturalness. In the second place, it promises to be able to do some of the theoretical work that Lewis has assigned to the latter concept. One of the most salient examples of this sort of work is provided by the 'best system' account of lawhood, as presented in Lewis's 'New Work' and elsewhere.[6]

As is well known, the basic idea of this account is to identify the laws of nature with the regularities that earn "inclusion in the ideal system" (*op. cit.*, p. 367), a set of true sentences that is closed under strict implication and which best balances strength (*i.e.*, information content) and simplicity of axiomatization. As is also well known, this basic idea faces the objection that maximum strength can be achieved by a very simple axiom '$\forall x\, Fx$' if the predicate 'F' picks out some condition that is (i) satisfied by all entities in the actual world and (ii) *not* satisfied by all entities in any other possible world.[7] Lewis's solution to this problem is to impose a constraint on the primitive vocabulary of the systems in question: they must be formulated in such a way that each undefined predicate corresponds to a perfectly natural attribute.[8]

I want to suggest that this role, which is in Lewis's account played by the notion of a perfectly natural attribute, can be profitably assigned, instead, to the concept of a logically simple attribute. (For the sake of brevity, the qualifier 'logically' will in the following often be omitted.) At least three reasons render such a move worth considering: First,

4. Cf. Taylor (1993) and Witmer, Butchard & Trogdon (2005, 329). The discussions in Schaffer (2004), Ainsworth (2009, §6), and Dorr & Hawthorne (2013) are also relevant.

5. This seems fair to say, at any rate, if talk about 'situations', 'goings-on', 'traits', 'features', etc., qualifies as 'fairly commonplace'. As Russell (1918/19, §3) discovered, people may not *like* to admit into their ontologies states of affairs that are negations or disjunctions of other states of affairs. But this does not mean that talk of, *e.g.*, the negation of a state of affairs is unintelligible. If anything, it seems to support the contrary claim.

6. The account has roots in the work of Mill (1882, bk. 3, ch. 4) and Ramsey (1928). Lewis's first version of it is given in his (1973, 73), without the refinements introduced in 'New Work' and 'Humean Supervenience Debugged'. The account has been criticized by, *e.g.*, Tooley (1977), Armstrong (1983), Carroll (1994), and Maudlin (2007). Generally sympathetic, though still critical, discussions can be found in, *e.g.*, Loewer (2007), Cohen & Callender (2009), and Eddon & Meacham (2015). The account has been endorsed by Sider (2011, 22f.).

7. In Lewis's own words (*ibid.*): "Given system S, let F be a predicate that applies to all and only things at worlds where S holds. Take F as primitive, and axiomatise S (or an equivalent thereof) by the single axiom $\forall x\, Fx$."
Hence one could let S be a system that holds only at the *actual* world (thereby making S maximally strong, as well as true), and then let F be a predicate that applies to all and only the things in the actual world. But evidently, if we are going to take '$\forall x Fx$' as our axiomatization of S, then we needn't be so restrictive in our choice of F: it is enough to choose F in such a way that it is true of all entities in the actual world and not true of all entities in any other world.

8. Cf. Lewis (1983, 367f.). An alternative solution, which Lewis also hints at, is to "replace strict implication by deducibility in some specified calculus" (p. 368). He then dismisses it on the ground that it "seems unnecessary given [the adopted solution], and seems incapable of solving our problem by itself". Whether or not this is so, Lewis's preferred solution certainly constitutes a natural approach, provided that an 'elite vocabulary' of the sort he envisions is in principle available.

as already mentioned, the concept of a logically simple attribute is amenable to analysis, which should alleviate any worries about the intelligibility of this concept. Second, the account to be offered in this paper will give the logically simple attributes a legitimate claim to be regarded as *fundamental*, owing in part to a strong analogy between that account and an adequate analysis of mereological simplicity. And third, the account will allow us to show (on the basis of the framework to be described in section 2) that the logically simple attributes are in at least one important way *sparse*. (We will see in the Conclusion why this matters.)

Apart from the analysis of lawhood, the concept of a logically simple attribute may also be usefully employed in the analysis of duplication. According to Lewis (1986, 61), two entities are duplicates iff "(1) they have exactly the same perfectly natural properties, and (2) their parts can be put into correspondence in such a way that corresponding parts have exactly the same perfectly natural properties, and stand in the same perfectly natural relations". Here again, the concept of a perfectly natural attribute can profitably be replaced with that of a logically simple attribute. For it is plausible to say (though I shall not argue for it in this paper) that simple attributes are both intrinsic and purely qualitative (*i.e.*, not 'impure'). Consequently, if two entities have exactly the same simple properties and their parts can be put into correspondence in such a way that corresponding parts have exactly the same simple properties and stand in the same simple relations, then those entities will be perfectly alike in their intrinsic qualitative profiles, which in turn makes it plausible to regard them as duplicates.

It is less clear whether the concept of a logically simple attribute can replace that of a perfectly natural attribute also in all other applications that Lewis has proposed for the latter. But this need not necessarily speak against the usefulness of our concept, since those proposed applications are not always uncontroversial. It is questionable, for instance, whether the problem of drawing the distinction between intrinsic and extrinsic properties is best addressed by relying (as in Langton & Lewis [1998] and Lewis [2001]) on the notion of naturalness to sort the 'disjunctive' properties from the non-disjunctive ones. And at any rate, if the concept of a simple attribute can fill the roles that the concept of a perfectly natural attribute plays in Lewis's accounts of lawhood and duplication, it will thereby already be able to do a considerable amount of important philosophical work—as witness, for example, the crucial place that the concept of duplication occupies in Lewis's (1983, 368f.) counterfactual account of causation. This should be enough to justify our interest in it.

The structure of this paper is as follows. Section 2 provides a sketch of the framework on which the account of logical simplicity will be based, and section 3 presents the account itself, together with some results that follow from it, in particular with regard to the sparseness of the simple attributes. The Conclusion then briefly reflects on the prospects of a best-system account of lawhood in which perfect naturalness is replaced with (logical) simplicity. Since the sketch of the framework in section 2 is going to form a vital part of the paper, it will be useful to add here a few remarks on the framework itself.

In a nutshell, the framework consists of a rudimentary ontology of attributes and states of affairs, together with a formal language whose main function will be to provide a way of constructing expressions by which attributes and states of affairs can be referred to, in such a way that these expressions reveal (to a greater or lesser extent) the 'logical structure' of their respective referents. What is meant by 'logical structure' will become clearer as we go, but it will be worth keeping in mind that it should not be thought of as anything *mereo*logical, except on a very broad conception of the part–whole relation.[9] Another noteworthy feature of the framework lies in the fact that it treats attributes and states of affairs as *sui generis*, rather than to reduce them, à la Montague (1969) or Lewis (1986), to set-theoretic constructions over a space of possible worlds. This departure from possibilist treatments is mainly motivated by familiar skepticism about possible worlds.[10] Instead of

9. Cf. Fine (2010).

10. In addition, an argument by Eddon (2011) suggests that *non-hyperintensional*

an abundance of possible worlds, the framework commits itself to an abundance of attributes and states of affairs, insofar as it admits states of affairs that are negations, conjunctions, etc., of other states of affairs, and allows for attributes to be rather freely 'generated' from states of affairs via λ-abstraction. No detailed defense of this ontology will be attempted in this paper, but the following two considerations may be worth flagging.

First, it is true that, *if* we could assume that attributes in general (as opposed to this or that subclass of them) are very sparse, so that no attribute is the negation of any other, or the conjunction of any two attributes, etc.,[11] then we might just say that any attribute whatsoever is logically simple, and spare ourselves any further trouble of working out an account of what it is for an attribute to be logically simple. But the assumption of an extremely sparse ontology of attributes is not uncontroversial, and so it should be at least somewhat worthwhile to see how we might make sense of the notion of a (logically) simple attribute on the assumption of an abundant ontology of attributes.[12]

conceptions of attributes, such as those of Lewis and Montague, suffer from the defect of conflating intrinsic with extrinsic properties. Her argument can be straightforwardly adapted to the case of simple vs. complex attributes. Thus, let P be any property that is had by some set-sized number of entities. On Lewis's conception, on which properties are sets of possibilia, P could then be identified with the 'disjunctive' property of *being identical with* x_1 *or with* x_2 *or…*, where the x_i are all the possible entities that have P. On Montague's conception, on which properties are functions from possible worlds to extensions, P could similarly be identified with the property of *being identical with* x_1^1 *or* x_1^2 *or … in* w_1 *or identical with* x_2^1 *or* x_2^2 *or … in* w_2 *or…*, where the w_i are all the worlds in which P has a non-empty extension and, for any given world w_i, the x_i^j are all the entities that have P in w_i. Since such disjunctive properties do intuitively not count as simple, it thus appears that a conception of attributes along the lines of Lewis or Montague forces us to treat all properties (or more precisely, all properties that are instantiated by a set-sized number of possibilia; this qualification is needed if we cannot have disjunctions of greater-than-set-sized length) as complex. This consequence may well be considered problematic.

11. See Dorr (2002, §4.2) for a fully explicit formulation of a sparseness claim of this sort.

12. Also see the related remarks in section 3.1 below.

Of course, this is not a defense of an abundant ontology itself, but rather only a way of motivating our *use* of such an ontology.

Second, with a view to actually defending an abundant ontology of attributes, one might appeal to a deflationist meta-ontology, possibly along the lines of Carnap (1950), Eklund (2008), or Thomasson (2015). However, rather than to adopt any particular meta-ontological proposal, I would here only like to venture a very brief remark on how *natural* it is to believe in the existence of logically complex attributes. We can probably all agree that it is extremely natural to believe that there are *mereo*logically complex things (*i.e.*, things with proper parts), and that in fact most of the physical objects that we consciously deal with in our everyday lives are mereologically complex. But now, if one accepts this view, and moreover allows that there are such things as properties and relations, then it should seem equally natural to hold the analogous view that there are logically complex attributes, and that in fact most of the attributes that we consciously deal with in our daily thought and talk are logically complex.[13] For, just as it would be highly implausible to think of our utterances of expressions like 'this desk', 'that watch', etc., as referring to mereologically simple things, so it would also be implausible to think of our utterances of expressions like 'desk', 'watch', 'black', 'water', etc., as having as their semantic values logically simple—*i.e.*, to a very first approximation: metaphysically unanalyzable—attributes.

If this last consideration is correct, it will be all the more worthwhile to have a clear grasp of what it is for an attribute to be (or to fail to be) logically complex.

13. The 'consciously' should be taken with a grain of salt: roughly, the phrase is here to be understood in the sense that one *consciously deals* with an attribute A at a time t iff one employs at t a concept that corresponds to A, in the sense of 'corresponds' in which, *e.g.*, the concept expressed by 'is a horse' corresponds to the property of *being a horse*. For example, if I think of this desk that it is black (making use of the—or *a*—concept of blackness), then I thereby consciously deal with the property of *being black*.

2. Sketch of a Framework

2.1 *Preliminaries*

The framework to be sketched in this section can be thought of as falling into two parts: on the one hand a rudimentary ontology of attributes and states of affairs, and on the other hand a system of expressions (a 'language') by which attributes and states of affairs can be referred to.

With regard to the ontology, the main thing to note is that the concept of a state of affairs plays in it an absolutely central role, since the framework's commitment to an abundance of attributes depends directly on its commitment to an abundance of states of affairs (from which attributes are 'generated' by way of λ-abstraction).[14] In contrast to the way in which states of affairs would be conceived of in a possibilist framework, they are here not regarded as set-theoretic constructions over a space of possible worlds. Nor are they in a straightforward sense taken to be 'structured' entities, if this were to mean that each state of affairs is constructed in one and only one way from the set of its basic constituents. Instead, to borrow a term from Hodes (1982), states of affairs are here conceived of as 'polymorphous'.[15]

With regard to the formal language, it is worth foreshadowing that we will use *formulas* as names of states of affairs,[16] while *λ-expressions* will be used as names of attributes. Thus, we may for instance use '$\lambda x\,(x = x)$' as a name for the property of self-identity, and write '(Socrates = Socrates)' to refer to the state of affairs that Socrates is Socrates. This goes for the formal language as well as for the 'metalanguage' in which this paper is written (*i.e.*, English, *modulo* certain additions). Accordingly, formulas and λ-expressions will here function as noun-phrases: each of them can appear as the subject of a sentence or as an object of a verb-phrase.[17] Something analogous holds for the formal language, in that formulas and λ-expressions are there allowed to appear 'in subject-position', *i.e.*, as elements of an argument-list.[18]

In the following subsection, I will first describe the formal language, focusing solely on its syntactic aspects. In section 2.3, I will turn to the semantics of this language and to the ontological assumptions presupposed by the semantics. In section 2.4, I will then state a sufficient condition for the identity of states of affairs. This condition will be weak enough to allow for relatively fine-grained distinctions among states of affairs (as well as attributes), so that the overall framework can be fairly regarded as 'hyperintensional', at least in spirit. (This last point will be briefly discussed in section 2.5.) Readers who are not interested in the technical details may wish to read only section 2.3.1 before skipping ahead to section 2.4; the concepts of interpretation and equivalence introduced there will be relevant throughout section 3.

2.2 *The Formal Language*

The system of expressions that we will use in order to refer to attributes and states of affairs is an infinitary version of the standard language of first-order logic, with the following stock of primitive symbols:

1. *Variables*: italicized letters (lower- or uppercase) of the Latin alphabet, with or without subscripts.

2. *Non-logical constants*: unitalicized one-word English nouns, singular verbs (*e.g.*, 'sleeps'), and adjectives, as well as hyphenated phrases (such as 'teacher-of'), with or without subscripts. Hyphenated phrases are introduced on an *ad hoc* basis.

14. The legitimacy of the notion of a state of affairs has been defended by a number of philosophers. One way of defending it is along the lines of Pollock (1967). To defuse potential worries concerning 'the slingshot argument', one could in addition draw on Oppy (1997). There is also, of course, the need to avoid certain paradoxes, on which see footnote 36 below.

15. For related discussion, see Swoyer (1998, 299f.) and Bynoe (2011).

16. Cf. Suszko (1975).

17. In consequence of this, when expressions of the formal language will be 'mentioned' rather than 'used', I shall not follow the logicians' convention of omitting quotation marks. Rather, single quotes will be used for 'literal' quotation (as well as scare quotes), and corner-quotes will be used for quasi-quotation.

18. For a more detailed definition, see p. 7 below.

3. The *logical constant* '$\mathtt{I}$', which is used to denote the identity relation.

4. The operators '$\neg$', '$\wedge$', '$\exists$', and 'λ'.

5. The parentheses '(' and ')', and the comma.

In addition, we adopt several abbreviatory devices. In particular, for any terms $t_1, t_2, \ldots$, the expression $\ulcorner(t_1 \vee t_2 \vee \ldots)\urcorner$ will abbreviate $\ulcorner\neg(\neg t_1 \wedge \neg t_2 \wedge \ldots)\urcorner$, $\ulcorner(t_1 \to t_2)\urcorner$ will abbreviate $\ulcorner\neg(t_1 \wedge \neg t_2)\urcorner$, $\ulcorner(t_1 \leftrightarrow t_2)\urcorner$ will abbreviate $\ulcorner(\neg(t_1 \wedge \neg t_2) \wedge \neg(\neg t_1 \wedge t_2))\urcorner$, and the formula $\ulcorner\neg\exists v_1, v_2, \ldots \neg\varphi\urcorner$ may be abbreviated as $\ulcorner\forall v_1, v_2, \ldots \varphi\urcorner$. Finally, the formula $\ulcorner\mathtt{I}(t_1, t_2)\urcorner$ may be written as $\ulcorner(t_1 = t_2)\urcorner$ and the formula $\ulcorner\neg\mathtt{I}(t_1, t_2)\urcorner$ as $\ulcorner(t_1 \neq t_2)\urcorner$. As usual, outermost parentheses may be omitted.

Since the language is meant to be free from arbitrary restrictions of expressive power, we allow among other things for infinitely long formulas and an inexhaustible supply of variables. 'Inexhaustible' means here that there are more variables than can be contained in any one formula. We achieve this by allowing *ordinal numbers* to play the role of subscripts. Thus, a subscript of a variable need not be conceived of as a *bona fide* symbol but can instead be thought of as an ordinal number.[19] (Of course, to *write down* a variable that has an ordinal as a subscript, one will still have to use some symbol or other to represent the subscript.) Since not just variables but also constants can have subscripts, we will apply the same policy to the subscripts of constants, and thereby help ourselves to an inexhaustible supply of constants as well. The underlying set theory for all this is ZFCU (Zermelo–Fraenkel set theory with choice and urelements), supplemented by an open-ended list of large-cardinal axioms.[20]

Before we proceed, it will be useful to introduce some terminology. First of all, talk of *occurrences* and *containment*, as applied to expressions, should be understood in a sufficiently liberal sense that every expression contains an occurrence of itself and thus, by extension, contains itself. Occurrences should be thought of as 'bound' to the expressions that contain them, in the sense that each occurrence is contained in exactly one expression. By contrast, an occurrence may be contained in more than one *occurrence*. Thus, the second occurrence of 'x' in the formula '$\exists x\, \text{loves}(\text{Sam}, x)$' is contained in only one expression—viz., that same formula—but is contained in *two* formula-occurrences, viz., first, the occurrence of '$\exists x\, \text{loves}(\text{Sam}, x)$' in itself and, second, the occurrence of '$\text{loves}(\text{Sam}, x)$' in the former.

The concept of a variable-occurrence's being *bound* by a given operator-occurrence can be defined in the usual way, provided it is kept in mind that both 'λ' and '$\exists$' (and only these) are variable-binding operators. A variable-occurrence will be said to be *bound simpliciter* iff it is bound by some operator-occurrence; otherwise it will be said to be *free*. In addition, a variable-occurrence will be said to be bound *in* a given expression or occurrence γ iff it is bound by some operator-occurrence within γ. On this basis, the more general concept of a free term-occurrence can be introduced as follows: If o is an occurrence of a term and γ an expression—or occurrence of an expression—that contains o, then o is *bound* or *free in* γ according as o does or does not contain a variable-occurrence that is free in o but bound in γ. In the special case in which γ is not just an occurrence but rather the unique expression that contains o, we will say that o is bound or free *simpliciter*. For example, the first but not the second occurrence of '$F(x)$' in '$F(x) \wedge \exists x\, F(x)$' is free. Another consequence is that occurrences of constants are always free.

19. In fact, one can go even further and think of *all* expressions of the formal language as pure sets, in the manner of Gödel codes. The original Gödel codes—natural numbers used to represent expressions in a finitary language of arithmetic—would of course be far too few to serve an analogous purpose in the case of the present infinitary language. The pure sets present a natural alternative, since, for each set-sized cardinality κ, there are at least κ-many of them.

20. In addition, we will occasionally make use of the informal notion of a *class*. For more on this, see, *e.g.*, Jech (2002, 5f.).

We will further say that an occurrence of a term stands *at subject-position* just in case it is an element of an *argument-list*, *i.e.*, of a parenthesis-enclosed, comma-delimited list of one or more term-occurrences that is preceded by an occurrence of an atomic term or λ-expression.[21] For example, in the formula 'loves(x, y)', the variables 'x' and 'y' both occur at subject-position, but the constant 'loves' doesn't. The complementary concept is that of predicate-position: an occurrence of a term stands *at predicate-position* just in case it immediately precedes an argument-list.

The notions of term, formula, and λ-expression can be recursively introduced as follows:

(T) A *term* is anything that is a variable, constant, formula, or λ-expression. (In addition, we call a term *atomic* iff it is a variable or constant.)

(F1) An expression $\ulcorner t(t_1, t_2, \ldots)\urcorner$ is a formula iff (i) t is a variable, constant, or λ-expression and (ii) $t_1, t_2, \ldots$ are one or more terms.[22]

21. The notions of term and λ-expression that have been invoked here will be formally introduced in the next paragraph. This may seem to give rise to a certain circularity, since the definition of 'term' will make use of the notions of formula and λ-expression, and the definitions of 'formula' and 'λ-expression' will in turn (viz., in (F4) and (L)) make use of the concept of subject-position. But this circularity is harmless given the recursive character of these definitions. Readers who are not convinced may take the right-hand side of the present definition of 'stands at subject-position' and insert it at the appropriate places in (F4) and (L) below.

22. Unless otherwise indicated, an ellipsis (*i.e.*, '$\ldots$') is here always used, roughly speaking, to abbreviate a sequence of arbitrary and possibly infinite length. In particular, (F1) might alternatively have been written as follows:

(F1′) For any cardinality κ, an expression $\ulcorner t(\underbrace{t_1, t_2, \ldots}_{\kappa\text{-many}})\urcorner$ is a formula iff (i) t is a variable, constant, or λ-expression and (ii) $\underbrace{t_1, t_2, \ldots}_{\kappa\text{-many}}$ are one or more terms.

