

Chapter 2

Pleonastic Propositions

2.1 Introduction

The close of the preceding chapter left us realizing that if there is a successful version of the face-value theory, it will have to employ a notion of propositions which satisfies certain demands. Should I now rehearse those demands and then try to cook up an account of propositions to meet them? Not if I can arrive first at an account of propositions shown to be plausible on independent grounds, and then display its relevance to the face-value theory. That is what I'll try to do.

I call the account of propositions I'll offer an account of *pleonastic propositions*, because calling it that serves as a useful reminder of a crucial feature of the account, presently to be explained. But the label should not obscure what I'm up to. I aim to be revealing the nature of the propositions we believe, the actual referents of that-clauses, not inventing a new species of abstract entity especially assembled to do a certain job.

Pleonastic propositions are but one kind of pleonastic entity. In introducing the notion of a pleonastic entity, it will be helpful to start with a kind of pleonastic entity a little removed from more familiar arenas of dispute concerning the existence and nature of abstract entities—not that I'm so naive as to think this will give me an entry into the fray which escapes the axes of those protecting long-hardened positions. (Fortunately, inter-theoretical rivalries have cooled somewhat in recent times. Dugald Stewart, a nineteenth century Scottish philosopher, tells us that, after a dormant period, the dispute between Realism and Nominalism was revived in the fourteenth century by William of Occam:

From this time the dispute was carried on with great warmth in the universities of France, of Germany, and of England; more particularly on the two former countries, where the sovereigns were led, by some political views, to interest themselves deeply in the contest [between Realism and Nominalism], and even to

employ the civil power in supporting their favorite opinions. The Emperor Lewis of Bavaria, in return for the assistance which, in his disputes with the Pope, Occam had given to him by his writings, sided with the Nominalists. Lewis the Eleventh of France, on the other hand, attached himself to the Realists, and made their antagonists the objects of a cruel persecution.¹

We may be thankful current political leaders show less interest in metaphysics.) The pleonastic entities with which I'll begin, in the next section, are fictional entities, such as the fictional character Hamlet. Section 2.3 discusses the way in which properties are pleonastic entities, since that will have relevance to later discussions in the book, and in 2.4 I arrive at the way in which the propositions we believe are pleonastic entities. Section 2.5 discusses the way in which pleonastic propositions complete the face-value theory and considers a couple of challenges to the claim that they do give a successful completion of it. Section 2.6 clarifies my theory in relation to the notion of a possible world and to the modal notions of necessity and possibility; and 2.7 gives the direction of a more complete understanding of content, and thereby the further direction of this book.

2.2 Fictional Entities

James Joyce's novel *Ulysses* begins with the sentence

Stately, plump Buck Mulligan came from the stairhead, bearing a bowl of lather on which a mirror and a razor lay crossed.

This occurrence in the novel of the name 'Buck Mulligan' neither refers nor purports to refer to anything. Joyce was not trying to refer to a man named 'Buck Mulligan' and failing miserably; he was, in the way characteristic of fiction, *making as if* to refer to a man with that name and to tell us something about him, and we, in reading the novel, collude in this make believe when we read the novel as a novel. What is remarkable is that this *pretending use* of the name 'Buck Mulligan' should create the existence of something whose name *is* 'Buck Mulligan', thereby making it possible to use the name in a genuinely referential way in true statements about that referent. The thing brought into

¹ Stewart 1970, pp. 136-7.

existence is a certain abstract entity, the *fictional character* Buck Mulligan. Thus, from the fact that

(1) Joyce wrote a novel in which he used ‘Buck Mulligan’ in the pretending way characteristic of fiction

we may infer

(2) Joyce created the fictional character Buck Mulligan.

Since this thing, the fictional character Buck Mulligan, now exists in its own right, it will, like any existing thing, have numerous properties, and these may be ascribed to it in true statements. For example:

Buck Mulligan isn’t as well known as certain of Joyce’s other characters, such as Molly Bloom.

Buck Mulligan isn’t a fictional detective.

Buck Mulligan was based on someone Joyce knew.