The same pattern should be applied to (F3), (F4), and (L) below.

(F2) An expression $\ulcorner \neg t\urcorner$ is a formula iff t is a variable, constant, or formula.

(F3) An expression $\ulcorner(t_1 \wedge t_2 \wedge \ldots)\urcorner$ is a formula iff $t_1, t_2, \ldots$ are two or more variables, constants, or formulas. (Outermost parentheses may be omitted.)

(F4) An expression $\ulcorner \exists v_1, v_2, \ldots\ \varphi\urcorner$ is a formula iff $v_1, v_2, \ldots$ are one or more pairwise distinct variables and φ is a formula such that, for each v_i (with $i = 1, 2, \ldots$): v_i has at least one free occurrence in φ, and each free occurrence of v_i in φ stands at subject-position.

(F5) Every formula belongs to one of the four types of expression described in (F1)–(F4).

(F6) Every formula is of only finite 'depth', as measured by the number of its levels of parentheses.[23]

(L) A *λ-expression* is any expression $\ulcorner(\lambda v_1, v_2, \ldots\ \varphi)\urcorner$, where φ is a formula and $v_1, v_2, \ldots$ are one or more pairwise distinct variables such that, for each v_i (with $i = 1, 2, \ldots$): v_i has at least one free occurrence in φ, and each free occurrence of v_i in φ stands at subject-position. (The parentheses may be omitted if the λ-expression does not stand in predicate-position.)

For example, the expression '$\lambda x, y\,$loves(x, y)' is a λ-expression, but '$\lambda x, y\,$loves(x, x)', '$\lambda x\, \neg x$', and '$\lambda x, y\, x(y)$' are not. By (F4), analogous restrictions hold for quantified formulas (*i.e.*, formulas beginning with '$\exists$' or '$\neg\exists$').[24]

23. For example: the depth of 'x' is 0, the depth of '$P(x)$' is 1, and the depth of '$(P(x) \wedge Q(x))$' is 2.

24. Cf. (F4) above. The restriction that λ-variables should occur only in subject-position has a precedent in Menzel (1993) and is in part motivated by the need to avoid semantic paradox. (Cf. below, footnote 36.) While the restriction is not arbitrary, it still implies a significant limitation of expressive power. *E.g.*, the property of *being an instantiated property* might, if it exists, be perspicuously denoted by an expression '$\lambda P\, \exists x\, P(x)$'; but in the present framework that expression would be ill-formed. So it may be desirable to

With the help of the concepts just introduced, we can now define two further concepts that will be useful below:

- If $\ulcorner\lambda v_1, v_2, \ldots\ \varphi\urcorner$ is a λ-expression, then the formula φ will be referred to as the *body* or *matrix* of that expression, while $v_1, v_2, \ldots$ will be referred to as its *λ-variables*.
- A term-occurrence o will be said to stand *at sentence-position* iff o does not stand in predicate-position and at least one of the following three conditions is satisfied:
  (i) o is preceded by an occurrence of '$\neg$'.
  (ii) o precedes or is preceded by an occurrence of '$\wedge$'.[25]
  (iii) o is preceded, for some variables $v_1, v_2, \ldots$, by an occurrence of either $\ulcorner\exists v_1, v_2, \ldots\urcorner$ or $\ulcorner\lambda v_1, v_2, \ldots\urcorner$.

As can be seen from the above definitions of 'formula' and 'λ-expression', the present language deviates in a number of respects from the ordinary language of first-order logic. In particular, the following points are worth noting:

- Formulas and λ-expressions can occur in subject-position.
- There is no distinction between individual constants and predicates.
- The language is infinitary in the following respects: it allows for infinitely long argument-lists and infinitely long conjunctions, and occurrences of '$\exists$' and 'λ' can bind infinitely many variables.
- Variables can appear in sentence- and predicate-position (and so expressions like '$\neg x$', '$x(y)$', and even '$x(y(z))$' all count as formulas), even though the language does *not* allow for quantification into predicate- or sentence-position.

To make the last point more palatable, it is helpful to observe, first of all, that it makes perfectly good sense to allow constants to appear in sentence- as well as in predicate-position. For example, if a constant c denotes a certain state of affairs s, then it is natural to treat $\ulcorner\neg c\urcorner$ as denoting the negation of s. Second, there seems to be no compelling reason to treat the syntax of variables any differently from that of constants, except that variables but not constants can be bound by variable-binding operators. Hence, it is natural to allow variables to appear—just like constants—in both predicate- and sentence-position, regardless of whether it is possible to quantify into either of these two kinds of position.[26]

2.3 *Semantics and Ontology*

2.3.1 *Identity, interpretations, and variable-assignments*

We begin with an ontological assumption concerning the relation of identity, to which we add a corresponding 'meaning-postulate':

(O1) There exists exactly one dyadic relation of identity, aka the *identity relation*.

(M1) For any entities x and y, the instantiation of the identity relation by x and y obtains iff x is numerically the same entity as y.

The notions of relation, instantiation, and obtainment that are used in these two statements will be introduced shortly (in section 2.3.2).

extend the framework in a way that somehow circumvents the mentioned restriction while at the same time avoiding paradox.

25. Note that '$\vee$', '$\rightarrow$', and '$\equiv$' have been introduced as merely abbreviatory devices, and need therefore not be dealt with separately at this point.

26. To be sure, if a constant does *not* denote a state of affairs, it will admittedly make little sense to have it occur in sentence-position. This is for instance reflected in the fact that, under the semantics specified in the next subsection, the formula $\ulcorner\neg c\urcorner$ is *denotationless* relative to any interpretation under which c does not denote a state of affairs. Formulas are thus allowed to lack a denotation, and the same goes for λ-expressions. (This is another respect in which the present language differs from that of classical logic.) Furthermore, since we allow for denotationless formulas, there is no compelling reason to prohibit denotationless constants, and nor do we prohibit them. An adequate logic for the present language would therefore have to be a version of *free* logic.

The reason why (O1) and (M1) are presented at this early stage, rather than only after these concepts have been properly introduced, lies in the fact that the existence of the identity relation is presupposed by the following definition of 'interpretation':

(I) An *interpretation* is any partial function (understood as a set of ordered pairs) from constants to entities that maps the constant '$\mathtt{I}$' to the identity relation.

In other words, an interpretation is any set S of ordered pairs such that (i) for any ordered pair $(x, y) \in S$, x is a constant, and (ii) S contains the pair $(\text{'}\mathtt{I}\text{'}, \mathtt{I})$, where $\mathtt{I}$ is the identity relation.[27] Similarly (though more straightforwardly), we will say that a *variable-assignment* is any partial function from variables to entities.

In the following, all talk of denotation should be understood as relativized to an interpretation and a variable-assignment. That is, it should be kept in mind that terms in general have or lack a denotation only relative to an interpretation and variable-assignment. This relativization will often be left implicit. By way of abbreviation, if 'I' refers to an interpretation and 'g' to a variable-assignment, we will write 'A denotes$_{I,g}$ B' to mean the same as 'A denotes B relative to I and g'. With this notation in hand, we can formulate our first semantic stipulation:

(S1) A constant or variable α has a denotation relative to an interpretation I and a variable-assignment g iff α is in the domain of either I or g. In this case, α denotes$_{I,g}$ the entity to which it is mapped by either I (if it is a constant) or g (if it is a variable).

Finally, if two expressions denote the same entity relative to an interpretation I and a variable-assignment g, we will say that these expressions are *coreferential* relative to I and g, or more briefly: coreferential$_{I,g}$.

2.3.2 *Instantiations*

The first class of states of affairs that we have to be concerned with is that of *instantiations of attributes*. A typical example is the instantiation by Hypatia (or alternatively: Hypatia's instantiation) of the property of *being human*, *i.e.*, the state of affairs that Hypatia is human. Like other states of affairs, instantiations may exist even if they do not obtain. Thus, if we say that there exists the state of affairs that Hypatia is human, this should not be read as implying that Hypatia is human. To imply the latter, one would rather have to say that the state of affairs in question *obtains*. More generally, if A is an attribute and $x_1, x_2, \ldots$ are entities, we will say that $x_1, x_2, \ldots$ (in this order) *instantiate* A iff the instantiation of A by $x_1, x_2, \ldots$ obtains. The concepts of instantiation and obtainment that are used here will be treated as primitive, as will the concept of a state of affairs; nonetheless we shall below adopt a 'meaning-postulate' that aims to clarify the concept of instantiation by stipulating that each instantiation is a state of affairs.[28] Finally, the locution 'instantiation of ... by ...' should be read as *referentially transparent*. Thus, if the expressions ε_1, ε_2, ε_3, ... respectively refer to the same entities as the expressions ζ_1, ζ_2, ζ_3, ..., then the predicate $\ulcorner$is an instantiation of ε_1 by $\varepsilon_2, \varepsilon_3, \ldots\urcorner$ will be coextensive with $\ulcorner$is an instantiation of ζ_1 by $\zeta_2, \zeta_3, \ldots\urcorner$.

27. Traditionally, interpretations are allowed to vary not only with respect to the denotations they assign to non-logical constants, but also with respect to the *domains of quantification* they assign to the primitive quantifier(s) of the respective language. This complication is here omitted. Instead, any occurrence of the quantifier '$\exists$' is in the present framework intended to range over everything there is. (This quantification is to be understood as unrestricted; we can leave it open, however, whether it should also be understood as 'absolutely general'. On the question of absolutely general quantification, cf. the essays collected in Rayo & Uzquiano [2006], as well as, *e.g.*, Williamson [2003] and Studd [2015].)

28. It might be wondered why this postulate makes no use of modal notions. (In particular, it might be wondered why it does not state that *necessarily*, every instantiation is a state of affairs. Analogous questions can be raised about the other meaning-postulates in this section, starting with (M1) above.) Briefly put, the reason is that modal notions are in my view to be explicated on the basis of logical concepts (together with a concept of a state of affairs and a concept of essence), rather than *vice versa*.

On the basis of the concept of instantiation, we can define the notion of an attribute as well as that of a *κ-adic* attribute, where κ is any cardinal number greater than zero. In particular, we will say that something is an *attribute* iff it has an instantiation, and that an attribute A is *κ-adic* (or alternatively: that κ is an *adicity* of A) iff there exists an instantiation of A by some (not necessarily distinct) entities $\underbrace{x_1, x_2, \ldots}_{\kappa\text{-many}}$.

To avoid this two-dimensional notation, we may alternatively say that an attribute A is κ-adic iff A has an instantiation by a *sequence of entities* of length κ: these two definitions should be read as notational variants of each other. The notions of property and relation can now be (re-)defined as follows: a *property* is any monadic attribute, and a *relation* is any attribute that has an adicity greater than one.

In the next step, we can formulate our main ontological assumption regarding instantiations of attributes, accompanied by the meaning-postulate already hinted at:

(O2) For any κ-adic attribute A and any sequence of length κ of entities $x_1, x_2, \ldots$, there exists exactly one instantiation of A by $x_1, x_2, \ldots$ (in this order).

(M2) Every instantiation is a state of affairs.

Further, we adopt the following stipulation with respect to the semantics of formulas (more will be added below):

(S2) A formula $\ulcorner t(\underbrace{t_1, t_2, \ldots}_{\kappa\text{-many}})\urcorner$ has a denotation relative to an interpretation I and a variable-assignment g iff t denotes$_{I,g}$ a κ-adic attribute A and the terms $t_1, t_2, \ldots$ respectively denote$_{I,g}$ (not necessarily distinct) entities $x_1, x_2, \ldots$. In this case, the formula denotes$_{I,g}$ the instantiation of A by $x_1, x_2, \ldots$ (in this order).

For example, suppose that 'human' is a constant denoting$_{I,g}$ the property of *being human*, and that 'Hypatia' is a constant denoting$_{I,g}$ Hypatia. Then the formula 'human(Hypatia)' will denote$_{I,g}$ Hypatia's instantiation of the property of *being human*.

In connection with instantiations, and more generally concerning the relationship between attributes and states of affairs, we will make three additional assumptions. Like (M2), these may also best be regarded as meaning-postulates, since they help to specify how the relevant technical terms are to be understood:[29]

(P1) For any attribute A, there exists at least one cardinal number $\kappa \geq 1$ such that A is κ-adic.

(P2) No attribute has more than one adicity.[30]

(P3) If A and B are two κ-adic attributes, then there exists at least one sequence of length κ of entities $x_1, x_2, \ldots$ such that the instantiation of A by $x_1, x_2, \ldots$ (in this order) is distinct from the instantiation of B by $x_1, x_2, \ldots$ (in this order).

In a nutshell, this last postulate states that no attribute shares all its instantiations with any other. Obviously this allows that two attributes

29. The present usage of the word 'assumption' is fairly inclusive, covering any sort of postulate as well as assumptions in the ordinary sense of the word.

30. This postulate has something of a predecessor in Armstrong's (1978, 94) *Principle of Instantial Invariance*: "For all n, if a universal is n-adic with respect to a particular instantiation, then it is n-adic with respect to all its instantiations". (Note, however, that Armstrong's concept of an instantiation refers only to obtaining states of affairs, whereas the present concept is more general in that it also applies to *non*-obtaining states of affairs.)

It is true that (P2) denies the existence of 'multigrade attributes', such as might be thought to correspond to predicates like 'is surrounded by' or 'are arranged in a circle', and it is also true that the existence of multigrade attributes cannot easily be ruled out *a priori* (cf. MacBride [2005, §2]; the same goes for what MacBride calls 'varigrade relations'). However, I think that there is little positive reason to complicate our ontology by the admission of such attributes. As far as I can see, we do not lose anything of substance if we say that multigrade predicates fail to correspond to attributes, provided that we can draw on the resources of set theory. For example, instead of postulating some multigrade relation allegedly expressed by the verb 'is surrounded by', we can make do with the dyadic relation that is instantiated by an entity x and a set Y just in case (say) each possible route of escape for x is blocked by some member of Y.

may share *some* of their instantiations. For instance, if R is a non-symmetric relation and x some entity, then the property of *bearing R to x* is plausibly distinct from the property of *being borne R to by x*, and yet the instantiation by x of the first property may well be identical with x's instantiation of the second.[31]

2.3.3 Conjunction and negation

Apart from instantiations of attributes, we also assume that there are negations and conjunctions of states of affairs. As in the case of (O2), the relevant ontological assumptions are accompanied by meaning-postulates and stipulations regarding the semantics of formulas:

(O3) For any state of affairs s, there exists exactly one negation of s.

(M3) For any state of affairs s, a *negation* of s is a state of affairs that obtains iff s does not obtain.

(S3) A formula $\ulcorner \neg t \urcorner$ has a denotation relative to an interpretation I and a variable-assignment g iff t denotes$_{I,g}$ a state of affairs s. In this case, the formula denotes$_{I,g}$ the negation of s.

(O4) For any states of affairs $s_1, s_2, \ldots$, there exists exactly one conjunction of them.

(M4) For any states of affairs $s_1, s_2, \ldots$, a *conjunction* of them is a state of affairs that obtains iff each one of the s_i (with $i = 1, 2, \ldots$) obtains.

(S4) A formula $\ulcorner t_1 \wedge t_2 \wedge \ldots \urcorner$ has a denotation relative to an interpretation I and a variable-assignment g iff each one of the t_i (with $i = 1, 2, \ldots$) denotes$_{I,g}$ a state of affairs s_i. In this case, the formula denotes$_{I,g}$ the conjunction of the s_i.

In interpreting these statements, it should be kept in mind that the locutions 'negation of ...' and 'conjunction of ...' are referentially transparent. For example, if the expressions ε and ζ refer to the same state of affairs, then the predicate $\ulcorner$is a negation of $\varepsilon\urcorner$ will be coextensive with $\ulcorner$is a negation of $\zeta\urcorner$; and analogously for 'conjunction of'.

2.3.4 Quantification

We further assume the existence of 'quantificational' states of affairs:

(O5) For any attribute A, there exists an existential quantification of A.

(M5) For any attribute A, an *existential quantification* of A obtains iff there exist some entities $x_1, x_2, \ldots$ that instantiate A.

(S5) A formula $\ulcorner \exists v_1, v_2, \ldots \varphi \urcorner$ has a denotation relative to an interpretation I and a variable-assignment g iff the λ-expression $\ulcorner \lambda v_1, v_2, \ldots \varphi \urcorner$ denotes$_{I,g}$ an attribute. In this case, the formula denotes$_{I,g}$ the existential quantification of that attribute.

The locution 'existential quantification of ...' should of course be taken to be referentially transparent.

It is tempting to read '$\exists$' as 'there exists' and, correspondingly, to interpret quantified formulas as formal counterparts of existence claims in English. This would not be entirely correct. For instance, suppose it is stipulated that I be an interpretation that assigns to the constant 'Pegasus' the same entity (if any) that the word 'Pegasus' refers to in English. (In addition, let g be an arbitrary variable-assignment.) Then, assuming that 'Pegasus' does not in fact refer to anything in English, no entity will under I be assigned to that constant, and consequently, by (S1), 'Pegasus' will not have a denotation relative to I and g. By (S2), (S3), and (S5), it will further follow that the formula '$\neg\exists x\,(x = \text{Pegasus})$' does not denote$_{I,g}$ anything and accordingly fails to be true relative to I and g, despite its being apparently true in English that "Pegasus doesn't exist".[32] This should not be considered a defect, however, since it has not been our aim that the semantics of the existential quantifier should in all respects conform to the ordinary usage of 'exists'.[33]

31. In fact, in the present framework, x's instantiation of the first property *is* identical with its instantiation of the second. (Cf. p. 15 below.)

32. For the relevant definition of 'truth', see p. 13 below.

33. For related discussion, see, *e.g.*, Plantinga (1974, ch. 8) and Ray (2014).

2.3.5 λ-Abstraction

We now turn to the semantics of λ-expressions. To formulate the relevant stipulation, it will be useful to have some notation for the modification of variable-assignments. Thus, if g is some variable-assignment, $v_1, v_2, \ldots$ are κ-many pairwise distinct variables, and $x_1, x_2, \ldots$ κ-many pairwise distinct entities, we will write '$g[x_1/v_1, x_2/v_2, \ldots]$' to refer to the variable-assignment that is just like g except that v_1 is mapped to x_1, v_2 is mapped to x_2, and so on. With this notation in hand, we can stipulate that

(S6) A λ-expression $\ulcorner\lambda \underbrace{v_1, v_2, \ldots}_{\kappa\text{-many}} \varphi\urcorner$ has a denotation relative to an interpretation I and a variable-assignment g iff there exists a sequence of length κ of entities $x_1, x_2, \ldots$ such that φ has a denotation relative to I and $g[x_1/v_1, x_2/v_2, \ldots]$. In this case, the λ-expression denotes$_{I,g}$ a κ-adic attribute A such that, for any sequence of length κ of entities $y_1, y_2, \ldots$, the instantiation of A by $y_1, y_2, \ldots$ (in this order) is the state of affairs that is relative to I and $g[y_1/v_1, y_2/v_2, \ldots]$ denoted by φ.[34]

Note that the second sentence ("In this case...") implies an ontological assumption about the abundance of attributes.[35] To illustrate the import of this assumption, let φ be the formula 'black$(x) \wedge$ horse(x)', let I be an interpretation that maps the constant 'black' to the property of *being black* and the constant 'horse' to the property of *being a horse*, and let g be a variable-assignment that maps the variable 'x' to some entity x. Then φ denotes$_{I,g}$ the state of affairs that x is a black horse. By the second sentence of (S6), it now follows that there exists a monadic attribute (in other words, a property) A whose instantiation by any given entity y is the state of affairs that y is a black horse. Moreover, by (P3) (p. 10 above), there exists only one such attribute. It will thus be appropriate to refer to A as 'the property of *being a black horse*'. This point generalizes to all other λ-expressions, and in this way it is guaranteed that each denoting λ-expression has a definite referent.[36]

34. The existence of such a state of affairs is guaranteed by our ontological assumptions and the recursive semantics of formulas (including (S6) itself), together with the fact that (i) by hypothesis, there exists at least one sequence of entities $x_1, x_2, \ldots$ such that φ has a denotation relative to I and $g[x_1/v_1, x_2/v_2, \ldots]$ and (ii) by the syntax of λ-expressions, the variables $v_1, v_2, \ldots$ occur in φ only at subject-position. (To see the relevance of this latter restriction, note that, by (S2), a formula $\ulcorner t(u)\urcorner$ has a denotation relative to a given interpretation I and variable-assignment g *only* if t denotes$_{I,g}$ a property. So, if φ were a formula $\ulcorner t(u)\urcorner$, where u is some term and t is one of the v_i, then it would not be the case that for every κ-sequence $y_1, y_2, \ldots$, there exists a state of affairs that is relative to I and $g[y_1/v_1, y_2/v_2, \ldots]$ denoted by φ.)

35. To have a label for it, call it '(O6)'. I trust that there is no need to write up (O6) separately from (S6), since it can be straightforwardly derived from the latter.

36. In connection with the need to avoid paradox, it may be worth elaborating briefly on what (S6) does *not* entail. In particular, due to the restriction that the variables $v_1, v_2, \ldots$ should occur in φ only at subject-position, one of the things that do not follow from (S6) is the existence of a property P such that, for any entity x, the instantiation of P by x is identical with x itself. In other words, it does not follow from (S6) that there exists what one might call a property of *obtainment*, which (if our formation rule for λ-expressions had been more liberal) might have been denoted by '$\lambda x\, x$'. Likewise, it does not follow from (S6):

- that there exists a property of non-self-instantiation, which might have been denoted by '$\lambda x\, \neg x(x)$' and whose existence would have given rise to a form of Russell's paradox;
- that there exists, for any state of affairs s, a property of *being a property that instantiates itself only if s obtains*, which might have been denoted by $\ulcorner\lambda x\, (x(x) \rightarrow \varphi)\urcorner$ (where φ is a formula that denotes s and in which 'x' does not occur free) and whose existence would have given rise to a form of Curry's paradox;
- that there exists a property, which might have been denoted by '$\lambda x\, \exists P\, (\neg P(x) \wedge (x = \forall y\, P(y)))$', of *being an entity x such that, for some property P, it is both the case that x does not have P and that x is the state of affairs that everything has P*, whose existence would have threatened to give rise to a version of the Russell–Myhill paradox. (I say 'threatened' because the paradox requires the premise, which our framework in its present form does not provide, that no two properties P and Q are such that the state of affairs $\forall x\, P(x)$ is identical with $\forall x\, Q(x)$.)

Finally, it does not follow from (S6) that, for every arbitrary subclass of some more-than-set-sized collection (such as the ordinals), there exists the property of *being a member of that subclass*, which would have given rise to

With the help of λ-expressions, one can easily define concepts of negation, conjunction, and disjunction that apply to attributes. Thus, if A is a κ-adic attribute, then the likewise κ-adic attribute $\lambda x_1, x_2, \ldots \neg A(x_1, x_2, \ldots)$ will be referred to as the *negation* of A, and may also be written as '$\sim A$'.[37] Further, if $A_1, A_2, \ldots$ are some attributes, all of the same adicity, then $\lambda x_1, x_2, \ldots \big(A_1(x_1, x_2, \ldots) \wedge A_2(x_1, x_2, \ldots) \wedge \ldots\big)$ will be referred to as the *conjunction* of those attributes, and $\lambda x_1, x_2, \ldots \big(A_1(x_1, x_2, \ldots) \vee A_2(x_1, x_2, \ldots) \vee \ldots\big)$ as their *disjunction*. The conjunction may also be written as '$A_1 \,\&\, A_2 \,\&\, \ldots$', and the disjunction as '$A_1 \text{ v } A_2 \text{ v } \ldots$'. Finally, if A and B are two κ-adic attributes, then '$(A \equiv B)$' may be used as a name for the attribute $\lambda x_1, x_2, \ldots \big(A(x_1, x_2, \ldots) \leftrightarrow B(x_1, x_2, \ldots)\big)$.

2.4 *Coarse-Grainedness*

We now have to add one more postulate, which will state a sufficient condition for the identity of states of affairs. Although it will superficially apply only to states of affairs, it will also have consequences for the individuation of attributes.

To formulate this postulate, we will require some additional terminology. First, a term will be said to be *semantically well-formed* iff there exist an interpretation and a variable-assignment relative to which it has a denotation. Second, a formula will be said to be *true* (*false*) relative to a particular interpretation I and variable-assignment g iff it denotes$_{I,g}$ an obtaining (non-obtaining) state of affairs. Relatedly, a formula will be said to be *logically true* iff it is true relative to every interpretation and variable-assignment. A concept of entailment suitable for our purposes can now be defined as follows:

(E) A formula φ *entails* a formula ψ iff φ is semantically well-formed and, for any interpretation I and variable-assignment g, the following two conditions are satisfied:

 (i) If φ has a denotation relative to I and g, then so does ψ.

 (ii) If φ is true relative to I and g, then so is ψ.

As usual, two formulas φ and ψ will be said to be *equivalent* iff they entail each other.[38]

It may be instructive to consider in some detail how it follows from our assumptions that any formula φ is equivalent to its double negation $\ulcorner \neg\neg\varphi \urcorner$. To this end, let I and g be any interpretation and variable-assignment. If φ has relative to I and g any denotation at all, it will denote$_{I,g}$ a state of affairs, as can be seen from the formation rules in section 2.2 together with the semantic stipulations (S1)–(S6) and the meaning-postulates (M1)–(M5). (For example, (S2) and (M2) together entail that a formula $\ulcorner t(t_1, t_2, \ldots) \urcorner$ denotes, if anything, a state of affairs.) On the other hand, if φ does not denote$_{I,g}$ anything, then it follows from (S3) that $\ulcorner \neg\neg\varphi \urcorner$ will not denote$_{I,g}$ anything, either; and *vice versa*. So we now only have to consider the case in which φ denotes$_{I,g}$ a state of affairs s. From (S3), it then follows that $\ulcorner \neg\neg\varphi \urcorner$ denotes$_{I,g}$ the negation of the negation of s, which, by (M3), obtains iff the negation of s fails to obtain. But, again by (M3), the negation of s fails to obtain iff s itself obtains. So the negation of the negation of s obtains iff s itself

cardinality-related worries not unlike the paradoxes discussed by Kaplan (1995) and Uzquiano (2015).

37. The use of λ-expressions for the purpose of denoting relations has been criticized by Peter van Inwagen, who claims that "we do not understand abstraction names of dyadic relations unless those relations are symmetrical" (2006, 468). I agree with this assessment insofar as I think that our understanding of relations in general—symmetrical and otherwise—could stand improvement. (For relevant discussion, see, *e.g.*, Fine [2000] and Dorr [2004].) But fortunately, no deep understanding of the general metaphysics of relations will be needed for the purposes of this paper.

38. The present concept of entailment is to some extent related to William Parry's (1933; 1972) concept of analytic implication as well as to Casimir Lewy's (1976) concept of analytic entailment. (Thanks here to Kit Fine.) To see the analogy to Parry's concept, one has to bear in mind the constraints that (S1)–(S6) impose on a given term's having a denotation relative to a given interpretation and variable-assignment. Under these constraints, and given clause (i) of (E), a formula can entail another *only* if the second formula does not contain any free occurrence of any variable or non-logical constant that does not also have a free occurrence in the first.

does. Consequently, if φ has a denotation relative to I and g, then φ is true relative to I and g if and only if $\ulcorner\neg\neg\varphi\urcorner$ is. This argument generalizes to any interpretation and variable-assignment, and so it follows that φ is equivalent to $\ulcorner\neg\neg\varphi\urcorner$.

Using the concept of equivalence introduced above, we can now formulate our postulate as follows:

(P4) For any states of affairs s and t, for any interpretation I and variable-assignment g, and for any formulas φ and ψ: if s and t are respectively denoted$_{I,g}$ by φ and ψ, and φ is equivalent to ψ, then s is identical with t.

In other words, no two states of affairs are denoted (relative to the same interpretation and variable-assignment) by equivalent formulas. This stipulation is motivated by the thought that states of affairs are in a certain sense 'worldly': they are *states* of affairs, not mere representations of affairs—just as properties are ways for things to *be*, and not ways for things to be described. This 'worldly' aspect of the present conception of states of affairs (and attributes) calls for an identity criterion under which states of affairs are individuated in as coarse-grained a manner as possible. Correspondingly, it calls for as strong a version of (P4) as possible. Certain limiting considerations, however, favor the present version of (P4) over its even stronger alternatives. Consider for instance Socrates' self-identity, *i.e.*, the state of affairs that Socrates is self-identical, or in symbols, (Socrates $=$ Socrates): intuitively, it seems absurd to say that this state of affairs is identical with *Plato*'s self-identity. But the formulas 'Socrates $=$ Socrates' and 'Plato $=$ Plato' are classically equivalent and denote, relative to a suitable interpretation and variable-assignment, the two states of affairs just mentioned. Hence, a strengthened version of (P4), according to which no two states of affairs are denoted by *classically* equivalent formulas, would have had the intuitively absurd consequence that Socrates' self-identity is the same as Plato's.

For another example, consider on the one hand the state of affairs that everything black is black and, on the other, the state of affairs that every horse is a horse. Relative to a suitable interpretation and variable-assignment, these states of affairs are respectively denoted by '$\forall x\,(\text{black}(x) \rightarrow \text{black}(x))$' and '$\forall x\,(\text{horse}(x) \rightarrow \text{horse}(x))$'. But these two formulas are again classically equivalent, and so, under the strengthened version of (P4) mentioned in the previous paragraph, the state of affairs that everything black is black counts as identical with the state of affairs that every horse is a horse. This consequence, too, is intuitively absurd,[39] and it is avoided by (P4) only due to its reliance on the *non*-classical concept of entailment defined in (E).[40]

A noteworthy corollary of (P4) is the thesis that all logically true formulas, such as '$\mathtt{I} = \mathtt{I}$' and '$\forall x\,(x = x)$', denote one and the same state of affairs, given that all logical truths are equivalent. This consequence may at first seem unpalatable: shouldn't the self-identity of the identity relation be distinct from the state of affairs that *everything* is self-identical? This is a natural doubt to have if one focuses on the syntactic difference between the two expressions. However, given the

39. That is, as long as it is granted that horsehood and blackness are not analyzable in terms of each other, which, at least *prima facie*, seems to be a safe assumption. (On the notion of analyzability, see section 3.5 below.)

40. In view of the relative complexity of (E), it may be wondered whether we might not instead have adopted a simpler concept of entailment, which could have been defined as follows:

(E*) A formula φ *entails* a formula ψ iff φ is semantically well-formed and, for every interpretation I and variable-assignment g, the following condition is satisfied: if φ is true relative to I and g, then so is ψ.

Although this definition is simpler than (E), it should arguably not be preferred, since it conflicts with the intuition that Socrates' self-identity is distinct from Plato's. To see this, suppose we had adopted (E*) instead of (E). This would in the first place have had the result that any two semantically well-formed formulas that denote an obtaining state of affairs relative to *no* interpretation and variable-assignment, such as '$\neg\mathtt{I}(\text{Socrates}, \text{Socrates})$' and '$\neg\mathtt{I}(\text{Plato}, \text{Plato})$', would have counted as equivalent. Hence, the non-self-identity of Socrates would under (P4) have counted as identical with the non-self-identity of Plato. But by (O3), we would then have had to say that the negation of Socrates' non-self-identity is identical with the negation of the non-self-identity of Plato. By (S3) and (P4), which allow us to remove double negation signs, it would finally have followed that Socrates' self-identity is identical with Plato's, contradicting the mentioned intuition.

'worldly' aspect of the present conception of states of affairs, it would be less misleading to focus instead on what the truth of the respective formula 'requires of the world'.[41] And in the relevant sense, the truth neither of '$\mathtt{I} = \mathtt{I}$' nor of '$\forall x\,(x = x)$' requires *anything* of the world; and so they should not be taken to denote distinct states of affairs.[42] By contrast, the truth of 'Socrates = Socrates' requires the existence of Socrates but not that of Plato, and the truth of 'Plato = Plato' requires the existence of Plato but not that of Socrates.

By stating a sufficient condition for the identity of states of affairs, (P4) provides a sufficient condition for the identity of attributes as well. Thus, let A and B be any κ-adic attributes, and suppose that, relative to some interpretation I and variable-assignment g, A and B are respectively denoted by the λ-expressions L_1 and L_2, where, for some κ-many variables $v_1, v_2, \ldots$ and two *equivalent* formulas φ and ψ, L_1 is identical with $\ulcorner \lambda v_1, v_2, \ldots\ \varphi \urcorner$ and L_2 is identical with $\ulcorner \lambda v_1, v_2, \ldots\ \psi \urcorner$. In addition, let $\underbrace{y_1, y_2, \ldots}_{\kappa\text{-many}}$ be some arbitrary (and not necessarily distinct) entities. By (S6), we then have that A's instantiation by $y_1, y_2, \ldots$ (in this order) is the state of affairs denoted by φ relative to I and $g[y_1/v_1, y_2/v_2, \ldots]$. Analogously, B's instantiation by $y_1, y_2, \ldots$ will be the state of affairs denoted by ψ relative to I and $g[y_1/v_1, y_2/v_2, \ldots]$. By (P4), given that φ and ψ are equivalent, it now follows that these states of affairs are one and the same. But $y_1, y_2, \ldots$ was an arbitrary κ-sequence of entities. So, for any κ-sequence of entities, A's instantiation by that sequence is the same as B's. Hence, by (P3), A is identical with B.

41. Talk of 'what is required of the world' plays a prominent role in, *e.g.*, Rayo (2008; 2009; 2013) and Williams (2010; 2012).

42. If we were what Rayo (2009) calls 'traditional Platonists', we should say that the truth neither of '$\mathtt{I} = \mathtt{I}$' nor of '$\forall x\,(x = x)$' requires anything of the world *other than the existence of the identity relation*. (For recall that '$\forall x\,(x = x)$' is shorthand for '$\neg\exists x\, \neg \mathtt{I}(x, x)$', and would thus not have a denotation if '$\mathtt{I}$' did not denote a dyadic relation.) But it would then still turn out that these two formulas do not differ from each other with respect to what their truth requires.

2.5 *Hyperintensionality*

An assumption vaguely similar to (P4) can already be found in Wittgenstein's *Tractatus*: "If p follows from q and q from p, they are one and the same proposition" (5.141).[43] As long as the 'follows from' in this assertion is understood in the relatively weak sense of classical entailment, the corresponding conception of propositions will be (merely) intensional rather than hyperintensional. By contrast, (P4) is compatible with a hyperintensional conception of states of affairs (and thus also of attributes), given that it employs a stronger notion of entailment and thereby leaves open the possibility that classically equivalent formulas may denote distinct states of affairs.[44]

To be sure, the fact that our postulate *allows* for such finer-grained distinctions does not yet mean that the present framework also *commits* us to them. For example, it does not commit us to the claim that Socrates' self-identity is distinct from Plato's. For this reason, it may in the end not be appropriate to call the present framework hyperintensional 'in letter', but it certainly is in spirit. For, by employing a concept of equivalence that is stronger than that of classical logic, (P4) is deliberately formulated in such a way as to make room for distinctions of the (relatively) fine-grained sort just alluded to. The framework could be easily rendered hyperintensional also 'in letter' by adding to it a further postulate that entails at least some such distinctions. The task of formulating such a postulate will here be left for another occasion.[45]

43. Other authors who have adopted assumptions similar to (P4) include Carnap (1942, 92), Hacking (1967, 165), Stalnaker (1976, 72f.), Pollock (1984, 54), and Olson (1987, 91). Alvin Plantinga appears tacitly to assume a related principle—which, in his terminology, could be stated as, 'No two states of affairs include each other'—when he argues (1974, §4.1) that there is at most one actual world.

44. Usually, the contrast between 'merely intensional' and 'hyperintensional' is not spelled out with the help of the contrast between classical and non-classical concepts of entailment, but rather (at least in part) with the help of modal notions. (For example, see Cresswell [1975, 25] and Nolan [2014, 151].) However, this difference does not matter for present purposes.

45. For the relevance of the present non-classical concept of entailment for the project of this paper, see footnote 60 below (p. 23).

3. Simple Attributes

We are now in a position to approach the main task of this paper, viz., the analysis of the concept of a logically simple attribute.[46] The analysis itself will be presented in section 3.3; readers are welcome to skip ahead, as the first two subsections merely aim to forestall several objections. In particular, section 3.1 tries to forestall objections regarding the choice of framework and the general character of the to-be-proposed account of attribute simplicity, while section 3.2 develops an analysis of *mereological* simplicity in the same abstract or 'indirect' style that will also characterize the former account. The point of this exercise will be to show that our account of attribute simplicity is analogous to an adequate analysis of mereological simplicity. It is hoped that this analogy will persuade the reader that it won't be inappropriate to apply the term 'logically simple' to all and only those attributes that are classified as such by our account.

3.1 Preliminary Objections and Replies

Some readers may object to the very project of developing an *analysis* of the concept of a logically simple attribute, on the ground that the meaning of the term 'logically simple' (as applied to attributes) has as yet not been properly fixed. Accordingly, they may regard it simply as a technical term awaiting stipulative definition. There is certainly something to this charge, for the term in question *is* a technical term, and we are apparently *somewhat* free to give it a stipulative definition. But only somewhat: after all, we all have some rough idea of what it is for something to be 'simple', and the adverb 'logically' certainly also has some cognitive significance, however vague. So one would expect that their combination will likewise have some more or less tangible meaning. For this reason, I will here continue to represent the task of this section as one of analysis rather than stipulative definition. But at the same time, it should be clear that this project is not so much an exercise in descriptive semantics as an attempt at systematic philosophy: the ulterior aim is to carve out a philosophically valuable notion, and not to construct a faithful description of linguistic usage.

If it could be assumed that the 'logical' construction of attributes out of others—*e.g.*, by forming their negations or conjunctions—never leads in a circle, then the task of analyzing the concept of a logically simple attribute would be rather easy. For in that case, one might simply say that an attribute is logically simple iff it is not the negation, nor the conjunction or disjunction (etc.), of any other attributes.[47] This straightforward approach requires, however, an extremely fine-grained individuation of attributes. For instance, any property P would have to be distinguished from its 'double negation' $\sim\sim P$; and for any properties P and Q, the disjunction of $P \,\&\, Q$ and $P \,\&\, {\sim}Q$ would have to be distinguished from P itself. In the present framework, attributes are not this finely individuated, mainly as a result of the postulate (P4) (p. 14 above), according to which no two states of affairs are denoted by equivalent formulas.[48] A friend of the straightforward approach just

46. As above, the qualifier 'logically' will be treated as optional. In addition, rather than to speak of 'logical simplicity, as applied to attributes', I will typically use the term 'attribute simplicity'.

47. Given a suitable account of what it is for something to be a 'constituent' of an attribute, one could say, even more straightforwardly, that an attribute is simple iff it has no (proper) constituents. Analyses along these lines can be found, *e.g.*, in Armstrong (1978, 67) and Heil (2003, §14.6).

48. To see that every property P is in the present framework identical to its double negation, note that the formula '$P(x)$' is (in the sense specified on p. 13 above) equivalent to '$\neg\neg P(x)$'. Thus, by (P4) and the semantics of formulas, these two formulas will relative to any given interpretation and variable-assignment denote either the same state of affairs or nothing at all. The identity of P and $\sim\sim P$ then follows with the help of (P3) and the semantics of λ-expressions. (Cf. the proof at the end of section 2.4.) To see that, moreover, some properties P and Q are such that P is identical with the disjunction of $P \,\&\, Q$ and $P \,\&\, {\sim}Q$, let P be some arbitrary property and let Q be the property of *being the only entity* (in symbols: $\lambda x\, \neg\exists y\, (y \neq x)$). Since '$P(x)$' is equivalent to the disjunction

$$\big(P(x) \wedge \neg\exists y\, (y \neq x)\big) \vee \big(P(x) \wedge \neg\neg\exists y\, (y \neq x)\big),$$

it follows from the semantics of λ-expressions together with the assumptions of our framework that P is then identical with the disjunction of $P \,\&\, Q$ and $P \,\&\, {\sim}Q$.

alluded to might thus wonder whether it would not be better to reject (P4) and to adopt instead a more fine-grained conception of attributes and states of affairs. A similar point might be made by an adherent of Armstrong's (1978) theory of universals, according to which the only logically complex universals are conjunctions of other universals. For if the only logical relation that holds between distinct universals is the one that holds between a conjunctive universal and its conjuncts, then why not say that the simple attributes are precisely those universals that are not conjunctive? The task would be easier still if we did not believe even in conjunctive universals, for then we could say that the simple attributes are just the universals.

I would respond to these considerations as follows. To the Armstrongian, it can be replied that, even if her theory of universals is correct, it nonetheless has not yet been conclusively established, and nor is it easy to see on what basis it *could* be established.[49] (The same may be said about the 'hyper-Armstrongian' who rejects even the existence of conjunctive universals.) As long as this situation persists, it would surely be a good thing to have an account of attribute simplicity that does not rely on the assumption that the only way for a universal to be logically complex is to be a conjunction of other universals, or on the even stronger assumption that there are no logically complex universals at all.