In these statements, ‘Buck Mulligan’ occurs as a genuinely referential singular term whose referent is a certain fictional character, and each statement says something true or false about that character. We may call this use of fictional names the *hypostatizing use* of fictional names. I call valid inferences like that from (1) to (2) *something-from-nothing transformations* since they take one from a statement in which no reference is made to a thing of a certain kind (in this case, to a fictional entity) to a statement in which there is a reference to a thing of that kind. Pleonastic entities are entities whose existence is secured by something-from-nothing transformations (I call these things “pleonastic” entities because something-from-nothing transformations often take us to pleonastic equivalents of the statements from which they are inferred). Fictional entities, as revealed in the something-from-nothing transformation from (1) to (2), are pleonastic entities.

How is it possible to get something from nothing in this way, to infer from statements in which a fictional name is used in the pretending way to a true statement in which it’s used in the hypostatizing way? When Joyce wrote the displayed first sentence of his novel he was not referring to anything by ‘Buck Mulligan’, neither to a real person

nor to a fictional person. Yet, as a result of that use of the name, there now exists a certain thing whose existence supervenes on Joyce's using the name 'Buck Mulligan' in the pretending way. How is it possible for the existence of a fictional character to be entailed by, to supervene on, a fact that doesn't involve that fictional character and can thus be fully specified without reference to it?

Before trying to answer this and related questions, let's ask another one: How is it possible for us to have *knowledge of* fictional entities, abstract entities whose existence supervenes on the pretending use of words? We may put this question in proper perspective in the following way. Imagine a possible world β exactly like the actual world, α , except that no one in β has the concept of a fictional entity, and hence no one has knowledge of the existence of any fictional entities. They have the pretending use in α , and, by stipulation, all the fiction that exists in α exists in β , and therefore every fictional entity that exists in α also exists in β (for it belongs to our concept of a fictional entity that the existence of such an entity supervenes solely on the pretending use of its name). But while α is heavily populated with fictional entities, no one in β is aware of the existence of any fictional entity. What would it take to bring the people in β up to epistemological snuff with us? What would the people in β have to do in order to discover the existence of the fictional entities in their world?

The answer is easy: what they would have to do, and all that they would have to do, would be to play a certain language game—viz., to adopt out hypostatizing use of fictional names. But how can that be? How can adopting a certain linguistic, or conceptual, practice give one knowledge of things that exist independently of that practice? Because to have the practice is to have the *concept*, and *it's a conceptual truth*—a truth knowable *a priori* via command of the concept—that *the existence of fictional entities supervenes on the pretending use of their names*.

Maybe you feel like reading Kant to me. For Kant, in response to the ontological argument for the existence of God, famously held that "existence isn't a predicate," where by this he meant that no mere concept, however defined, can secure that there exist things that fall under the concept. Hartry Field, endorsing Kant's point, has succinctly restated it thus:

An investigation of conceptual linkages can reveal conditions that things must satisfy if they are to fall under our concepts; but it can't yield that there are things that satisfy those concepts (as Kant pointed out in his critique of the ontological argument for the existence of God).²

For consider the concept of a *wishdate*, which I hereby stipulatively introduce thus:

x is a *wishdate* =_{df} x is a person whose existence supervenes on someone's wishing for a date, every such wish bringing into existence a person to date

The point that Kant and Field are making implies that while this is a perfectly kosher definition, it doesn't result in its being true that there are any wishdates, no matter who wishes for a date. All that follows from the stipulative definition of a wishdate is that *if (per impossibile) wishdates exist, then their existence supervenes on the mere wish for a date.*

At the same time, isn't it *obvious* that, if there are fictional entities, then it's a conceptual truth that their existence supervenes on the pretending use of their names? I put the qualification "if there are fictional entities" not because the existence of fictional entities is in doubt but because there are philosophers who doubt it. By my lights, it's true that

The fictional spy James Bond is a lot more famous than the fictional detective Adam Dalgleish

and that in this displayed statement the names 'James Bond' and 'Adam Dalgleish' occur as genuinely referential singular terms whose referents are fictional characters.³ So, there must be an important difference between the concept of a fictional character and that of a wishdate such that by virtue of that difference it can be a conceptual truth that the existence of a fictional character supervenes on the pretending use of its name but it can't be any kind of truth that the existence of a wishdate supervenes on the wish for a date. It can hardly be any kind of truth, let alone a conceptual truth, that wishing for a date brings into existence a person to date. That sort of wishing does not have that power, and if it

² Field 1989, p. 5.

did, it would be a contingent causal power; the wishing would not *entail* that there springs into existence someone to date. Yet it seems clear that it is a conceptual truth that using the name ‘*n*’ in writing a fiction creates the fictional character *n*. How should we characterize the crucial difference by virtue of which it can be a conceptual truth that fictional entities supervene on the pretending use of fictional names, but the existence of a date does not supervene on the wish for a date? Perhaps in the following way.

There is a crucial difference between the concept of a wishdate and the concept of a fictional entity (or of any other kind of pleonastic entity, although for now my focus will be just on fictional entities). To a first approximation, the difference is that:

There are numerous theories *T* such that when we add to *T* the concept of a wishdate together with the claim that wishing for a date entails the existence of a wishdate, the resulting theory is *not* a *conservative extension* of *T*. But if we add to *any* theory *T* the concept of a fictional entity together with its attendant claim that using a name in the pretending way entails the existence of a fictional entity, the resulting theory *is* a conservative extension of *T*.

A theory *T'* is a *conservative extension* of a theory *T* provided that *T'* includes *T* and nothing storable in the vocabulary of *T* is logically entailed by *T'* but not by *T*. (The notion of “adding a concept” is vague; it should be taken to mean that the canonical expression for the new concept has its full meaning when introduced, and that as much of that meaning as can be made explicit via defining conditions is made explicit, especially whatever existence-entailing something-from-nothing conditionals are partly definitive of the concept.)

Suppose *T* is a fairly rich true physical theory that doesn’t employ the concept of a wishdate or the concept of a fictional entity, but does assert that someone wished for a date and that James Joyce wrote *Ulysses*, where this is understood to include the text of the novel. Adding to *T* the claim that wishing for a date entails the existence of a wishdate will entail that there exists a person—the person brought into existence by the

³ Saul Kripke gave convincing arguments for the existence of fictional entities in his unpublished John Locke Lectures, given at Oxford University in 1973. See also John Searle 1979, chapter 3.

wish for a date—whose existence wasn't recognized in T but was storable in the vocabulary of T . The resulting theory is therefore clearly not a conservative extension of T . Indeed, should the new person exist, he or she would be a substantial physical object that would enormously disturb the preexisting causal order. But when we add the concept of a fictional entity to T , and with it its something-from-nothing entailment claims, what can we get that we couldn't get from T alone? Well, we can get the statement that the fictional character Buck Mulligan was created by James Joyce, but (subject to a small qualification) we can assert nothing new that can be said in the language of T . Adding fictional entities to one's ontology via legitimate something-from-nothing entailments does nothing to disturb the preexisting causal order. That is why adding those entailment claims to T yields a conservative extension of T .

The first shot needs refinement; as it stands it's too strong to account for what makes the truths expressed by instances of (roughly) 'If someone uses the name " n " in the pretending way, then there is a fictional entity n ' conceptual truths; and it's also arguably too weak, that it lets in things that shouldn't be let in.

The first shot is too strong because there are ways of adding the concept of a fictional entity together with its something-from-nothing entailment claims and getting a theory that isn't a conservative extension of the original theory. For example:

- The new theory, but not the original theory, may entail that more than such-and-such many things exist, which was storable in the original theory.
- In addition to asserting that Joyce wrote *Ulysses*, the original theory may also assert that if there are abstract entities, then it will snow in Miami in August. When we add the concept of a fictional entity to that theory, along with its something-from-nothing entailment claim, we get a theory which entails that fictional entities—and therefore abstract entities—exist, in which case the resulting theory will also entail that it will snow in Miami in August, which was assertible but unasserted in the original theory.
- In addition to asserting that Joyce wrote *Ulysses*, the original theory may assert that there are no abstract entities, so that when

we add the concept of a fictional entity and its something-from-nothing entailment claim, we get an inconsistent theory from which everything follows.