Essentially the same response can be given to what one might call the non-Armstrongian proponent of a (very) fine-grained conception of attributes. *Qua* non-Armstrongian (as the term will be used here), such a theorist believes that there are not only conjunctive attributes but also, *e.g.*, disjunctive and/or negative ones. But by virtue of her fine-grained conception of attributes (where 'fine-grained' is to be understood in the sense of implying a denial of (P4)), she would for example deny that any properties P and Q are such that the disjunction of $P \,\&\, Q$ and $P \,\&\, {\sim}Q$ is identical with P. Arguments for such a view can be produced, but are not conclusive.[50] As long as this is the case, it

49. This is perhaps reflected in the fact that Armstrong himself has, albeit reluctantly, come to allow the existence of negative and disjunctive universals in his (1997, 28). (Also cf. his [1978, 119].)

50. The most prominent argument is due to Sober (1982), who discusses a certain hypothetical device that would be more aptly described as a 'triangularity detector' than as a 'trilaterality detector', even though triangularity and trilaterality are "mathematically equivalent in standard geometric theories" (p. 183). In a somewhat similar vein, Ralf Bader (2013, 538) maintains that, for some properties P and Q (where P is intrinsic and Q extrinsic), $P \vee Q$ should be distinguished from $P \vee ({\sim}P \,\&\, Q)$, on the ground that the former but not the latter can be "had both intrinsically and extrinsically" by one and the same entity. And finally, a closely related argument, which has been suggested to me by Alex Skiles, is based on considerations of grounding: Let s_1, s_2, and s_3 be three obtaining states of affairs (or 'facts') that may respectively be denoted by the formulas 'cube(Max)', '$\text{cube(Max)} \wedge \exists y\,(y \neq \text{Max})$', and '$\text{cube(Max)} \vee (\text{cube(Max)} \wedge \exists y\,(y \neq \text{Max}))$'. In English, s_1 may be described as the fact that Max is a cube, s_2 as the fact that Max is an accompanied cube, and s_3 as the fact that Max is either a cube or an accompanied cube. By (P4), s_1 is identical with s_3. But it seems plausible to say that, while s_1 is not grounded in s_2, s_3 is. In other words, it seems true to say that the fact that Max is either a cube or an accompanied cube, unlike the 'simple' fact that Max is a cube, is grounded in the fact that Max is an accompanied cube. It would then follow that s_1 and s_3, as described, are not in fact one and the same state of affairs, contrary to (P4).

An adequate discussion of these three arguments (and others like them) would here lead too far afield, but at least *prima facie*, they can all be answered by roughly the same kind of consideration. Thus, to begin with the first argument, we can defensibly maintain that a triangularity detector is *eo ipso* a trilaterality detector (on a fairly literal reading, at least); for there does not appear to be any compelling reason to think that the justification for treating the terms 'triangularity detector' and 'trilaterality detector' as non-interchangeable has anything to do with the relatively arcane matter of how to individuate properties, rather than merely with the differences in which the devices thereby designated function, which are respectively hinted at by those two terms. In other words, what makes us call a given device a 'triangularity detector' rather than a 'trilaterality detector' may well only have to do with, roughly speaking, how the behavior of the device might best be explained, which in turn has to do with the respective *Sinn* of 'triangularity' and 'trilaterality', and not only—that is, not exclusively—with the reference of these terms. Something analogous may be said about the ascription of predicates like 'has P intrinsically' and 'has P extrinsically' (where 'P' is to be replaced by a name of a property) and for the truth of sentences like 'X is grounded in Y', where 'X' and 'Y' are to be replaced by names of states of affairs. Thus, insofar as it is true to say of some entity that it has the property $P \vee Q$ both intrinsically and extrinsically, but *not* true to say of the same entity that it has the property $P \vee ({\sim}P \,\&\, Q)$ both would certainly be desirable to have an analysis of attribute simplicity that harmonizes with a coarser-grained conception of attributes, and in particular with a conception under which (P4) is true. More precisely: it would be desirable to have an analysis that, for at least one plausible coarser-grained conception C, does not yield, when combined with the postulates characterizing C (such as, say, (P4)), any consequences that are either wildly implausible or imply that the distinction between simple and complex attributes is theoretically useless.[51] Moreover, even if some fine-grained conception of attributes should turn out to be more theoretically useful than any of its coarser-grained competitors, this would not yet mean that none of those competitors has enough theoretical utility to make it worthwhile to see how that distinction might be drawn if we have adopted one of them instead.

To forestall a further objection, I should start by conceding that (as has already been hinted at) the analysis of attribute simplicity to be proposed below will be constructed in a somewhat abstract or 'indirect' fashion. More specifically, the analysis will explicate its target not—or at least not exclusively—by talking directly about the respective attributes and their 'constituents', but rather by talking about the λ-expressions by which those attributes are denoted. The reason for this is twofold and entirely pragmatic. First, a λ-expression provides a very compact, and hence convenient, way of describing how a given attribute is logically related to other attributes.[52] And second, the indirect approach allows us to 'abstract away' from the ways in which complex attributes are metaphysically (as opposed to logically) related to their respective constituents. For example, in saying that a given property P is denoted by the λ-expression '$\lambda x\,(Q_1(x) \wedge Q_2(x))$' (relative to an interpretation and variable-assignment relative to which 'Q_1' and 'Q_2' respectively denote the two properties Q_1 and Q_2), we are in no way required to specify *how exactly* P is 'made up' from Q_1 and Q_2. For example, we do not have to specify whether P stands in some fundamental relation to *both* Q_1 and Q_2, or whether, instead, it is fundamentally related to some state of affairs or other kind of entity that has Q_1 and Q_2 as constituents.

intrinsically and extrinsically, this need not be taken to mean that '$P \vee Q$' and '$P \vee (\sim P \,\&\, Q)$' denote distinct properties. Instead, the truth or falsity of those predications may depend merely on the way in which the property $P \vee Q$ is respectively represented. And similarly, nothing compels us to treat the putative difference in truth-value between

(1) The fact that Max is a cube is grounded in the fact that Max is an accompanied cube

and

(2) The fact that Max is either a cube or an accompanied cube is grounded in the fact that Max is an accompanied cube

as indicating a difference between what is respectively *denoted* by 'the fact that Max is a cube' and 'the fact that Max is either a cube or an accompanied cube', rather than merely a difference in representation. (For considerations in favor of the representational relativity of grounding-talk, see Schnieder [2010] and Jenkins [2011]. Also cf. Schnieder [2006] and Krämer & Roski [2015].)

Having said all this, I should also note that there are cases in which it would be *absurd* to explain away a putative distinction between two attributes (or states of affairs) as a mere difference in representation. For example, suppose the argument to be countered by this sort of move has the conclusion that the property of *being red* is distinct from the property of *being green*, on the premise that there are things that have the one property but not the other. If we were to apply here the above strategy with a view to defending the (absurd) claim that 'being red' and 'being green' really pick out one and the same property P, we would have to say that, whenever we commonly regard a given entity x as having the property of *being red* but not the property of *being green*, all that is really going on is that x has the property P relative to the description 'being red' but not relative to the description 'being green'. But this latter claim is of course glaringly implausible. After all, the way we are struck by a green thing really *is* different from the way we are struck by a red thing, and this difference apparently requires explanation by the 'things themselves', rather than by mere differences in description.

51. The latter would be the case if, *e.g.*, one could derive from those postulates that under the proposed analysis there are no logically simple attributes at all.

52. Cf. Swoyer (1998, 300): "[T]he role of the syntactic structure of a complex predicate is not to exhibit the internal structure of a relation [where 'relation' applies in Swoyer's usage also to properties]; it is to disclose that relation's niche in the logical network of relations."

Related to the previous objection, it might be thought that under the to-be-proposed analysis, given its heavy use of λ-expressions, the logical simplicity or complexity of any particular attribute will be a matter of some highly contingent linguistic factors. For instance, since the analysis will presuppose the existence of 'linguistic items' such as constants, variables, and operators, it may seem to have the consequence that the simplicity of a given attribute depends on what constants and variables there are. But—the worry continues—what constants or variables there are should have *no bearing at all* on whether a particular attribute is complex or simple. This objection rests on the natural assumption that, where the analysis talks of variables and constants, etc., it will commit itself to the existence of 'linguistic items' in the usual sense of the phrase. However, to assume this is to overlook that we can also adopt a much more abstract conception of our formal language and its expressions. On such a conception, there will be no obstacle to identifying variables, constants, and other expressions with *pure sets*, roughly in the manner of Gödel codes.[53] It is this more abstract conception that should be taken to be operative in the following.

To guard against yet another objection, the next subsection develops an account of *mereological* simplicity that, rather than to talk directly about mereological fusions and their parts, is framed in terms of 'mereological descriptions'. The close analogy between this account and the analysis of attribute simplicity on offer in section 3.3 will hopefully assuage any suspicions that the attributes satisfying the to-be-proposed analysans cannot rightly be referred to as 'logically simple'.

3.2 *An 'Indirect' Account of Mereological Simplicity*

A natural way to frame a mereological description of a particular entity is to write down a list of some of its parts, such that the entity in question is the fusion of those parts. Since a "list of parts" is obviously only a list of *names* of parts, it may sometimes happen (unless the underlying language is 'Lagadonian') that a given list contains two names of one and the same entity. Also, since a list is not a set, it may happen that it contains the same name twice. As for the names themselves, there are various possibilities, but let us suppose that each name can take one of two forms: it can be either atomic (*i.e.*, a variable or constant), or it can itself be a mereological description. Finally, to avoid ambiguity, it is desirable that the list be preceded or otherwise accompanied by some operator whose job is to signal that the overall expression—*i.e.*, the list together with the operator—should be read as denoting the fusion of the listed objects.[54]

To develop a simple system of descriptions along these lines, let us make the following recursive stipulations:

- An *m-term* is either a variable, a constant, or a Σ-expression. (In addition, we call an m-term *atomic* iff it is either a variable or a constant.)
- An *m-list* is a parenthesis-enclosed, comma-delimited list of m-terms, of length ≥ 1 and finite 'depth' (as determined by the number of its levels of parentheses).
- A Σ-*expression* is an expression $\ulcorner \Sigma\, \varphi \urcorner$, for some m-list φ. (This m-list will also be referred to as the Σ-expression's *matrix*.)

Σ-expressions are designed to be the mereological counterparts of λ-expressions. That is, just as λ-expressions may be said to function as more or less detailed 'logical' descriptions of the attributes denoted by them, so Σ-expressions provide more or less detailed *mereo*logical descriptions of their respective referents.

Let us further stipulate that, for any interpretation I, variable-assignment g, and m-terms $t_1, t_2, \ldots$, a Σ-expression $\ulcorner \Sigma\, (t_1, t_2, \ldots) \urcorner$ has a denotation relative to I and g just in case each one of the t_i has a denotation relative to I and g, in which case that expression denotes$_{I,g}$

53. Cf. footnote 19 above.

54. The *fusion* of some entities $x_1, x_2, \ldots$ may here be taken to be the unique entity that has each one of the x_i as a part and is such that each one of its parts overlaps at least one of the x_i. (Cf. Hovda [2009].)

the fusion of those entities that are respectively denoted$_{I,g}$ by the t_i. For example, relative to an interpretation that assigns Socrates to 'Socrates' and Plato to 'Plato', the expression '$\Sigma\,(\text{Socrates}, \text{Plato})$' denotes the fusion of Socrates and Plato. We will adopt the usual axioms of classical mereology, and accordingly assume that there is no 'null individual'. Under this assumption, an entity x is *mereologically simple* (*i.e.*, atomic) iff x has no proper parts, or in other words: iff each part of x is identical with x. An adequate analysis of mereological simplicity in terms of Σ-expressions can therefore be stated as follows:

(MS) An entity x is *mereologically simple* iff, for any interpretation I, variable-assignment g, and Σ-expression S: if S denotes$_{I,g}$ x, then each m-term in S denotes$_{I,g}$ x.

So, for instance, if a Σ-expression $\ulcorner\Sigma\,(t_1, t_2)\urcorner$ is to denote a mereologically simple entity relative to an interpretation I and a variable-assignment g, then both t_1 and t_2 have to denote$_{I,g}$ x.

Although (MS) may admittedly not be a *cognitively* adequate analysis of mereological simplicity, it can nonetheless claim to be *a priori extensionally adequate*, provided that talk of m-terms and Σ-expressions is given a suitably abstract reading.[55] When the term 'adequate' is in the following used without qualification, it should be understood in precisely this sense. This is not to suggest, however, that all we can reasonably want of our analyses is that they should be *a priori* extensionally adequate. Among other things, we will ordinarily also want them to be free of circularity.

(MS) is useful in that it serves as a straightforward example of an 'indirect' analysis that is nonetheless adequate. However, it does not yet provide us with a way of seeing how one might construct an adequate account of what it is for an attribute to be *logically* simple.[56] To arrive at an analysis of mereological simplicity that is more suitable for our purposes, we may instead begin with the following, somewhat naïve account:

(MS_0') An entity x is *mereologically simple* iff, for any interpretation I, variable-assignment g, and Σ-expression S: if S denotes$_{I,g}$ x, then, for some atomic m-term t, the matrix of S is identical with $\ulcorner(t)\urcorner$.

By stipulating that t be an *atomic* m-term (so that, in particular, t cannot be a Σ-expression), this analysis tries to ensure that its right-hand side

55. See above, p. 19. To a first approximation, the phrase '*a priori* extensionally adequate' can be unpacked exactly as one would expect, viz., in such a way that it applies to a given analysis A iff it can be ascertained *a priori* that A's analysandum is coextensive with its analysans. (The role that is here given to the *a priori* is one of the things that distinguish the 'philosophical' type of analysis currently at issue from the 'logico-metaphysical' type that will be discussed in section 3.5 below.) The notion of the *a priori* is however notoriously unclear, and moreover it is open to doubt whether the extensional adequacy of a proposed philosophical analysis can be strictly speaking *ascertained* unless one already has a definition of the respective analysandum to hand (in which case extensional adequacy can be established by proving the equivalence of the definiens and the proposed analysans). For this reason, it would presumably be better to understand that phrase in a more liberal way, viz., as applying to exactly those analyses that are (i) extensionally adequate and (ii) such that their extensional adequacy does not depend on any empirical fact. (The relevant notion of dependence may, at least roughly, be taken to be that of counterfactual dependence, while an 'empirical fact' may be understood to be any fact outside of logic, mathematics, and semantics.) This eliminates any epistemological requirement: an analysis can count as '*a priori* extensionally adequate' even if its extensional adequacy is impossible to ascertain.

56. If one were to hew *very* close to (MS) in an attempt to analyze the notion of a logically simple attribute, one would arrive at something like the following:

($*$) An attribute A is *logically simple* iff, for any interpretation I, variable-assignment g, and λ-expression L: if L denotes$_{I,g}$ A, then each of the constants and variables that occur free in L denotes$_{I,g}$ A.

However, this analysis can be quickly seen to fail. For, given (P4) and the other assumptions of our framework, any property P is relative to some interpretation I and variable-assignment g denoted by a λ-expression $\ulcorner\lambda x\,(F(x) \vee G(x))\urcorner$, where F and G are constants or variables that respectively denote$_{I,g}$ $P\,\&\,Q$ and $P\,\&\,{\sim}Q$, for some suitable property Q (cf. footnote 48 above); and it will usually not be the case that P is identical with both $P\,\&\,Q$ and $P\,\&\,{\sim}Q$ (or indeed either of them). An analogous point can be made for attributes of any higher adicity. Hence, given (P4), every attribute whatsoever is under ($*$) classified as logically non-simple, which is clearly undesirable.

is not satisfied by just any old entity. With or without this restriction, however, the analysis suffers in fact from the *opposite* defect: its right-hand side fails to be satisfied by anything at all. This is because, for any given entity x, there exist an interpretation I, a variable-assignment g, and an atomic m-term t such that x is for example denoted$_{I,g}$ by $\ulcorner\Sigma\,(t,t)\urcorner$, whose matrix is distinct from $\ulcorner(t)\urcorner$. This is a very basic example of what will in the following be referred to as a **redundancy problem**: the misclassification of an entity as in some sense complex due to the redundant complexity of some of its representations.

At first blush, the present redundancy problem can be solved by replacing the 'identical with' in (MS_0') with an 'equivalent to', where two m-lists φ and ψ are considered *equivalent* just in case they contain exactly the same atomic m-terms. So, for example, the m-lists '(Socrates, Socrates)' and '(Socrates)' are equivalent in this sense, and the same goes for '(Socrates, Σ (Plato, Aristotle))' and '(Plato, Σ (Socrates, Aristotle))'. But this maneuver constitutes only a partial remedy. For suppose that u and v are two atomic m-terms that, relative to some interpretation I and variable-assignment g, both denote a certain entity x. As a result of the semantics of Σ-expressions, the substitution of coreferential m-terms within a given Σ-expression's matrix leaves the denotation of the overall expression unchanged. Hence, x will be denoted$_{I,g}$ not only by $\ulcorner\Sigma\,(u,u)\urcorner$ but also, *e.g.*, by $\ulcorner\Sigma\,(u,v)\urcorner$; and for any atomic m-term t (including u and v), $\ulcorner(u,v)\urcorner$ is evidently not equivalent to $\ulcorner(t)\urcorner$, given that u and v are distinct. This is an example of what I shall call an **obfuscation problem**: while the complexity of $\ulcorner\Sigma\,(u,v)\urcorner$ is no less redundant than that of $\ulcorner\Sigma\,(u,u)\urcorner$, the redundancy is in the former case obfuscated by the use of the additional m-term v. In general, an obfuscation problem can be regarded as a complication of the corresponding redundancy problem: to solve it, we have to refine our solution of the latter.

Let us say that a Σ-expression is *problematic* just in case there is no atomic m-term t such that the expression's matrix is equivalent to $\ulcorner(t)\urcorner$. The fact that any given entity is, as we have just seen, denoted by a problematic Σ-expression shows that (MS_0') is not an adequate analysis of mereological simplicity, even in the amended version in which the 'identical with' is replaced with 'equivalent to'. Fortunately, it is easy to see how (MS_0') can be repaired: before we classify a given entity as composite (*i.e.*, not mereologically simple) on account of its being denoted by a certain problematic Σ-expression, we must first check whether that expression cannot be transformed into an *un*problematic expression by way of replacing some of the m-terms in its matrix with other m-terms that are coreferential with the former. Whenever the expression can in this way be transformed into an unproblematic one, the fact that a given entity is denoted by it should not count against that entity's simplicity. Thus, for instance, the fact that in the previous paragraph the entity x is denoted$_{I,g}$ by $\ulcorner\Sigma\,(u,v)\urcorner$ should not count against the simplicity of x, for by replacing the occurrence of v with the coreferential$_{I,g}$ m-term u, we obtain the unproblematic expression $\ulcorner\Sigma\,(u,u)\urcorner$.

A plausible restriction on this sort of substitution requires that the to-be-replaced m-terms should always be *atomic*. This serves to rule out the replacement of Σ-expressions by atomic m-terms, which would otherwise make for an all-too easy way of turning problematic Σ-expressions into unproblematic ones. For example, if t is a constant that is relative to a given interpretation and variable-assignment coreferential with the Σ-expression 'Σ (Socrates, Plato)', then, by replacing the latter with t, the problematic nested Σ-expression 'Σ (Σ (Socrates, Plato))' could be turned into the unproblematic $\ulcorner\Sigma\,(t)\urcorner$. But it clearly seems unjustified to say that an entity's being denoted by 'Σ (Σ (Socrates, Plato))' should not count against that entity's simplicity, given that Socrates is distinct from Plato.

Based on the above discussion, we can now formulate an improved version of (MS_0') that constitutes an adequate analysis of mereological simplicity. To solve the obfuscation problem in the way suggested, it will be convenient first to introduce the auxiliary concept of an m-term's *reduction*. Roughly speaking, a reduction of a given m-term t is the result of replacing in t zero or more occurrences of atomic m-terms with coreferential other m-terms. More formally:

(MR) An m-term t' is a *reduction* of an m-term t relative to an interpretation I and a variable-assignment g iff t' is the result of replacing in t zero or more occurrences of atomic m-terms by occurrences of other m-terms in such a way that the following condition is satisfied:

($*$) For each replaced occurrence o and every m-term τ: if o is an occurrence of τ, then the m-term that replaces o is coreferential$_{I,g}$ with τ.

For example, if I is an interpretation and g a variable-assignment such that 'Bob' is coreferential$_{I,g}$ with 'Sam', then '$\Sigma\,(\mathrm{Sam}, \mathrm{Sam})$' will relative to I and g be a reduction of '$\Sigma\,(\mathrm{Bob}, \mathrm{Sam})$'. If t is an atomic m-term, then, by condition ($*$), any reduction of t relative to I and g will be either t itself or some other m-term that is coreferential$_{I,g}$ with t. More generally, condition ($*$) ensures that any reduction of an m-term t relative to I and g is coreferential$_{I,g}$ with t, provided that t itself has a denotation relative to I and g.

With the help of this concept of reduction, the notion of mereological simplicity can be analyzed as follows:

(MS'_1) An entity x is *mereologically simple* iff, for any interpretation I, variable-assignment g, and Σ-expression S: if S denotes$_{I,g}$ x, then, for some atomic m-term t, S has relative to I and g a reduction whose matrix is equivalent to $\ulcorner(t)\urcorner$.

Or equivalently (though slightly more complicated):

(MS'_2) An entity x is *mereologically simple* iff, for any interpretation I, variable-assignment g, and Σ-expression S: if S denotes$_{I,g}$ x, then there exist an interpretation $I' \supseteq I$, a variable-assignment $g' \supseteq g$, and an atomic m-term t such that S has relative to I' and g' a reduction whose matrix is equivalent to $\ulcorner(t)\urcorner$.

It is easy to see that both of these analyses are adequate.

Proof sketch: For the sake of brevity, we will focus on (MS'_1). Recall that an analysis counts as adequate just in case it is '*a priori* extensionally adequate'. (See footnote 55 for clarification.) So, to show that (MS'_1) is adequate, it suffices to give an *a priori* proof (more precisely: a proof whose soundness does not depend on any empirical fact) of the thesis that an entity is mereologically simple iff it satisfies the right-hand side of (MS'_1). The reader may convince herself that the following is such a proof.

For the left–to-right direction, let x be any mereologically simple entity, and let I, g, and S be (respectively) an interpretation, variable-assignment, and Σ-expression such that S denotes$_{I,g}$ x. Since x, being mereologically simple, is the fusion of no other things than x itself, it follows that

(1) Any atomic m-term in the matrix of S denotes$_{I,g}$ x.[57]

Further, from the syntax of Σ-expressions (in particular, from the fact that any m-list has to be of only finite 'depth' and contain at least one element), it can be inferred that

(2) The matrix of S contains at least one atomic m-term.

But from (1) and (2), we can conclude that there exists at least one atomic m-term t that denotes$_{I,g}$ x and is contained in the matrix of S. Again using (1), it follows that S has relative to I and g a reduction whose matrix is equivalent to $\ulcorner(t)\urcorner$. So any mereologically simple entity satisfies the right-hand side of (MS'_1).

For the right-to-left direction, we prove the contrapositive. Let x be any entity that is not mereologically simple, which is to say that x has a proper part y. There will then exist atomic m-terms t and u, as well as an interpretation I and a variable-assignment g, such that x and y are respectively denoted$_{I,g}$

57. This step of the argument makes essential use of the somewhat contentious assumption that there is no null individual, which may raise doubts about its *a priori* character. To address this worry, the framework described in this section would have to be altered in three respects. First, the notion of mereological simplicity would have to be redefined, to the effect that an entity is mereologically simple iff it has no proper parts other than the null individual. (Assuming that parthood is antisymmetric, which is presumably not in question, and given that any null individual would by definition be a part of everything whatsoever, there can be only one null individual, since any two of them would have to be part of each other.) Second, we would have to introduce a special constant, say, '0', and redefine the notion of interpretation by saying that an *interpretation* is a function F from constants to entities that satisfies the following constraint: if there is a null individual, then F assigns that individual to '0', and if there is none, then F assigns nothing to '0'. Third, we would have to redefine the notion of equivalent m-lists, to the effect that two m-lists are equivalent iff, ignoring occurrences of '0', they contain exactly the same atomic m-terms.

With these changes in place, it could again be shown, by an argument fairly similar to the present one (except for some additional steps) and without relying on the assumption that there is no null individual, that the analyses (MS'_1) and (MS'_2) are extensionally adequate.

by t and u. Moreover, since y is a part of x, the Σ-expression $\ulcorner\Sigma\,(t,u)\urcorner$ will also denote$_{I,g}$ x. Now let τ be any atomic m-term that denotes$_{I,g}$ x. Since y is only a *proper* part of x, no m-term coreferential$_{I,g}$ with u will contain an m-term denoting$_{I,g}$ x. *A fortiori*, no m-term coreferential$_{I,g}$ with u will contain τ. But by the syntax of Σ-expressions, any m-term must contain at least one atomic m-term. So any m-term coreferential$_{I,g}$ with u must contain some atomic m-term *other than* τ. This in turn means that $\ulcorner\Sigma\,(t,u)\urcorner$ will relative to I and g have no reduction whose m-list is equivalent to $\ulcorner(\tau)\urcorner$. From the way in which this has just been shown, we can now draw the following, more general conclusion: For *any* atomic m-term τ that denotes$_{I,g}$ x, the Σ-expression $\ulcorner\Sigma\,(t,u)\urcorner$ has relative to I and g no reduction whose matrix is equivalent to $\ulcorner(\tau)\urcorner$. But x is of course denoted$_{I,g}$ by $\ulcorner\Sigma\,(t,u)\urcorner$. Hence x fails to satisfy the right-hand side of (MS_1'), as required.

3.3 An Analogous Analysis of Attribute Simplicity

The development of (MS_2') provides a useful guideline for the analysis of the concept of a logically simple attribute. As above, we can begin with a fairly naïve approximation (notice the analogy to (MS_0')):

(AS_0) An attribute A is *logically simple* iff, for any interpretation I, variable-assignment g, and λ-expression L: if L denotes$_{I,g}$ A, then, for some atomic term F, the matrix of L is identical with $\ulcorner F(v_1, v_2, \ldots)\urcorner$, where $v_1, v_2, \ldots$ are (in this order) the λ-variables of L.[58]

By stipulating that F be an *atomic* term (so that, in particular, F cannot be a λ-expression), this analysis tries to ensure that its right-hand side is not satisfied by just any old attribute. With or without this restriction, however, the analysis suffers in fact from the *opposite* defect: under the assumptions of the present framework—notably (P4)—its right-hand side fails to be satisfied by anything at all. This is because, for any given property P (and analogously for any given relation), there exist an interpretation I, a variable-assignment g, and an atomic term F such that P is for example denoted$_{I,g}$ by the λ-expression $\ulcorner\lambda x\,(F(x) \wedge F(x))\urcorner$. The reason for this, in turn, lies in the fact that the formula $\ulcorner F(x) \wedge F(x)\urcorner$ is equivalent to $\ulcorner F(x)\urcorner$ (in the sense of 'equivalent' introduced in section 2.4). For, by (P3) and (P4) together with the semantics of λ-expressions, it follows from the equivalence of $\ulcorner F(x) \wedge F(x)\urcorner$ and $\ulcorner F(x)\urcorner$ that, relative to any given interpretation and variable-assignment, the λ-expressions $\ulcorner\lambda x\,F(x)\urcorner$ and $\ulcorner\lambda x\,(F(x) \wedge F(x))\urcorner$ will denote either the same property or nothing at all.[59] What this shows is that we have here again a redundancy problem, just as in the above discussion of (MS_0').

At first blush, this defect can be repaired by replacing the 'identical with' in (AS_0) with an 'equivalent to', where the relevant notion of equivalence is the one that is used in (P4).[60] However, this maneuver constitutes only a partial remedy. For suppose that G and H are two constants that, relative to some interpretation I and variable-

58. The specification "in this order" is here meant to imply that, for some formula φ, L is identical with $\ulcorner\lambda v_1, v_2, \ldots\,\varphi\urcorner$.

59. Cf. footnote 48 above.

60. Why not choose instead a concept of equivalence that is stronger or weaker than the one used in (P4)? The reason for not choosing a stronger concept (*i.e.*, one under which some pairs of formulas that are equivalent in the sense at issue in (P4) fail to count as equivalent) is relatively straightforward: choosing a stronger concept would only mean—if it would have any effect at all—that the proposed solution would fail to address all the cases that give rise to the present redundancy problem.

The reason for not choosing a *weaker*, *i.e.*, more liberal, concept (such as that of classical equivalence) is a little more subtle. In the first place, the use of a weaker concept would not leave our account of attribute simplicity with any fatal defects. However, it would (if it would have any effect at all) have the consequence that under the resulting account an *unnecessarily broad* class of attributes would be classified as logically simple. Although this would not be fatal, it would render the corresponding notion of simplicity markedly less significant. To use a crude analogy, it would be a bit as if we had decided to adopt a concept of mereological simplicity under which an entity counts as mereologically simple iff it has at most two proper parts. Moreover, if the alternative equivalence relation were chosen to be sufficiently weak, then a certain desirable sparseness result for logically simple attributes would no longer hold. (See section 3.5 below, in particular footnote 76.) For these reasons, the arguably best choice at this point is to use the same concept of equivalence as is used in (P4).

It should be noted that these considerations are independent from the question of which particular notion of equivalence is at play in (P4). Thus, if that postulate had made use of a different concept C, then the result of these considerations would have been that our account of attribute simplicity should similarly employ C.

assignment g, both denote a certain property P. As a result of the semantics of λ-expressions, the substitution of coreferential terms within the matrix of a given λ-expression leaves the denotation of the overall expression unchanged (provided that certain restrictions concerning bound variables are observed—see below). Hence, P will be denoted$_{I,g}$ not only by $\ulcorner \lambda x\,(G(x) \wedge G(x))\urcorner$ but also, *e.g.*, by $\ulcorner \lambda x\,(G(x) \wedge H(x))\urcorner$; and for any atomic term F (including G and H), the formula $\ulcorner G(x) \wedge H(x)\urcorner$ is obviously not equivalent to $\ulcorner F(x)\urcorner$, given that G and H are distinct. So we have here an obfuscation problem, and to solve it, we will have to refine our solution of the present redundancy problem.

Let us call a λ-expression L *problematic* just in case there is no atomic term F such that L's matrix is equivalent to $\ulcorner F(v_1, v_2, \ldots)\urcorner$, where $v_1, v_2, \ldots$ are (in this order) the λ-variables of L. The fact that any given attribute is, as we have just seen, denoted by a problematic λ-expression shows that (AS_0) should not be regarded as an adequate analysis of attribute simplicity, even in the amended version in which the 'identical with' is replaced with 'equivalent to'. Fortunately, it is easy to see how (AS_0) can be repaired: before we classify a given attribute as logically complex (*i.e.*, non-simple) on account of its being denoted, relative to some interpretation I and variable-assignment g, by a certain problematic λ-expression, we must first check whether that expression cannot be transformed into an *un*problematic expression by way of replacing some of the terms in its matrix with other terms that are coreferential$_{I,g}$ with the former. Whenever the expression can in this way be transformed into an unproblematic one, the fact that a given attribute is denoted by it should not count against that attribute's simplicity. Thus, for instance, the fact that in the previous paragraph the property P is denoted$_{I,g}$ by $\ulcorner \lambda x\,(G(x) \wedge H(x))\urcorner$ should not count against the simplicity of P, for by replacing the occurrence of H with the coreferential$_{I,g}$ constant G we obtain the unproblematic expression $\ulcorner \lambda x\,(G(x) \wedge G(x))\urcorner$.

A plausible restriction on this sort of substitution requires that the to-be-replaced terms should always be atomic. This serves to rule out the replacement of non-atomic terms—most notably λ-expressions—by atomic terms, which would otherwise make for an all-too easy way of turning problematic λ-expressions into unproblematic ones. For example, if F is a constant that is relative to a given interpretation and variable-assignment coreferential with the λ-expression '$(\lambda x\,(\text{black}(x) \wedge \text{horse}(x)))$', then, by replacing the latter with F, the problematic nested λ-expression '$\lambda x\,(\lambda x\,(\text{black}(x) \wedge \text{horse}(x)))(x)$' could be turned into the unproblematic expression $\ulcorner \lambda x\, F(x)\urcorner$. But it seems quite unjustified to say that an attribute's being denoted by '$\lambda x\,(\lambda x\,(\text{black}(x) \wedge \text{horse}(x)))(x)$' should not count against that attribute's simplicity.

Based on the above discussion, we can now formulate an improved version of (AS_0) that constitutes an adequate analysis of attribute simplicity. To solve the present obfuscation problem in the way suggested, it will be convenient first to introduce the auxiliary concept of a term's *reduction*. Roughly speaking, a reduction of a given term t is the result of replacing in t zero or more occurrences of atomic terms with coreferential other terms. More precisely:

(R) A term t' is a *reduction* of a term t relative to an interpretation I and a variable-assignment g iff t' is the result of replacing in t zero or more occurrences of atomic terms by occurrences of other terms, in such a way that the following three conditions are satisfied:

- (i) For each replaced occurrence o and any term τ: if o is an occurrence of τ, then the term with which o is replaced is coreferential$_{I,g}$ with τ.
- (ii) There is no substitution of bound variable-occurrences.
- (iii) No replacing term-occurrence contains a variable-occurrence that is 'captured' as a result of the replacement.[61]

61. Less succinctly, this condition can be put as follows: There is no term-occurrence o in t' that (i) replaces an occurrence of an atomic term in t and (ii) contains a variable-occurrence that is bound by an occurrence of a variable-binding operator (*i.e.*, either 'λ' or '$\exists$') within t' but outside of o.

For example, if I is an interpretation and g a variable-assignment such that 'Bob' is coreferential$_{I,g}$ with 'Sam', then '$R\,(\text{Sam}, \text{Sam})$' will relative to I and g be a reduction of '$R\,(\text{Bob}, \text{Sam})$'. The conditions (i)–(iii) ensure that any reduction of a term t relative to I and g is coreferential$_{I,g}$ with t itself, provided that t itself has a denotation relative to I and g. In this respect, (R) is exactly analogous to (MR), with the differences being merely a result of the fact that λ-expressions have a more complex syntax and semantics than Σ-expressions.

A few other consequences of (R) are also worth noting. First, relative to any interpretation and variable-assignment, every term is a reduction of itself. Second, if a term t' is a reduction of a term t, then any reduction of t' (relative to the same interpretation and variable-assignment) is likewise a reduction of t. We can thus say that, for any given interpretation and variable-assignment, the corresponding 'reduction relation' is both reflexive and transitive.[62]

With the help of the concept of reduction defined in (R), the notion of a logically simple attribute can now finally be analyzed as follows:

(AS) An attribute A is *logically simple* iff, for any interpretation I, variable-assignment g, and λ-expression L: if L denotes$_{I,g}$ A, then there exist an interpretation $I' \supseteq I$, a variable-assignment $g' \supseteq g$, and an atomic term F such that L has relative to I' and g' a reduction whose matrix is equivalent to $\ulcorner F(v_1, v_2, \ldots)\urcorner$, where $v_1, v_2, \ldots$ are (in this order) the λ-variables of L.

Thus, if P is a simple property, and if I, g, and L are (respectively) an interpretation, a variable-assignment, and a λ-expression such that L denotes$_{I,g}$ P, then there must exist an interpretation $I' \supseteq I$, a variable-assignment $g' \supseteq g$, and a λ-expression L' such that (i) L' is a reduction of L relative to I' and g', and (ii) for some atomic term F, the matrix of L' is equivalent to $\ulcorner F(v)\urcorner$, where v is the λ-variable of L (and hence also of L').

For a more concrete example, let P be the property of *being an unmarried man*, let g be some variable-assignment, and let I be an interpretation that respectively maps the two constants 'man' and 'married' to the property of *being a man* and the property of *being married*. As a result, P will relative to I and g be denoted by the λ-expression '$\lambda x\,(\text{man}(x) \wedge \neg\text{married}(x))$'. Let L be this λ-expression, and suppose we wish to know whether P might be logically simple. Given (AS), the answer depends (among other things) on whether, for some interpretation $I' \supseteq I$ and variable-assignment $g' \supseteq g$, the terms 'man' and 'married' can in L be replaced with other terms in such a way that (i) the resulting λ-expression is relative to I' and g' a reduction of L, and (ii) the matrix of this reduction is equivalent to $\ulcorner F(x)\urcorner$, for some atomic term F. If such a replacement is possible, then the fact that P is denoted by L will not count against the simplicity of P (which is not to say that P may not still fail to be simple); but if *no* such replacement should be possible, then P is indeed not simple. One way to paraphrase this is to say that the question of whether P is simple depends on the 'metaphysics' of the properties of *being a man* and *being married*. More specifically, as we will see in section 3.5 below, the question depends on whether these two properties are analyzable in terms of the identity relation and/or P itself.

To adjudicate whether (AS) is adequate, it is useful to assess the extent to which (AS) is analogous to (MS$'_2$). In comparing the two analyses, it will be apparent that the disanalogies are indeed few and limited, and can all be accounted for by the fact that λ-expressions have an in several ways more complicated syntax and semantics than Σ-expressions.[63] Due to the strong analogy between the two analyses, the attributes that satisfy the right-hand side of (AS) can be appropriately

62. This talk of 'reduction relations' should be understood in the mathematical sense of 'relation' (*i.e.*, as referring simply to classes of ordered pairs), rather than in the metaphysical sense that is at issue in the rest of this paper.

63. To wit, there are the following disanalogies:

referred to as in some sense simple, while all others can be appropriately referred to as in the corresponding sense complex. And the particular sense in which these attributes *are* simple or complex will presumably best be referred to as 'logical' (or perhaps 'logico-metaphysical'), given that the operators with the help of which λ-expressions are constructed are themselves probably best characterized as *logical* operators. As a result, (AS) can reasonably be regarded as extensionally adequate: for, by what has just been said, the attributes that satisfy its right-hand side are appropriately referred to as logically simple (and can thus be regarded as falling under its analysandum), while all others are appropriately referred to as logically complex. Moreover, since we have reached this result by *a priori* reasoning, (AS) can be regarded as adequate *tout court* (in the sense specified on p. 20 above).

- The definition of 'reduction' for terms, *i.e.*, (R), differs from its mereological counterpart (MR) (p. 22 above) by two additional clauses. The need for these clauses arises quite straightforwardly from the fact that λ-expressions can (and do) contain variable-binding operators.
- On the right-hand side of (MS_2'), the Σ-expression S is required to have a reduction whose matrix is equivalent to $\ulcorner(t)\urcorner$, but on the right-hand side of (AS), the λ-expression L is instead required to have a reduction whose matrix is equivalent to $\ulcorner F(v_1, v_2, \ldots)\urcorner$, where $v_1, v_2, \ldots$ are the λ-variables of L. The need for this additional complication stems simply from the fact that 'λ' is, unlike 'Σ', a variable-binding operator.

In addition, the reader will have noticed that, in developing (AS), we have left out the step that would have been analogous to the construction of (MS_1'). If this step had not been omitted, we would first have constructed an analysis of attribute simplicity whose right-hand side (in analogy to the right-hand side of (MS_1')) would not have quantified over any supersets of the respective interpretation I and variable-assignment g. But that analysis would have been inadequate, due to the following disanalogy between Σ-expressions and λ-expressions. If a Σ-expression denotes a given mereologically simple entity x relative to an interpretation I and a variable-assignment g, then that Σ-expression will have to contain at least one atomic term that also denotes$_{I,g}$ x. (This is, in a nutshell, the reason why (MS_1') is equivalent to (MS_2').) By contrast, due to our postulate (P4), a λ-expression that denotes a given attribute A need not contain any atomic term that also denotes A. (Cf. footnote 56 above.) So, to prevent (AS) from having the consequence that no attribute is logically simple, the right-hand side of (AS) must in its consequent quantify over interpretations and/or variable-assignments that are supersets of, respectively, I and g. (For reasons of symmetry, I have here chosen to quantify over supersets *both* of I and of g.)

Apart from being adequate, (AS) has the desirable feature of being flexible enough to yield sensible results when combined with different conceptions of attributes. For instance, suppose we wish to adopt a very fine-grained conception on which attributes are distinct from their double negations. The proper way to do this with minimal changes to the rest of the framework would be to modify our definition of 'entails' (p. 13 above) in such a way that, under the revised definition, any formula φ is inequivalent to its double negation $\ulcorner\neg\neg\varphi\urcorner$. Such a change would in the first place lead to a weakening of (P4), with the desired result that any attribute A would count as distinct from $\sim\sim A$. In addition, since the term 'equivalent' is also used in (AS), and the sense it has there should be taken to be the same as the sense it has in (P4) (for reasons described in footnote 60 above), the revised definition of 'entails' would have the further consequence that the double negation of a simple attribute would no longer count as simple.[64]

Of course, no matter how coarse- or fine-grained a conception of attributes (AS) is combined with, it will not tell us anything very concrete as to *which* attributes are logically simple. For unless we supply some concrete information about the metaphysics of specific attributes, we will not receive any verdict as to whether, say, the property of *being an electron* is logically simple, any more than we can expect to learn from an adequate account of mereological simplicity whether cats or horses are mereologically simple. However, it *can* reasonably be asked whether the account, when combined with the framework of section 2, allows us to show that the class of the simple attributes is closed under certain operations, and also whether it allows us to show that the

64. As may be expected, the same goes for the *single* negation of a simple attribute. For with (P4) weakened as indicated, the proof for the thesis (C2) in section 3.4 below, according to which the negation of a simple attribute (other than the identity relation) is again simple, will not go through.

simple attributes are in some ways sparse. These questions will be addressed in the following two subsections.