That the first shot may also be too weak is implied by a type of example independently suggested to me by Kit Fine and by Josh Schechter. For consider the concept of an anti-fictional entity, where that is stipulated to be the concept of an abstract object whose existence both supervenes on anything one likes and rules out the existence of fictional entities. It's arguable that the concept of an anti-fictional entity is on a par with the concept of a fictional entity as regards conservative extension, so that adding the concept of an anti-fictional entity to a theory will conservatively extend it to whatever extent adding the concept of a fictional entity would, thereby frustrating hopes of accounting for pleonastic concepts in terms of conservative extension (since if the concepts of a fictional entity and of an anti-fictional entity were both pleonastic concepts, then that would imply that fictional entities did and did not exist).⁴

These problematic examples motivate a strategy employed by Hartry Field when he wanted to say that the result of adding a mathematical theory to a nominalistic theory conservatively extends the nominalistic theory but had to contend with a nominalistic theory that says things that rule out the existence of abstract entities.⁵ Field's solution to this problem was, roughly speaking, to say that a mathematical theory can't be assured of conservatively extending every nominalistic theory, but can be assured of conservatively extending the result of restricting the quantifiers of any given nominalistic theory to things that aren't mathematical entities. (We restrict T 's quantifiers to things that satisfy a predicate ' Px ' thus: if T asserts ' $\forall xGx$ ', then we replace that with ' $\forall x(Px \rightarrow Gx)$ ', and if T asserts ' $\exists xGx$ ', then we replace that with ' $\exists x(Px \& Gx)$ '.) I believe that a suitable deployment of Field's strategy will enable me to put together a definition that will capture what I want from the notion of a pleonastic entity.⁶ The definition should be

⁴ The parasitic nature of the concept of an anti-fictional entity raises a problem familiar from other problems created by parasitic concepts. So you think that hypothesis H is true because it provides the best explanation of all observable evidence. But how can you justifiably reach that conclusion when you have done nothing to rule out another hypothesis which explains the same evidence—viz., the hypothesis Anti- H , which says that H is false but that all the observable evidence is as it would be if H were true?

⁵ Field 1980, pp. 11-12.

⁶ The idea of accommodating certain difficulties for a conservative extension test by restricting a theory's quantifiers to things in its recognized ontology is also used by Bob Hale and Crispin Wright in their

regarded as stipulative, substantive claims coming in the form of claims about what kinds of things are pleonastic entities. For now, the definition needs only to get right the distinction between wishdates and fictional entities. I start with some preliminary definitions.

Where ‘ \Rightarrow ’ expresses metaphysical entailment,⁷ ‘ $S \Rightarrow \exists xFx$ ’ is a *something-from-nothing F-entailment claim* iff (i) its antecedent is metaphysically possible but doesn’t *logically* entail either its consequent or any statement of the form ‘ $\exists x(x = \alpha)$ ’, where ‘ α ’ refers to an F , and (ii) the concept of an F is such that if there are F s, then $S \Rightarrow \exists xFx$. (I’ll say that the concept of an F “implies” a something-from-nothing F -entailment claim if it satisfies (ii).)

A *pleonastic entity* is an entity that falls under a *pleonastic concept*; and a *pleonastic concept* is the concept of an F which implies *true* something-from-nothing F -entailment claims.

For any theory or sentence T , $T^{\sim F}$ is the theory or sentence that results from restricting each quantifier of T to things that aren’t F .

I now offer the following conservative-extension criterion for being a pleonastic concept.

(CE) The concept of an F implies true something-from-nothing F -entailment claims—and is therefore a *pleonastic concept*—iff (i) it implies something-from-nothing F -entailment claims, and (ii) for any theory T and sentence S

discussion of conservativeness in connection with Hume’s Principle. Their way of putting the qualification is that the extending theory “must not introduce fresh commitments which (i) are expressible in the language as it was prior to the introduction of its [additional material] and which (ii) concern the previously recognized ontology of concepts, objects, and functions, etc., whatever in detail they may be” (Hale and Wright 2000, p. 302).

⁷ A *metaphysically entails* B just in case the material conditional $A \rightarrow B$ is metaphysically necessary. As I understand the notion, metaphysical necessity is that strong form of necessity such that whatever is logically, arithmetically, or conceptually necessary is *ipso facto* metaphysically necessary, but something can be metaphysically necessary without being necessary in any of those ways (such as, perhaps, the proposition that water is composed of H_2O molecules). Something can be “physically necessary”—e.g., the proposition that nothing travels faster than the speed of light—without being metaphysically necessary. I don’t think metaphysical necessity can be defined. This is touched on again in 2.6.

expressible in T ,⁸ if the theory obtained by adding to T^F the concept of an F , together with its something-from-nothing F -entailment claims, logically entails S^F , then T^F logically entails S^F .⁹

In other words, adding pleonastic entities to any theory conservatively extends that theory, relative to the restriction on quantification. We can see how CE handles the problems for the first shot in the following way (in considering each example, ' T ' will designate the original theory).

The first problem was that the theory obtained when we add to T the concept of a fictional entity may entail that more than such-and-such many things exist, when this isn't entailed by T . But since the new theory doesn't entail that more than such-and-such many *non-fictional* things exist, it's a conservative extension of T^F (where ' Fx ' abbreviates ' x is a fictional entity'), whose claims about how many things exist are limited to claims about how many *non-fictional* things exist.

The second problem for the first shot was that T may assert that if there are abstract entities, then it will snow in Miami in August. But T^F merely says that if any non-fictional entities are abstract entities, then it will snow in Miami in August, and thus nothing storable but unstated in T^F is forthcoming when we add the concept of a fictional entity to T^F .

The third problem was that T may assert that there are no abstract entities. But since T^F merely asserts that nothing exists that is both a non-fictional entity and an abstract entity, there is again no problem for CE.

Nor is CE threatened by Fine and Schechter's anti-fictional entities. If there is a threat here, it's either (a) that the concept of an anti-fictional entity would, unacceptably, be a pleonastic concept if CE were correct, or (b) that the result of adding fictional entities to the required restriction of a theory that asserts the existence of anti-fictional entities would not conservatively extend that theory. It's clear (b) is not a threat.

Suppose $T = (S \ \& \ \exists xAx)$, where ' Ax ' abbreviates ' x is an anti-fictional entity'. Now, ' $\exists xAx$ ' for all intents and purposes is equivalent to

⁸ S 's being expressible in T doesn't imply its being asserted in T ; it merely means that S can be formulated in the vocabulary of T .

$$\exists x(x = x \ \& \ \forall y(\sim Fy))$$

where ‘ $\sim Fy$ ’ abbreviates ‘ y is not a fictional entity’. Thus, when we restrict T with the formula ‘ $\sim Fy_i$ ’ (for appropriate variable ‘ y_i ’), $T^{\sim F} =$

$$S^{\sim F} \ \& \ \exists x((\sim Fx \ \& \ x = x) \ \& \ \forall y(\sim Fy \rightarrow \sim Fy))$$

which is conservatively extended by the theory that adds fictional entities. Nor is (a) a threat, since the result of adding anti-fictional entities to the required restriction of a theory that asserts the existence of fictional entities is an inconsistent theory from which everything follows. For let $T = (S \ \& \ \exists xFx)$, where ‘ Fx ’ abbreviates ‘ x is a fictional entity’. When we restrict each quantifier in T with the formula ‘ $\sim Ax_i$ ’, $T^{\sim A} =$

$$S^{\sim A} \ \& \ \exists x(\sim Ax \ \& \ Fx)$$

so that when we add to this the claim that anti-fictional entities exist, we get the inconsistent

$$S^{\sim A} \ \& \ \exists x((\sim Ax \ \& \ Fx) \ \& \ \exists y(y = y \ \& \ \forall z(\sim Fz)))$$

It might be thought that CE still faces problems. An example of Hartry Field’s (Cian Dorr raised a similar problem) might again suggest that CE is too weak, that it lets in things I don’t want to let in. Field’s example concerns the concept of an undetectable non-interfering god. Won’t that conservatively extend any theory that the concept of a fictional entity will conservatively extend? I don’t think so. What are we to make of the concept of a “god”? If the concept of an undetectable and non-interfering god is simply the concept of an undetectable and non-interfering thing, then there is no problem, since that is what fictional entities are. If the concept of a god is to be non-vacuous, it must entail that gods have propositional attitudes. But now consider a theory that entails the existence of propositional attitudes. When we add to that theory the claim that an undetectable and non-interfering god exists, we get a theory that introduces new propositional attitudes and therefore isn’t a conservative extension of the original theory.¹⁰

⁹ CE supercedes the criterion proposed in Schiffer (2001).