3.4 *Closure*

A relatively straightforward consequence of (AS) is to the effect that the converse of any simple dyadic relation is also simple.

Proof sketch: Let R be any simple dyadic relation; we have to show that the converse of R is also simple. Suppose first that R is identical with the identity relation, $\mathtt{I}$. Given that, for any variables u and v, $\ulcorner u = v \urcorner$ is equivalent to $\ulcorner v = u \urcorner$, it follows from (P3) and (P4) (together with the semantics of λ-expressions) that $\mathtt{I}$ is identical with its own converse. Hence, given that R is identical with $\mathtt{I}$, R is identical with its own converse. But R is simple, and so it follows that the converse of R is simple.

Suppose next that R is distinct from $\mathtt{I}$. We will use the following corollary of (AS):

(ASC) Any attribute A other than the identity relation is logically simple iff, for any interpretation I, variable-assignment g, constant F, and λ-expression L: if A is denoted$_{I,g}$ by both F and L, then L has relative to I and g a reduction whose matrix is equivalent to $F(v_1, v_2, \ldots)$, where $v_1, v_2, \ldots$ are (in this order) the λ-variables of L.

That this follows from (AS) can be seen as follows. Consider the biconditional $\ulcorner \varphi$ iff $\psi \urcorner$, where φ is the right-hand side of (AS) and ψ the right-hand side of (ASC). To show that (ASC) follows from (AS), it is enough to show that this biconditional follows from the assumption that A is an attribute other than the identity relation. So let A be any such attribute. The left-to-right direction of the mentioned biconditonal can then easily be verified using the transitivity of reduction together with the fact that, relative to any interpretation, the constant '$\mathtt{I}$' denotes no attribute other than the identity relation (and hence does not denote A), and the right-to-left direction is trivial.[65]

Let now I, g, F, and L be, respectively, an interpretation, a variable-assignment, a constant, and a λ-expression such that the converse of R, *i.e.*, $\lambda x, y\, R(y, x)$, is denoted$_{I,g}$ by both F and L. Let φ be the matrix of L, and let v_1 and v_2 (in this order) be L's λ-variables, so that $L = \ulcorner \lambda v_1, v_2\, \varphi \urcorner$. By (AS), it is sufficient to show that L has relative to I and g a reduction whose matrix

65. To see why (ASC) has to include the qualification "other than the identity relation", note that, for any variables u and v, the formula $\ulcorner \mathtt{I}(u, v) \urcorner$ is equivalent to $\ulcorner \mathtt{I}(u, v) = \mathtt{I}(u, u) \urcorner$. As a result, it follows from the assumptions of our framework that the identity relation $\mathtt{I}$ is denoted by '$\lambda x, y\, (\mathtt{I}(x, y) = \mathtt{I}(x, x))$' (or, written without the use of '=': '$\lambda x, y\, \mathtt{I}(\mathtt{I}(x, y), \mathtt{I}(x, x))$'). Let L be this λ-expression, let I and g be an interpretation and a variable-assignment, and let F be a non-logical constant. Clearly, L does not have a reduction relative to I and g whose matrix is equivalent to $\ulcorner F(x, y) \urcorner$, but it *does* have a reduction (viz., itself) whose matrix is equivalent to '$\mathtt{I}(x, y)$'. Hence, without the requirement that A be distinct from $\mathtt{I}$, the right-hand side of (AS) will not imply the right-hand side of (ASC).

Given that $\mathtt{I}$ is denoted by '$\lambda x, y\, \mathtt{I}(\mathtt{I}(x, y), \mathtt{I}(x, x))$', it follows from the semantics of λ-expressions that the negation of the identity relation, *i.e.*, $\sim\mathtt{I}$ or the 'distinctness relation', is denoted by '$\lambda x, y \neg \mathtt{I}(\mathtt{I}(x, y), \mathtt{I}(x, x))$'. By what has just been said, this latter λ-expression, call it '$L^{\sim}$', does not have a reduction (relative to any interpretation and variable-assignment) whose matrix is equivalent to $\ulcorner F(x, y) \urcorner$, for any atomic term F. Hence, the distinctness relation is not logically simple in the sense of (AS). It would be otherwise if, for some constant D, we had defined the concept of interpretation in such a way that an interpretation is a (partial) function from constants to entities that maps '$\mathtt{I}$' to $\mathtt{I}$ and D to $\sim\mathtt{I}$; for in this case, the matrix of $L^{\sim}$ would have been equivalent to $\ulcorner D(x, y) \urcorner$. We can thus see that the logical simplicity of the distinctness relation is under (AS) made to depend on an arbitrary feature of our framework or, more precisely, on an arbitrary feature of our definition of 'interpretation', which is arguably undesirable. To remedy this defect, we could weaken the right-hand side of (AS) to obtain the following, more complicated analysis:

(AS$'$) An attribute A is *logically simple* iff, for any interpretation I, variable-assignment g, and λ-expression L: if L denotes$_{I,g}$ A, then there exist an interpretation $I' \supseteq I$, a variable-assignment $g' \supseteq g$, a formula φ, and an atomic term F that satisfy the following two conditions:

(i) L has relative to I' and g' a reduction whose matrix is equivalent to φ.

(ii) For any variables $v_1, v_2, \ldots$: if $v_1, v_2, \ldots$ are (in this order) the λ-variables of L, then $\ulcorner \lambda v_1, v_2, \ldots \varphi \urcorner$ has relative to I' and g' a reduction whose matrix is equivalent to $\ulcorner F(v_1, v_2, \ldots) \urcorner$.

To see how this repairs the mentioned defect of (AS), consider again $L^{\sim}$, and let φ be the formula $\ulcorner \neg \mathtt{I}(x, y) \urcorner$: the matrix of $L^{\sim}$ is then equivalent to φ, so that, *a fortiori*, $L^{\sim}$ has a reduction (relative to any interpretation and variable-assignment) whose matrix is equivalent to φ; and for a suitable interpretation I, variable-assignment g, and atomic term F (where F denotes$_{I,g}$ $\sim\mathtt{I}$), the λ-expression $\ulcorner \lambda x, y\, \varphi \urcorner$ has relative to I and g a reduction $\ulcorner \lambda x, y \neg(\lambda x, y \neg F(x, y))(x, y) \urcorner$, whose matrix is equivalent to $\ulcorner F(x, y) \urcorner$. Thus, the fact that the distinctness relation is denoted by $L^{\sim}$ does not pose an obstacle to its being classified as logically simple under (AS$'$).

Insofar as (AS$'$) fails to be sensitive to an arbitrary feature of our framework, it is arguably superior to (AS). On the other hand, it is also significantly more complicated, and the flaw of (AS) that is here under discussion only seems to affect (in the present framework) the classification of the distinctness relation. For this reason, I propose that we stick with (AS).

is equivalent to $\ulcorner F(v_1, v_2)\urcorner$. To this end, let φ' be the result of replacing in φ the free occurrences of v_1 with occurrences of v_2 and *vice versa*, if necessary after a renaming of bound variables, so that any replacing occurrences of v_1 and v_2 are also free.[66] Since L denotes$_{I,g}$ the converse of R, it follows from the semantics of λ-expressions that $\ulcorner \lambda v_2, v_1\, \varphi'\urcorner$ also denotes$_{I,g}$ the converse of R; and so $\ulcorner \lambda v_1, v_2\, \varphi'\urcorner$ denotes$_{I,g}$ R itself. Let L' be this latter λ-expression, let G be a constant not in the domain of I (and hence not occurring in L or L', nor identical with '$\mathtt{I}$'), and let I' be an interpretation that is just like I, except that it maps G to R. (Conceiving of interpretations as classes of ordered pairs, we thus have $I' = I \cup \{(G, R)\}$.) Since L' denotes R relative to I and g, it follows that L' denotes R also relative to I' and g. Moreover, since R is simple and distinct from $\mathtt{I}$, it follows from (ASC) that L' has relative to I' and g a reduction whose matrix is equivalent to $\ulcorner G(v_1, v_2)\urcorner$. It is then clear, given how L' results from L, that L has relative to I' and g a reduction L'' whose matrix is equivalent to $\ulcorner G(v_2, v_1)\urcorner$. Let now φ'' be the matrix of L'', and let ψ be the result of replacing in φ'' every occurrence of G with $\ulcorner(\lambda x, y\, F(y, x))\urcorner$. Since G is coreferential$_{I',g}$ with $\ulcorner(\lambda x, y\, F(y, x))\urcorner$, and since F is a constant, it follows that $\ulcorner \lambda v_1, v_2\, \psi\urcorner$ is a reduction of L''—and hence of L—relative to I' and g. Since G occurs neither in ψ nor in L, $\ulcorner \lambda v_1, v_2\, \psi\urcorner$ is a reduction of L also relative to I and g. Moreover, since G is distinct from the logical constant '$\mathtt{I}$', and since φ'' is equivalent to $\ulcorner G(v_2, v_1)\urcorner$, it can be seen that ψ is equivalent to $\ulcorner(\lambda x, y\, F(y, x))(v_2, v_1)\urcorner$ and hence to $\ulcorner F(v_1, v_2)\urcorner$. So L has relative to I and g a reduction whose matrix is equivalent to $\ulcorner F(v_1, v_2)\urcorner$, as required.

This result can be straightforwardly generalized using the concept of an attribute's 'permutations', which is itself a generalization of the concept of a dyadic relation's converse:

(P) A *permutation* of a κ-adic attribute A is any attribute that is, relative to some interpretation I and variable-assignment g, denoted by $\ulcorner \lambda v_1, v_2, \ldots\, F(u_1, u_2, \ldots)\urcorner$, where F is a constant that denotes A relative to I and g, and where the v_i and u_i are, respectively, κ-many pairwise distinct variables such that $\{u_1, u_2, \ldots\} = \{v_1, v_2, \ldots\}$.

In the special case in which A is dyadic, it has two permutations, viz., A itself and its converse. The above result can now be given the following, more general form, whose proof can be easily obtained from the previous one:

66. For example, if φ is $\ulcorner(G(v_1) \wedge \exists v_1\, H(v_1, v_2))\urcorner$, for two constants G and H, then φ' will be $\ulcorner(G(v_2) \wedge \exists u\, H(u, v_1))\urcorner$, for some variable u distinct from v_1.

(C1) For any logically simple relation R, any permutation of R is logically simple.

Similarly it can be shown that:

(C2) For any logically simple attribute A other than the identity relation, the negation of A is logically simple.

Proof sketch: The proof is structurally nearly identical to the second part of the previous one. To avoid clutter, we will consider only the case in which A is a simple dyadic relation; the argument will easily generalize to other adicities.

Let I, g, F, and L be, respectively, an interpretation, a variable-assignment, a constant, and a λ-expression such that the negation of A is denoted$_{I,g}$ by both F and L. Let φ be the matrix of L, and let v_1 and v_2 (in this order) be L's λ-variables, so that $L = \ulcorner \lambda v_1, v_2\, \varphi\urcorner$. By (AS), it is sufficient to show that L has relative to I and g a reduction whose matrix is equivalent to $\ulcorner F(v_1, v_2)\urcorner$. Since L denotes$_{I,g}$ the negation of A, it follows straightforwardly from the assumptions of our framework (in particular (P3) and (P4)), together with the semantics of λ-expressions, that $\ulcorner \lambda v_1, v_2 \neg\varphi\urcorner$ denotes A itself. Let now G be a constant not in the domain of I (and hence neither occurring in φ nor identical with '$\mathtt{I}$'), and let I' be an interpretation that is just like I, except that it maps G to A. Since the λ-expression $\ulcorner \lambda v_1, v_2 \neg\varphi\urcorner$ denotes A relative to I and g, it denotes A also relative to I' and g. Moreover, since A is simple and distinct from $\mathtt{I}$, it follows from (ASC) (see previous page) that $\ulcorner \lambda v_1, v_2 \neg\varphi\urcorner$ has relative to I' and g a reduction whose matrix is equivalent to $\ulcorner G(v_1, v_2)\urcorner$. It is then clear that L, *i.e.*, $\ulcorner \lambda v_1, v_2\, \varphi\urcorner$, has relative to I' and g a reduction L' whose matrix is equivalent to $\ulcorner \neg G(v_1, v_2)\urcorner$. Let φ' be the matrix of L', and let ψ be the result of replacing in φ' every occurrence of G with $\ulcorner(\lambda x, y \neg F(x, y))\urcorner$. Since G is relative to I' and g coreferential with $\ulcorner(\lambda x, y \neg F(x, y))\urcorner$ and F a constant, it follows that $\ulcorner \lambda v_1, v_2\, \psi\urcorner$ is a reduction of L'—and hence of L—relative to I' and g. Given that G occurs neither in ψ nor in L, $\ulcorner \lambda v_1, v_2\, \psi\urcorner$ is a reduction of L also relative to I and g. Moreover, since G is distinct from '$\mathtt{I}$' and φ' is equivalent to $\ulcorner \neg G(v_1, v_2)\urcorner$, it can be seen that ψ is equivalent to $\ulcorner \neg(\lambda x, y \neg F(x, y))(v_1, v_2)\urcorner$ and hence to $\ulcorner F(v_1, v_2)\urcorner$. So L has relative to I and g a reduction whose matrix is equivalent to $\ulcorner F(v_1, v_2)\urcorner$, as required.

Thus, the class of the logically simple attributes is closed under permutation, and the class of the logically simple attributes (other than identity) is closed also under negation.[67] This latter result marks a significant contrast between the concept of a logically simple attribute and

67. N.B.: This talk of the "logically simple attributes other than identity" is not meant to presuppose that the identity relation is logically simple.

Lewis's concept of perfect naturalness, for Lewis was rather strongly committed to the view that the negation of a perfectly natural attribute is not itself perfectly natural.[68] It should be kept in mind, however, that (C2) is not a consequence of (AS) alone, but rather a consequence of (AS) in combination with our moderately coarse-grained conception of attributes. If Lewis had embraced a suitably fine-grained conception of attributes, for instance along the lines sketched in his (1986, 56f.), then he *might* conceivably have found (AS) to be an adequate account of perfect naturalness; for under such a conception it would no longer be true that the double negation of an attribute is identical with the attribute itself, so that the proof for (C2) could not be completed. But in the end, Lewis did not adopt a fine-grained conception of attributes, while he *was* consistently committed to the view that the negations of perfectly natural attributes are not themselves perfectly natural. So it remains problematic to think of Lewis's notion of a perfectly natural attribute as coextensive with that of a simple attribute (even apart from the fact that, in Lewis's system, properties and relations are not *sui generis* entities but rather set-theoretic constructions over a space of possibilia).

An objector might at this point argue that, since (AS) in combination with the assumptions of our framework leads to the result that the negation of a simple attribute (other than identity) is again simple, there must be something wrong either with (AS) or with our framework.[69] As far as I can see, this objection may take one of two forms.

First, it may be based on the idea that the negation of a simple attribute *just cannot* be simple. To this, it can be replied that it is not clear on what grounds one could hold that "the negation of a simple attribute just cannot be simple", unless the objector takes it to be independently plausible that some Armstrongian (or hyper-Armstrongian[70]) theory of universals is true, and identifies the simple attributes with those universals. But whether any such theory *is* true is itself very much an open question. It accordingly seems to be good policy, where possible, not to rely on such a theory—a policy that Lewis adhered to, for example, when he tried to explicate his concept of perfect naturalness.[71]

Second and more interestingly, the objection may rest on the premise that there is a very deep and important distinction to be drawn between 'positive' and 'negative' attributes, and that it would be desirable if this distinction were reflected by our account of attribute simplicity (*e.g.*, in such a way that only 'positive' attributes can count as simple). I suspect that something like this worry will be at the root of at least some philosophers' skepticism about (C2). An effective way to respond to it would be to ask where the conviction comes from that "there is a very deep and important distinction to be drawn between positive and negative attributes". Possibly, it comes again from a commitment to something like an Armstrongian theory of universals. In this case, it can once more be replied that it would be good to see how far we can get without relying on such a theory. But alternatively (and perhaps additionally), the conviction may come from some empirical observations that militate in favor of a sharp divide between positive and negative attributes.

The thesis that there must be some such distinction to be drawn is admittedly hard to deny, for it seems clear that there must be *something* that makes it natural to regard properties such as *being blue*, *being an electron*, and *being human* as 'positive' while their respective negations are regarded as 'negative'. But it is not clear how metaphysically deep this distinction really cuts. The very fact that the divide between positive and negative attributes is such a pretheoretically obvious feature of our *Lebenswelt* should probably be taken as a sign that the distinction is not exactly fundamental. (For how likely is it that truths of fundamental metaphysics can be so easily read off the phenomena?[72])

68. *E.g.*, see his (2009, 204).

69. Thanks to an anonymous reviewer for pressing this issue.

70. Cf. p. 17 above.

71. For similar considerations, see section 3.1 above.

72. There is also the question of how widely the distinction is applicable: the

Moreover, one can arguably give an account of this distinction that is at least *prima facie* more plausible, insofar as it does not presuppose that we have such easy access to deep metaphysical facts. On this account, what makes a given attribute 'positive' or 'negative' is simply the role it plays in the natural order of things (to the extent that we have epistemic access to it). For example, as a matter of physical law or possibly because of the metaphysics of the relevant properties, it never happens that something that is blue all over is also yellow all over, but things that are non-blue can as easily be yellow as they can be non-yellow. Likewise, it never happens that an electron is also a human being, but there are many non-electrons that are human and even more that are non-human. What examples like these suggest is that the attributes we are tempted to call 'positive' are just those that are more restrictive than their respective negations as to what other attributes they can—within the bounds of nomological necessity—be coinstantiated with. If this points at least roughly in the direction of an adequate account of the distinction between positive and negative attributes, then there seems to be no compelling reason to think that the distinction has anything to do with logical simplicity.

3.5 Analyzability and Sparseness

Whereas (C2) implies that the simple attributes are in at least one way more abundant than one might have expected, a further consequence of (AS) implies that the simple attributes are in at least one important way *sparse*. To formulate this consequence, we first have to introduce some further terminology. It will be useful to begin with a basic notion of analyzability:

(A) An entity x is *analyzable in terms of* an entity y iff there exist an interpretation I, a variable-assignment g, and terms t and u such that t contains a free occurrence of u and x and y are respectively denoted$_{I,g}$ by t and u.

Thus, if the property of *being a bachelor* is (relative to some interpretation) denoted by '$\lambda x\,(\text{man}(x) \wedge \neg\text{married}(x))$', where the constant 'married' denotes (relative to that same interpretation) the property of *being married*, then we can derive from this that the property of *being a bachelor* is analyzable in terms of the property of *being married*.

The concept of analyzability, as defined in (A), can be fairly described as 'logico-metaphysical'. For if two entities x and y are respectively denoted by two terms t and u, and t contains a free occurrence of u, then t will have to be some formula or λ-expression, unless it is identical with u. This contrasts with any notion of analyzability that can be given (in some framework other than the present one) a definition that is on the one hand analogous to (A) but on the other hand such that t is in it allowed to be, *e.g.*, a set-expression $\ulcorner\{\ldots, u, \ldots\}\urcorner$ or some expression $\ulcorner F(\ldots, u, \ldots)\urcorner$, where F may denote an arbitrary function. It also contrasts, for obvious reasons, with any *aprioristic* notion of analyzability, *i.e.*, with any notion under which an entity x counts as 'analyzable' in terms of an entity y just in case it is somehow *a priori* that x can be adequately analyzed in terms of y. For, to take a standard example, it is hardly an *a priori* fact that the property of *being water* is analyzable in terms of such properties as *being oxygen*, *being a proton*, *being a quark*, etc.[73]

predicate 'has eternal life' rings positive while 'is immortal' rings negative—but apparently, both predicates correspond to the very same property. (Cf. Frege [1919, 150f.].)

73. Nominalists might here object that *science* has never told us anything at all about the alleged property of *being water*; instead, it has merely told us that 'to be water is to be H_2O', or perhaps even only that every portion of water is at the same time a portion of H_2O. This latter hypothesis, however—that all we've been told is that 'water' and 'H_2O' are coextensive—can presumably be dismissed on the basis of considerations like those given by Rayo (2013, 18ff.). This leaves us with the first hypothesis: that we've learned nothing more than that 'to be water is (or: "just is") to be H_2O'. This hypothesis is certainly correct insofar as science does not explicitly commit us to the existence of properties. However, *given* that we accept an abundant ontology of attributes, and in particular assuming that there exists such a thing as the property of *being water*, the at first blush most natural way in which a claim such as 'To be water is to be H_2O' may be understood (or at

It may be worth pointing out that the present concept of logico-metaphysical analyzability is rather broad. For instance, every entity is under (A) analyzable in terms of itself; and given (P3) and (P4), it further follows that every attribute and every state of affairs is analyzable in terms of its own negation as well as in terms of the identity relation.[74] The notion is thus considerably broader than our ordinary usage of 'analyzable' may suggest. Unsurprisingly, it is *transitive*: if an entity x is analyzable in terms of some entity that is again analyzable in terms of some entity y, then x is likewise analyzable in terms of y.

In the next step, we have to introduce the notion of an entity's being 'fully analyzable' in terms of a set of entities:

(FA) An entity x is *fully analyzable* in terms of a set S iff the following two conditions are satisfied:

(i) x is analyzable in terms of each member of S.

(ii) There exist an interpretation I, a variable-assignment g, and a term t such that: x is denoted$_{I,g}$ by t, and every variable or constant that has a free occurrence in t denotes$_{I,g}$ a member of S.

Trivially, every entity x is fully analyzable in terms of the set $\{x\}$; and no entity whatsoever is fully analyzable in terms of the empty set. (The latter fact is in part due to the syntax of formulas and λ-expressions, which requires that each formula or λ-expression contain at least one free occurrence of a variable or constant.) Further, if x is an attribute or state of affairs that is fully analyzable in terms of some set S, then x will also be fully analyzable in terms of $S \cup \{\mathtt{I}\}$; but the converse need not hold.[75]

Finally, on the basis of the concept just introduced, we can define the concept of a 'purely logical' entity:

(PL) An entity is *purely logical* iff it is fully analyzable in terms of $\{\mathtt{I}\}$.

any rate, the best way in which it may be assimilated into our framework) is to understand it as saying that the property of *being water* is identical with the property of *being H_2O*, which is another way of saying that *being water* is the property of *consisting of some number of oxygen atoms and twice that number of hydrogen atoms, bound together in such-and-such ways*. (This is a simplification; but of course, to say that to be water is to be H_2O is itself fairly inaccurate—cf. Tahko [2015] and references therein.)

This procedure of assimilating 'just is' statements (to use Rayo's term) to 'property identities' is not the only way in which we can understand such claims. It had better not; for, as Dorr (MS) points out (essentially repeating an argument by Correia [2006, 761f.]), trying to understand the following claim in terms of properties leads straight into trouble:

($*$) To be a non-self-instantiator is (or "just is") to fail to instantiate oneself.

If one were to interpret this as stating that the property of *being a non-self-instantiator* is identical with the property of *failing to instantiate oneself*, one would be led into Russell's paradox. Dorr takes this to be a reason for adopting a primitive 'just is' operator. But arguably, this move will not be forced upon us as long as we have at our disposal the concept of a state of affairs. For it seems that, with the help of this concept, ($*$) can tolerably well be paraphrased as follows:

(†) For any entity x and any state of affairs s: x is a non-self-instantiator in s iff x fails to instantiate itself in s.

The locution 'x φs in s' that is employed in this paraphrase can be read as shorthand for 's is a sufficient condition for its being the case that x φs'. It is then a further question whether this notion of sufficiency should be regarded as a primitive (presumably of one's semantic framework, along with, *e.g.*, the concept of denotation), or whether it can be explicated in other terms. At any rate, the notion can be expected to be useful in providing something like a semantics for the monadic predicates of the language under consideration. For instance, we can say that, for any entity x and any state of affairs s: s is a sufficient condition for its being the case that x *is a non-self-instantiator* iff either (a) x is not a property and s is identical with $s \wedge (x = x)$, or (b) x is a property and s is identical with $s \wedge \neg x(x)$.

To someone who has adopted a 'just is' operator as a piece of primitive terminology, any proposal to paraphrase ($*$) by something along the lines of (†) may appear misguided from the outset. But such a way of looking at the matter seems to be warranted *only* if one has sufficient reason to think that a primitive 'just is' operator is in fact needed.

74. For example, any state of affairs s is identical with $s \wedge (s = s)$, and any property P is identical with $\lambda x\,(P(x) \wedge (x = x))$.

75. To forestall a possible misunderstanding, note that '$\mathtt{I}$' but not '$\{\mathtt{I}\}$' is an expression of the formal language described in section 2. Like all expressions of that language, '$\mathtt{I}$' should here be interpreted in accordance with the semantics specified in section 2.3, and should thus be taken to denote the identity relation. Consequently, '$\{\mathtt{I}\}$' should be taken to denote the set that has the identity relation as its only member.

The purely logical entities include, for example, the identity relation, the property of *being self-identical*, and the state of affairs that everything is self-identical. Clearly, every purely logical entity is either an attribute or a state of affairs.

Our sparseness claim can now be stated as follows:

(S) For any attribute A and for any entity x: if A is analyzable in terms of x, but x is neither purely logical nor fully analyzable in terms of $\{A, \mathtt{I}\}$, then A is logically complex.

Proof sketch: Let A and x be as described in the antedecent of (S), and suppose for *reductio* that A is logically simple. Since A is analyzable in terms of x, there will exist an interpretation I, a variable-assignment g, a λ-expression L, and a constant c such that the following three conditions are satisfied:

(1) L denotes$_{I,g}$ A.

(2) c denotes$_{I,g}$ x.

(3) L contains a free occurrence of c.

Given that A is simple, it follows from (1) by (AS) that there exist an interpretation $I' \supseteq I$, a variable-assignment $g' \supseteq g$, a λ-expression L', and an atomic term F such that:

(4) L' is relative to I' and g' a reduction of L, and

(5) The matrix of L' is equivalent to $\ulcorner F(v_1, v_2, \ldots)\urcorner$, where $v_1, v_2, \ldots$ are the λ-variables of L (as well as of L').[76]

From (1) and (4), it now follows that L' denotes$_{I',g'}$ A. Hence, given (5) and our postulates (P3) and (P4) (together with the semantics of λ-expressions), we have that

(6) F denotes$_{I',g'}$ A.

Meanwhile, from (2)–(4), we can infer that there exists a term t (not necessarily distinct from c) such that:

(7) t denotes$_{I',g'}$ x, and

(8) L' results from L by replacing (possibly among other things) a free occurrence of c with t.

Given (4) and (8), we further have:

(9) For any term τ: if τ occurs free in t, then τ occurs free also in the matrix of L' and is not identical with any of L's λ-variables.

And from (5) we can infer that:

(10) For any atomic term τ: if τ occurs free in the matrix of L', then τ is identical with '$\mathtt{I}$' or F or one of L's λ-variables.

Together with (9), this yields:

(11) For any atomic term τ: if τ occurs free in t, then τ is identical with either '$\mathtt{I}$' or F.

From (6), (7), and (11), it now follows that

(12) x is either purely logical or fully analyzable in terms of $\{A\}$ or $\{A, \mathtt{I}\}$.

But by hypothesis, x is neither purely logical nor fully analyzable in terms of $\{A, \mathtt{I}\}$. Hence,

(13) x is fully analyzable in terms of $\{A\}$.

Now this means that x is either identical with A or denoted (relative to a suitable interpretation and variable-assignment) by a formula or λ-expression. In either case, x is an attribute or state of affairs that is fully analyzable in terms of $\{A\}$. But then it follows that x is fully analyzable in terms of $\{A, \mathtt{I}\}$, contrary to hypothesis. This completes the *reductio*.

As a corollary, we can note the following:

(SC) For any attribute A and any entities x and y: if A is analyzable in terms of both x and y, neither x nor y is purely logical, and at least one of x and y is not analyzable in terms of the other, then A is logically complex.

Proof: Let A, x, and y be as described in the antecedent of (SC), and suppose for *reductio* that A is logically simple. Given that A is simple and analyzable in terms of both x and y, it follows from (S) that both x and y are either purely logical or fully analyzable in terms of $\{A, \mathtt{I}\}$. But by hypothesis, neither x nor y is purely logical, and so it follows that both x and y must be analyzable in terms of A. By the transitivity of analyzability, we then have that x and y are analyzable in terms of each other, contrary to hypothesis.

76. As in the rest of this paper, the relevant notion of equivalence is the one specified in section 2.4 above. (Cf. also footnote 60, p. 23.) This is of importance here because the inference that yields (10) below would not go through if we were instead operating with the classical concept of equivalence.

Thus, if A and B are any two κ-adic attributes of which neither is purely logical, and at least one of which is not analyzable in terms of the other, then their various Boolean combinations—such as $(A \,\&\, B)$, $(A \vee B)$, and $(A \equiv B)$—are all logically complex. For example, suppose (as seems plausible) that the properties of *being unmarried* and *being a man* both fail to be purely logical, and that at least one of them fails to be analyzable in terms of the other. It then follows from (SC) that the property of *being an unmarried man*, *i.e.*, the conjunction of the two properties in question, is complex.[77]

As this example illustrates, (S) does not deliver any verdict on the complexity of a given attribute without some input as to what entities are analyzable in terms of which attributes. An objector might now insist that something more needs to be said as to *how one can tell* whether a given entity is analyzable in terms of such-and-such attributes. She may grant that we already have some fairly good idea of how scientific inquiry can lead to the insight that the property of *being water* is analyzable in terms of the properties of *being hydrogen* and *being oxygen*, or to the realization that the property of *being green* is analyzable in terms of certain attributes having to do with the reflection (as well as emission and transmission) of visible light. But on the face of it, the scientific methods that yield such results do not seem likely to allow us to *rule out* any thesis to the effect that a given attribute or state of affairs is analyzable in terms of such-and-such attributes. For example, it is far from clear what scientific methods could be used to rule out the thesis that the property of *being green* is identical with the 'gruesome' disjunction of, on the one hand, the property of *being green and observed before the year 2050 AD* and, on the other hand, the property of *being green and observed during or after the year 2050 AD*. The putative problem lies of course in the fact that this disjunction is *a priori* coextensive with the property of *being green*, so that it is on the face of it unclear how anyone might, by scientific methods or otherwise, drive an epistemic wedge between the two.

In response to this worry, it is tempting to suggest that it sees an epistemological problem where there is only a problem of semantics. For it might be replied that the question is not: 'How can we tell whether the property of *being green* is analyzable in terms of such-and-such other properties?', but rather: 'What would be a reasonable property for our predicate "green" to pick out?'. And a reasonable semantics for a given natural language will presumably not assign to the predicates of that language any needlessly complex semantic values. In particular, a reasonable semantics of English will not assign to the word 'green' a property $(P \,\&\, Q) \vee (P \,\&\, {\sim}Q)$ where P alone would do just fine.[78]

77. The same conclusion can also be reached by a slightly different route. For one may plausibly suppose that the property of *being unmarried* (or alternatively, the property of *being a man*) is neither purely logical nor fully analyzable in terms of $\{\text{bachelor}, \mathtt{1}\}$, where 'bachelor' denotes the property of *being an unmarried man*. It then follows directly from (S) that *being an unmarried man* is complex.

78. Some readers may suspect that the interpretive maxim at play here—viz., that one should avoid unmotivated complexity in the semantic values one ascribes to the expressions of a language—can also be appealed to in arguing that a naïve student of arithmetic should normally be regarded as engaged in addition rather than 'quaddition'. (For discussion of related ideas, see Kripke [1982, 38f.] and, *e.g.*, Lewis [1983, 374–77], Humphrey [1999], and Williams [2007]. The issue is also drawn attention to by Sider [2011, §3.2].) But the matter is hardly straightforward. In the first place, it is not obvious what kind of semantic value is under discussion. If the semantic value of 'plus' is the class of all tuples (x, y, z) such that $z = x + y$, then this value is on the face of it far more complex than, *e.g.*, the class of all tuples (x, y, z) such that: $x, y \in \{1, \ldots, 500\}$; $z = x + y$ for $x, y \in \{1, \ldots, 100\}$; and $z = 5$ for $x, y \in \{101, \ldots, 500\}$. For this second class has only finitely many members. Nor is the situation any clearer if 'plus' is instead taken to refer to an algorithm, for what goes on in a person's head when doing arithmetic can presumably be quite complicated (even on the 'algorithmic level'): for instance, he or she may employ different strategies depending on whether the numbers involved are large or small, etc. For these reasons, the idea that one should interpret the naïve student of arithmetic as engaged in addition rather than quaddition simply "because addition is simpler" strikes me as dubious. Still, it is arguable that considerations of simplicity have some role to play in this connection. In particular, they may plausibly be appealed to in deciding which features of a given cognitive process should be considered 'malfunctions'. (This might be further elaborated into something like a Lewisian best-system account of cognitive semantics.) But this

I think that this response is by and large correct—*but only* if the objector is right in assuming that the property of *being green* is distinct from the mentioned disjunction. On reflection, however, it is not so obvious that this assumption is true. After all, our physical knowledge is far from complete. So it is conceivable (albeit somewhat implausible) that, as we dig deeper into the physics and metaphysics of *being green* and *being observed before the year 2050 AD*, it will eventually emerge that any entity in terms of which the latter property is analyzable is also an entity in terms of which *being green* is analyzable. And if this should be the case, then, given the assumptions of our framework (notably, again, (P3) and (P4)), it will have to be conceded that *being green* is indeed identical with the 'gruesome' disjunction mentioned above.[79]

To conclude this section, let us very briefly consider whether there might be 'logically gunky' attributes and states of affairs. Since, by definition, an entity is mereologically gunky iff it has no atomic (*i.e.*, mereologically simple) parts,[80] we should analogously say that an attribute or state of affairs is *logically gunky* iff it is not analyzable in terms of any logically simple attribute or state of affairs.[81] *Prima facie* it may well seem epistemically possible that there should exist some attributes or states of affairs that are gunky in this sense. In the present framework, however, the hypothesis that there exists such 'logical gunk' conflicts with the rather plausible assumption that the identity relation is logically simple.[82] In fact, it even conflicts with the weaker assumption that the identity relation is logically non-gunky. For, as already mentioned, every attribute and state of affairs is in the present framework analyzable in terms of the identity relation. So, if the identity relation is logically non-gunky, then, by the transitivity of 'is analyzable in terms of', every attribute or state of affairs is logically non-gunky as well.

4. Conclusion

In the previous section, we have developed an account of logical simplicity as applied to attributes. With the help of the assumptions of the framework described in section 2, we have then derived some consequences of this account. In particular, we have shown that the permutations and negations of logically simple attributes (at least where the latter are distinct from the identity relation) are again simple, and we have also derived a useful sparseness result.

From here, we can go on in a number of directions. Most pressingly, we can (and should) investigate in detail the various philosophical applications to which the distinction between simple and complex attributes lends itself. Most *temptingly*, we can (and should) explore what other distinctions can be drawn by following a similar approach. To begin with the latter, a particularly low-hanging fruit is the distinction between simple and complex *states of affairs*. This distinction can be drawn in almost perfect analogy to the way in which we have (in section 3.3) drawn the distinction between simple and complex attributes. For we can say that

(SAS) A state of affairs s is *logically simple* iff, for any interpretation I, variable-assignment g, and formula φ: if φ denotes$_{I,g}$ s, then there

---

would clearly not be an application of the maxim in question, which only concerns the simplicity of semantic values.

79. The word 'gruesome' arguably constitutes something of an exaggeration here, since the disjunction in question is *a priori* coextensive with the (relatively) well-behaved property of *being green*. It is thus a far cry from such paradigmatically 'gruesome' attributes as the property of *being grue*.

80. Cf. Lewis (1991, 20).

81. An analysis of the concept of a logically simple state of affairs will be given at the beginning of the next section.

82. To see the plausibility of this assumption, suppose it were false. Then the identity relation would have to be denoted, relative to some interpretation I and variable-assignment g, by some λ-expression L that does relative to I and g *not* have a reduction whose matrix is equivalent to $\ulcorner \mathtt{I}(v_1, v_2) \urcorner$, where v_1 and v_2 are L's λ-variables. But what might L be? Or, more pointedly: what could the attributes or other entities be that are made reference to in L's matrix—entities in terms of which the identity relation is supposedly analyzable? It is extremely difficult to see what those entities could be, and this perplexity can hardly be attributed to some lack of knowledge as to how identity can be analyzed. After all, it does not seem as if the present concept of identity were particularly thick and mysterious.

exist an interpretation $I' \supseteq I$, a variable-assignment $g' \supseteq g$, and an atomic term t such that φ has relative to I' and g' a reduction that is equivalent to $\ulcorner \neg\neg t \urcorner$.[83]

Somewhat further afield, we can try to put our framework to use in developing analyses of a range of concepts that correspond to different kinds of complex attribute, such as that of a disjunctive attribute, a relational property, or an extrinsic property.

But more urgent than these further projects is (arguably) an investigation into the philosophical work that can be done with the help of the distinction between simple and complex attributes. The probably most salient question is whether our distinction can do the same work as Lewis's distinction between perfectly natural and less-than-perfectly natural attributes. Two potential applications that fall under this rubric have already been hinted at in the Introduction: first, a best-system account of lawhood, and second, an account of duplication along the lines of Lewis (1986, 61). A third potential application, in meta-semantics, has been alluded to in the previous section (p. 33). Let us consider the first of these applications in somewhat greater detail.

83. For the relevant concept of reduction, see p. 24 above; for the relevant concept of equivalence, see p. 13. The double negation sign on the last line should be taken to be nothing more than a technical trick, needed here merely to obtain from t a formula that is guaranteed to denote the same state of affairs as t itself (as long as t itself denotes a state of affairs). To be sure, part of what guarantees this is our postulate (P4), and if (P4) were weakened to allow for a very fine-grained conception of attributes, then some other solution would have to be found. For instance, one might stipulate that, for any term t, the expression $\ulcorner [t] \urcorner$ is a formula, and has a denotation relative to an interpretation I and a variable-assignment g if and only if t denotes$_{I,g}$ a state of affairs, in which case $\ulcorner [t] \urcorner$ denotes$_{I,g}$ that same state of affairs. The '$\neg\neg t$' on the last line of (SAS) could then be replaced with '$[t]$'.

It may also be worth pointing out in this connection that what Russell (1918/19, §1) calls 'atomic facts'—*i.e.*, facts "which consist in the possession of a quality by some particular thing"—are *not* logically simple in the present sense. For all we know, there might not exist any simple states of affairs at all; but it is hard to imagine what form (short of an oracle) any *evidence* for or against their existence might take.

As mentioned in the Introduction, Lewis's appeal to the concept of a perfectly natural attribute in the context of his account of lawhood was motivated by the need to avoid the consequence that an axiom like '$\forall x\, F(x)$' would be included in an ideal system, where the atomic predicate 'F' stands, *e.g.*, for a property P that is instantiated by all and only those things that exist in the actual world. This motivation has force regardless of whether one buys into an ontology of possible worlds or accepts the existence of a property P as just described. For essentially the same problem will also arise if P is merely a 'mighty' property, where a κ-adic attribute A is *mighty* just in case there exists a λ-expression L such that (i) A is relative to some interpretation and variable-assignment denoted by L and (ii) the matrix of L, or some variant of it, provides an extremely detailed description of the universe.[84,85] (Just how detailed the description would have to be is admittedly no precise matter.) The axiom '$\forall x\, F(x)$' would then still be very strong, and may strictly imply various regularities that we would not wish to regard as laws of nature. In Lewis's account, this problem is dealt with by the requirement that the respective system's primitive vocabulary should only refer to perfectly natural attributes. Suppose now that we instead adopt the analogous requirement in which 'perfectly natural' is replaced with 'logically simple'. If this modified requirement is to solve the problem, it will have to turn out that all mighty properties are logically complex.

84. Since L is a λ-expression, each of its λ-variables will have to occur free in its matrix (see p. 7 above), and so any 'description' that L's matrix provides of the world will have to take the form of an open sentence. Accordingly, the term 'description' should here be understood in a somewhat loose sense.

85. If φ is the matrix of L, then a 'variant' of φ should be taken to be any formula that is entailed by either $\ulcorner \varphi \wedge \psi \urcorner$ or $\ulcorner \neg\varphi \wedge \psi \urcorner$, where ψ is some comparatively simple further formula. (As a special case, $\ulcorner \neg\varphi \urcorner$ is a variant of φ.) To see the need for this generalization, suppose that φ is equivalent to a certain conditional $\ulcorner \psi \rightarrow \chi \urcorner$, where ψ is comparatively simple and χ (rather than φ itself) provides a "detailed description of the universe". The property denoted by $\ulcorner \lambda v\, \varphi \urcorner$ (where v is L's λ-variable) will then pose essentially the same problem for a best-system account of lawhood as would be posed if φ itself provided that description.