¹⁰ Kit Fine has wondered about changing Field’s non-interfering god to a prime mover, or first cause. But I don’t think that adding a first cause can be assured of conservatively extending all relevant theories. A theory that doesn’t say anything was caused but in which causal claims are expressible won’t be conservatively extended by the addition of the assertion that there’s a first cause, since the resulting theory will entail that some event was caused.

CE, as I said, should be read as a stipulative definition of ‘pleonastic concept’, with substantive claims coming in the form of claims that such-and-such concept is, or isn’t, a pleonastic concept. The concept of a fictional entity is a pleonastic concept, fictional entities therefore pleonastic entities. Fictional entities are that thin and inconsequential. We over-generalize Kant’s insight when we lump concepts that pass the CE test with the concept of a wishdate and the concept of God. Fictional entities are mere shadows of the pretending use of their names;¹¹ they come softly into existence, without disturbing the preexisting causal order in any way. That is why claims that they exist may be conservatively added to the truths we had before those claims were added. There is nothing more to the nature of fictional entities than is determined by the hypostatizing language game that recognizes them in our ontology. This is in contrast with cats, islands, electrons, and whatever else enjoys the highest degree of ontological and conceptual independence from our linguistic or conceptual practices. Their essential natures must be discovered by scientific investigation, not by an armchair perusal of our talk about them. But there can be nothing more to the nature of fictional entities than is determined by our hypostatizing use of fictional names. The “science” of them may be done in an armchair by reflective participants in the hypostatizing practice; fictional entities can have “no hidden and substantial nature for a theory to uncover. All we know and all we need to know about [them] in general”¹² is determined by our hypostatizing use of fictional names. As we shall see, this sort of thinness is an important feature of the pleonastic entities that most interest us.

Let’s briefly take stock. Sherlock Holmes is a much more famous fictional detective than Adam Dalgleish. The statement I just made is literally true, and its literal truth requires the existence of the fictional characters Sherlock Holmes and Adam Dalgleish. Given the existence of these fictional characters, it ought to be clear that even slightly reflective people have *a priori* knowledge of the conceptual truth that the existence of any fictional entity supervenes on the existence of a work of fiction in which the character’s name is used in the pretending way characteristic of fiction. This explains how we are well enough *en rapport* with certain fictional entities to have knowledge

¹¹ The “shadows” metaphor is borrowed from David Armstrong’s (1989) metaphor of properties as shadows of predicates.

about them. It is, however, one thing for us, as philosophers, to know that it's a conceptual truth that the existence of fictional entities supervenes on the pretending use of their names and another, more difficult thing to discern and state clearly what it is about the concept of a fictional entity that allows the relevant something-from-nothing entailment claims to be conceptual truths. I'm more confident that these are conceptual truths than I am that the stuff about conservative extension captures what's essential to that, and I'm more confident that the general idea about conservative extension is on the right track than I am about whether CE has everything exactly right. Anyway, I hope I've succeeded in getting at what is crucial with CE, but, philosophy and the effort to state necessary and sufficient conditions being what they are, I wouldn't be completely shocked to learn that what I've written needs further qualification. Even approximate truth may do pretty well for my further purposes, and at this point it should be plain what those purposes are. My strategy has been to begin with things—fictional entities—subject to something-from-nothing entailments that are fairly undeniably conceptual truths; next, to discern what it is about the concept of a fictional entity that permits its something-from-nothing entailment claims to be conceptual truths; next, to state the result of that discernment as a general principle; and eventually to show that the propositions we believe, along with other things of interest, satisfy that principle. I turn now, in the next two sections, to other things, closer to what matters in this book, which, I claim, also fall under concepts that satisfy CE.

2.3 Pleonastic Properties

Pleonastic entities are entities whose existence is typically secured by something-from-nothing transformations—“secured” not necessarily in the sense that they are brought into existence (like fictional entities) but in the sense that their existence supervenes on the premises of something-from-nothing transformations. We have a something-from-nothing transformation when from a statement involving no reference to an *F* we can deduce a statement that does refer to an *F*. The property of being a dog is a pleonastic entity. From the statement

(3) Lassie is a dog

¹² Johnston 1988, p. 38.

whose only singular term is 'Lassie', we can validly infer the pleonastic equivalent

(4) Lassie has the property of being a dog

which contains the new singular term 'the property of being dog', whose referent is the property of being a dog.