Does our account have this consequence? In light of the sparseness result (S) of section 3.5, the answer seems to be 'yes'; for that result suggests that mighty properties are indeed complex. Recall that, according to a corollary of (S), no simple attribute is analyzable in terms of any two entities of which neither is purely logical and at least one fails to be analyzable in terms of the other.[86] Meanwhile, if a given attribute is to count as 'mighty', then, by definition, that attribute will have to be denoted (relative to some interpretation and variable-assignment) by a λ-expression whose matrix is such that either it or some variant of it provides an extremely detailed description of the universe. It is reasonable to expect that most formulas of this sort will make reference to at least two attributes that are such that neither of them is purely logical and at least one of them fails to be analyzable in terms of the other. (A *prima facie* plausible example for such a pair of attributes would be the property of *being a photon* and the relation of *being one meter apart*.) If L is a λ-expression whose matrix is a formula of this sort, then (S) entails, via the mentioned corollary, that the attribute denoted by L is logically complex.

In this way, (S) establishes the logical complexity of a very broad range of mighty attributes. Although there is more to be said here, these considerations hopefully begin to make it seem plausible that the concept of a logically simple attribute can serve as a workable substitute for the concept of a perfectly natural attribute in the context of a Lewis-style account of lawhood. And if it is workable, it is arguably superior: for as we have seen, the concept of a logically simple attribute is amenable to analysis, and that analysis requires neither a primitive notion of naturalness or fundamentality, nor an Armstrongian (or otherwise 'sparse') ontology of universals. In addition, there is an epistemological argument in the offing. For as long as we operate with a primitive concept of perfect naturalness, it will be hard to see why scientific inquiry should ever tend to yield theories whose basic predicates pick out perfectly natural attributes.[87]

By contrast, the concept of a logically simple attribute fits seamlessly into the (at least *prima facie*) attractive conception of scientific inquiry as a quest for theories that are at the same time powerful, elegant, and true. Granted, we may never be able to know with certainty that a given attribute is in fact logically simple, since we may never be able definitely to exclude the possibility of further analysis. But if a scientist (or team of scientists) hits upon a theory that is not only elegant, powerful, and empirically adequate, but also—like the famous Eightfold Way classification of elementary particles—formulated in terms of attributes that, for all we know, might well be logically simple (*vulgo*: unanalyzable), then it should be reasonable to expect that those attributes will tend to be on the logically simple side, rather than to be extremely complex.[88] The rationale for this is simply that simpler

86. For the definitions of 'analyzable' and 'purely logical', see the first half of section 3.5.

87. Cf. Cohen & Callender (2009, 13):

> The fundamentalist [*i.e.*, the adherent of Lewis's best-system account of lawhood with its reliance on the concept of perfect naturalness] might [defend her view] by adapting an argument used to defend scientific realism. [...] Consider, for instance, the famous "no miracles" argument for realism that infers from the remarkable success of science to its truth. Even if one accepts this inference, this doesn't give the fundamentalist what she needs. For it is consistent with the correctness of the inference that the generalization $\ulcorner(x)(Fx \supset Gx)\urcorner$ is successful and true in both possible worlds considered, but a law in only one of them. The argument must be not only that success and truth are correlated (as per the standard no-miracles argument) but that success and perfect naturalness are correlated. We don't see any reason to believe this is so.

Apart from this consideration, I borrow from Cohen and Callender also the example of the 'Eightfold Way' classification of particles that will be made use of in the next paragraph.

88. This formulation evidently makes use of a concept of *degrees* of logical complexity. Fortunately, given the analysis of logical simplicity presented in section 3.3, it is now not difficult to see how this concept might be explicated. A natural first step would be to construct an analysis of the 'binary' notion of a logically complex attribute, which can be obtained by simply negating the right-hand side of that earlier analysis:

theories of the universe are *ceteris paribus* to be preferred over more complex ones.

The best-system account of lawhood that results from replacing Lewis's notion of perfect naturalness with the concept of a logically simple attribute can be regarded as a first step toward answering the larger question of what role the latter concept should play in a general theory of fundamentality. Plausibly enough, logically simple attributes are in some sense fundamental. But what exactly is their place in the overall picture? Should we aim for some unifying account of fundamentality in which a single notion takes center stage?[89] Or can we live equally well (if not better) with a more heterogeneous account that is instead based on several more or less disparate notions?[90] If we can help it, we should arguably avoid the use of new and esoteric primitives, but it remains to be seen whether a comprehensive account of fundamentality can be constructed without them.

Finally, there is the question of how best to deal with Lewis's concept of a perfectly natural attribute if it should indeed turn out not to serve any theoretical purpose that is not equally well or better served by our concept of a logically simple attribute.[91] Ignoring minor variations, there seem to be two possible views that one might take. On the one hand, one might hold that in the imagined scenario, given the relative obscurity of Lewis's concept, there would simply be no further need for it, so that the best way forward would be to let it fall into disuse. On the other hand, one might hope that the concept can be 'saved' through explication. In this vein, one might for instance propose that we recognize the perfectly natural attributes as precisely those that are both logically simple and 'positive' in the sense hinted at in section 3.4. If this or some alternative analysis of the concept of a perfectly natural attribute were to become generally accepted, then there would presumably be no reason not to use the term 'perfectly natural' in accordance with that analysis. But if, on the contrary, *no* analysis were to become generally accepted, then the concept would apparently retain its accustomed obscurity, and it would again seem best—as on the first view—to abandon its use.[92]

(AC) An attribute A is *logically complex* iff there exist an interpretation I, a variable-assignment g, and a λ-expression L such that (i) L denotes$_{I,g}$ A, and (ii) for any interpretation $I' \supseteq I$, variable-assignment $g' \supseteq g$, and atomic term F: L has relative to I' and g' no reduction whose matrix is equivalent to $\ulcorner F(v_1, v_2, \ldots) \urcorner$, where $v_1, v_2, \ldots$ are (in this order) the λ-variables of L.

In addition, one would need a suitable quantitative concept of complexity for formulas. A fairly straightforward option would be to say that, for any cardinality κ, a formula φ is *of complexity* κ iff κ is the number of occurrences of logical operators within φ. In the next step, one could adapt the right-hand side of (AC) to obtain the following account:

(DC) An attribute A *has a logical complexity of at least* κ iff there exist an interpretation I, a variable-assignment g, and a λ-expression L such that (i) L denotes$_{I,g}$ A, and (ii) for any interpretation $I' \supseteq I$, variable-assignment $g' \supseteq g$, and atomic term F: L has relative to I' and g' no reduction whose matrix is equivalent to a formula of a complexity less than κ.

On this basis, the *logical complexity* of an attribute A may be defined as the least upper bound of those cardinalities κ that are such that A has a logical complexity of at least κ.

89. Such a view has recently been advanced, *e.g.*, by Schaffer (2009) and Sider (2011). Along with Rosen (2010) and others, Schaffer maintains that the notion of metaphysical grounding should be regarded as primitive. By contrast, Correia (2013) explores ways in which metaphysical grounding might be analyzed in terms of essence.

90. Cf., *e.g.*, Koslicki (2015).

91. Although this is at the present stage a rather hypothetical question, it can arguably still be useful to consider, if only to shed some light on the dialectical thrust of this paper.

92. The central ideas of section 3.3 and a version of the framework described in section 2 can already be found in my dissertation, written at Washington University in St. Louis under the supervision of Thomas Sattig. I am further indebted to Claudio Calosi, Nikk Effingham, Kit Fine, Dan Marshall, Gideon Rosen, Nathan Wildman, the participants of an *eidos* seminar at which an earlier version of this paper was discussed (in particular, Akiko Frischhut, Ghislain Guigon, and Jan Walker), and especially to Fabrice Correia and Alex Skiles. I am also grateful to six anonymous referees (two each for *Erkenntnis*, *Mind*, and this journal) and to audiences at a workshop in Cambridge and at the Eighth European Congress of Analytic Philosophy in Bucharest. The writing of this paper was generously supported by the DAAD during a stay at the University of Tübingen, as well as by the Swiss National Science Foundation (project CRSII1–147685).

References

Ainsworth, P. M. (2009). Newman's Objection. *The British Journal for the Philosophy of Science*, *60*, 135–71.

Armstrong, D. M. (1978). *A Theory of Universals*. Cambridge: Cambridge University Press.

Armstrong, D. M. (1983). *What Is a Law of Nature?* Cambridge: Cambridge University Press.

Armstrong, D. M. (1997). *A World of States of Affairs*. Cambridge: Cambridge University Press.

Bader, R. M. (2013). Towards a Hyperintensional Theory of Intrinsicality. *The Journal of Philosophy*, *110*, 525–63.

Bartley, C. (2001). Udayana. In O. Leaman (ed.), *Encyclopedia of Asian Philosophy* (pp. 542–44). London: Routledge.

Bynoe, W. (2011). Against the Compositional View of Facts. *Australasian Journal of Philosophy*, *89*, 91–100.

Carnap, R. (1942). *Introduction to Semantics*. Cambridge, MA: Harvard University Press.

Carnap, R. (1950). Empiricism, Semantics, and Ontology. *Revue internationale de philosophie*, *4*, 20–40. Reprinted with minor changes in Carnap, *Meaning and Necessity*, 2nd edition, Chicago: The University of Chicago Press, 1956.

Carroll, J. W. (1994). *Laws of Nature*. Cambridge: Cambridge University Press.

Cohen, J. & Callender, C. (2009). A Better Best System Account of Lawhood. *Philosophical Studies*, *145*, 1–34.

Correia, F. (2006). Generic Essence, Objectual Essence, and Modality. *Noûs*, *40*, 753–67.

Correia, F. (2013). Metaphysical Grounds and Essence. In M. Hoeltje, B. Schnieder, & A. Steinberg (eds.), *Varieties of Dependence* (pp. 271–96). Munich: Philosophia.

Cresswell, M. J. (1975). Hyperintensional Logic. *Studia Logica*, *34*, 25–38.

Dorr, C. (2002). *The Simplicity of Everything*. PhD thesis, Princeton University.

Dorr, C. (2004). Non-Symmetric Relations. In D. W. Zimmerman (ed.), *Oxford Studies in Metaphysics*, volume 1 (pp. 155–92). Oxford: Oxford University Press.

Dorr, C. (MS). To Be F Is to Be G. Unpublished manuscript.

Dorr, C. & Hawthorne, J. (2013). Naturalness. In K. Bennett & D. W. Zimmerman (eds.), *Oxford Studies in Metaphysics*, volume 8 (pp. 3–77). Oxford: Oxford University Press.

Eddon, M. (2011). Intrinsicality and Hyperintensionality. *Philosophy and Phenomenological Research*, *82*, 314–36.

Eddon, M. & Meacham, C. J. G. (2015). No Work for a Theory of Universals. In B. Loewer & J. Schaffer (eds.), *A Companion to David Lewis* (pp. 116–37). Malden, MA: Wiley Blackwell.

Eklund, M. (2008). The Picture of Reality as an Amorphous Lump. In T. Sider, J. Hawthorne, & D. W. Zimmerman (eds.), *Contemporary Debates in Metaphysics* (pp. 382–96). Malden, MA: Blackwell.

Fine, K. (2000). Neutral Relations. *The Philosophical Review*, *109*, 1–33.

Fine, K. (2010). Towards a Theory of Part. *The Journal of Philosophy*, *107*, 559–89.

Frege, G. (1919). Die Verneinung. *Beiträge zur Philosophie des deutschen Idealismus*, *3/4*, 143–57.

Goodman, N. (1955). *Fact, Fiction, and Forecast*. Cambridge, MA: Harvard University Press.

Hacking, I. (1967). Possibility. *The Philosophical Review*, *76*, 143–68.

Heil, J. (2003). *From an Ontological Point of View*. Oxford: Oxford University Press.

Hodes, H. T. (1982). The Composition of Fregean Thoughts. *Philosophical Studies*, *41*, 161–78.

Hovda, P. (2009). What Is Classical Mereology? *Journal of Philosophical Logic*, *38*, 55–82.

Humphrey, J. A. (1999). Quine, Kripke's Wittgenstein, Simplicity, and Sceptical Solutions. *The Southern Journal of Philosophy*, *37*, 43–55.

van Inwagen, P. (2006). Names for Relations. *Philosophical Perspectives*, *20*, 453–77.

Jech, T. (2002). *Set Theory* (2nd ed.). Berlin: Springer.

Jenkins, C. S. (2011). Is Metaphysical Dependence Irreflexive? *The Monist*, *94*, 267–76.

Kaplan, D. (1995). A Problem in Possible-World Semantics. In W. Sinnott-Armstrong, D. Raffman, & N. Asher (eds.), *Modality, Morality, and Belief: Essays in Honor of Ruth Barcan Marcus* (pp. 41–52). Cambridge: Cambridge University Press.

Koslicki, K. (2015). The Coarse-Grainedness of Grounding. In K. Bennett & D. W. Zimmerman (eds.), *Oxford Studies in Metaphysics*, volume 9 (pp. 306–44). Oxford: Oxford University Press.

Krämer, S. & Roski, S. (2015). A Note on the Logic of Worldly Ground. *Thought*, *4*, 59–68.

Kripke, S. A. (1982). *Wittgenstein on Rules and Private Language*. Cambridge, MA: Harvard University Press.

Langton, R. & Lewis, D. K. (1998). Defining 'Intrinsic'. *Philosophy and Phenomenological Research*, *58*, 333–45.

Lewis, D. K. (1973). *Counterfactuals*. Malden, MA: Blackwell.

Lewis, D. K. (1983). New Work for a Theory of Universals. *Australasian Journal of Philosophy*, *61*, 343–77.

Lewis, D. K. (1986). *On the Plurality of Worlds*. Oxford: Blackwell.

Lewis, D. K. (1991). *Parts of Classes*. Oxford: Blackwell.

Lewis, D. K. (1994). Humean Supervenience Debugged. *Mind*, *103*, 473–90.

Lewis, D. K. (2001). Redefining 'Intrinsic'. *Philosophy and Phenomenological Research*, *63*, 381–98.

Lewis, D. K. (2009). Ramseyan Humility. In D. Braddon-Mitchell & R. Nola (eds.), *Conceptual Analysis and Philosophical Naturalism* (pp. 203–22). Cambridge, MA: MIT Press.

Lewy, C. (1976). *Meaning and Modality*. Cambridge: Cambridge University Press.

Loewer, B. (2007). Laws and Natural Properties. *Philosophical Topics*, *35*, 313–28.

MacBride, F. (2005). The Particular–Universal Distinction: A Dogma of Metaphysics? *Mind*, *114*, 565–614.

Maudlin, T. (2007). *The Metaphysics within Physics*. Oxford: Oxford University Press.

Menzel, C. (1993). The Proper Treatment of Predication in Fine-Grained Intensional Logic. *Philosophical Perspectives*, *7*, 61–87.

Mill, J. S. (1882). *A System of Logic*. New York: Harper.

Montague, R. (1969). On the Nature of Certain Philosophical Entities. *The Monist*, *53*, 159–94.

Nolan, D. (2014). Hyperintensional Metaphysics. *Philosophical Studies*, *171*, 149–60.

Olson, K. R. (1987). *An Essay on Facts*. Stanford, CA: CSLI.

Oppy, G. (1997). The Philosophical Insignificance of Gödel's Slingshot. *Mind*, *106*, 121–42.

Parry, W. T. (1933). Ein Axiomensystem für eine neue Art von Implikation (analytische Implikation). *Ergebnisse eines mathematischen Kolloquiums*, *4*, 5–6.

Parry, W. T. (1972). Comparison of Entailment Theories. *Journal of Symbolic Logic*, *37*, 441–42.

Plantinga, A. (1974). *The Nature of Necessity*. Oxford: Clarendon.

Plate, J. (2012). *Essence and Ontology*. PhD thesis, Washington University in St. Louis.

Pollock, J. L. (1967). The Logic of Logical Necessity. *Logique et Analyse*, *10*, 307–23.

Pollock, J. L. (1984). *The Foundations of Philosophical Semantics*. Princeton, NJ: Princeton University Press.

Ramsey, F. P. (1928). Universals of Law and of Fact. In D. H. Mellor (ed.), *F. P. Ramsey: Philosophical Papers* (pp. 140–4). Cambridge: Cambridge University Press.

Ray, G. (2014). The Problem of Negative Existentials Inadvertently Solved. In M. García-Carpintero & G. Martí (eds.), *Empty Representations* (pp. 262–74). Oxford: Oxford University Press.

Rayo, A. (2008). On Specifying Truth-Conditions. *The Philosophical Review*, *117*, 385–443.

Rayo, A. (2009). Towards a Trivialist Account of Mathematics. In O. Bueno & Ø. Linnebo (eds.), *New Waves in Philosophy of Mathematics* (pp. 239–60). New York: Palgrave MacMillan.

Rayo, A. (2013). *The Construction of Logical Space*. Oxford: Oxford University Press.

Rayo, A. & Uzquiano, G. (2006). *Absolute Generality*. Oxford: Oxford University Press.

Rosen, G. (2010). Metaphysical Dependence: Grounding and Reduction. In B. Hale & A. Hoffmann (eds.), *Modality* (pp. 109–35). Oxford: Oxford University Press.

Russell, B. (1918/19). The Philosophy of Logical Atomism. *The Monist*, *28–29*. Reprinted in R. C. Marsh (ed.), *Logic and Knowledge*, London: Routledge, 1956.

Schaffer, J. (2004). Two Conceptions of Sparse Properties. *Pacific Philosophical Quarterly*, *85*, 92–102.

Schaffer, J. (2009). On What Grounds What. In D. J. Chalmers, D. Manley, & R. Wasserman (eds.), *Metametaphysics* (pp. 347–83). Oxford: Oxford University Press.

Schnieder, B. (2006). 'By Leibniz's Law': Remarks on a Fallacy. *The Philosophical Quarterly*, *56*, 39–54.

Schnieder, B. (2010). A Puzzle about 'Because'. *Logique et Analyse*, *53*, 317–43.

Shoemaker, S. (1980). Causality and Properties. In P. van Inwagen (ed.), *Time and Cause* (pp. 109–35). Dordrecht: Reidel.

Sider, T. (2011). *Writing the Book of the World*. Oxford: Oxford University Press.

Sober, E. (1982). Why Logically Equivalent Predicates May Pick out Different Properties. *American Philosophical Quarterly*, *19*, 183–89.

Stalnaker, R. C. (1976). Possible Worlds. *Noûs*, *10*, 65–75.

Studd, J. P. (2015). Absolute Generality and Semantic Pessimism. In A. Torza (ed.), *Quantifiers, Quantifiers, and Quantifiers: Themes in Logic, Metaphysics, and Language* (pp. 339–66). Heidelberg: Springer.

Suszko, R. (1975). Abolition of the Fregean Axiom. In R. Parikh (ed.), *Logic Colloquium*, volume 453 of *Lecture Notes in Mathematics* (pp. 169–239). Heidelberg: Springer Verlag.

Swoyer, C. (1998). Complex Predicates and Logics for Properties and Relations. *Journal of Philosophical Logic*, *27*, 295–325.

Tahko, T. E. (2015). Natural Kind Essentialism Revisited. *Mind*, *124*, 795–822.

Taylor, B. (1993). On Natural Properties in Metaphysics. *Mind*, *102*, 81–100.

Thomasson, A. L. (2015). *Ontology Made Easy*. Oxford: Oxford University Press.

Tooley, M. (1977). The Nature of Laws. *Canadian Journal of Philosophy*, *7*, 667–98.

Uzquiano, G. (2015). Recombination and Paradox. *Philosophers' Imprint*, *15*(19).

Williams, J. R. G. (2007). Eligibility and Inscrutability. *The Philosophical Review*, *116*, 361–99.

Williams, J. R. G. (2010). Fundamental and Derivative Truths. *Mind*, *119*, 103–41.

Williams, J. R. G. (2012). Requirements on Reality. In F. Correia & B. Schnieder (eds.), *Metaphysical Grounding* (pp. 165–85). Cambridge: Cambridge University Press.

Williamson, T. (2003). Everything. *Philosophical Perspectives*, *17*, 415–65.

Witmer, D. G., Butchard, W., & Trogdon, K. (2005). Intrinsicality without Naturalness. *Philosophy and Phenomenological Research*, *70*, 326–50.

Wittgenstein, L. (1921/1961). *Tractatus Logico-Philosophicus*. London: Routledge & Kegan Paul.