Just as we asked about the license for the something-from-nothing transformations that yield fictional entities, so we may repeat the question for properties: What explains the validity of the something-from-nothing transformation that takes us from (3) to (4)? How *can* (3) entail (4), thereby entailing that the property of being a dog exists? The answer, of course, mirrors that given for fictional entities; it's the answer that applies to all pleonastic entities. It's a *conceptual truth* that if Lassie is a dog, then Lassie has the property of being dog. The intuitive rightness of this can be glossed in the same way we glossed the intuitive rightness of its being a conceptual truth that the existence of a fictional entity supervenes on the pretending use of a fictional name. For how, one might ask, are we able to have knowledge about properties, mind- and language-independent abstract entities that are wholly incapable of causally interacting with us? This question, as before, is made more vivid in the following way. Suppose there is a possible world exactly like ours except that our counterparts in that world don't have any property-hypostatizing linguistic or conceptual practices, and hence have no concept of a property. These people can think that Lassie is a dog, but they can't infer from this that Lassie has the property of being a dog, even though in that world, as in every world in which Lassie exists, it's necessarily true that if Lassie is a dog, then Lassie has the property of being a dog. Lacking the concept of a property, these people are entirely ignorant of properties, even though they live in a world as rich in properties as the actual world. What would it take to bring these people up to epistemological snuff with us?

What it would take, and all that it would take, would be for them to engage in a certain manner of speaking, a certain language game—namely, our property-hypostatizing practices, in particular our property-yielding something-from-nothing transformations. But how can this be? They certainly couldn't discover the existence of *volcanoes* by engaging in any language game. How can merely engaging in a linguistic, or conceptual, practice give one knowledge of things that exist independently of that practice? Because to engage in the practice is to have the concept of a property, and to

have the concept of a property is to know *a priori* the conceptual truths that devolve from that concept, such as the conceptual truth that every dog has the property of being a dog.

We should expect these conceptual truths to be explained by the fact that the concept of a property is a pleonastic concept, properties therefore pleonastic entities. And it's clear that the concept of a property passes the CE test. For take any theory that doesn't employ the notion of a property. It says that Fido is a dog and that dogs bark, but it doesn't say that Fido has the property of being a dog or that things that have the property of being a dog also have the property of being barkers. When we add to this theory the concept of a property, which carries with it its something-from-nothing property-entailment claims, the theory that results is (modulo the point about restriction of quantifiers) a conservative extension of the original theory. Once again, properties are that thin and that inconsequential; properties, to borrow again Armstrong's metaphor, are mere shadows of predicates; adding them to a theory does nothing to disturb the output of that theory, does nothing to alter that theory's take on the preexisting causal order.

Thus, as with all pleonastic entities, properties have "no hidden and substantial nature for a theory to uncover."¹³ The essential truths about them are directly or indirectly determined by the hypostatizing practices constitutive of the concept of a property, together with those necessary *a priori* truths applicable to things of any kind, such as that if $x = y$, then whatever property x has, y has, and vice versa. As regards the principles by which properties are individuated, it means that if a question of individuation is left unsettled by the practices constitutive of the concept of a property, then that question has no determinate answer. The same applies to any other kind of pleonastic entity. This has important implications for the mind/body problem. Before illustrating this with respect to properties, let me first illustrate it with respect to pleonastic entities of another kind, events.

Events are pleonastic entities, for they, too, enter our ontology via something-from-nothing transformations. From

Jane was born on a Tuesday,
whose only singular term is 'Jane', we may validly infer the pleonastic equivalent,
Jane's birth was on a Tuesday,

¹³ Johnston, *op. cit.*

which contains the new singular term ‘Jane’s birth’, whose referent is Jane’s birth, an event. Leibniz’s law gives us a means for establishing numerous non-identities. If Jane’s birth occurred in 1850 and her death occurred in 1933, then Jane’s birth \neq Jane’s death. But there is a paucity of principles determined by our conceptual practices that enable us to establish interesting identifications; evidently, such principles aren’t required by the purposes for which we have the concept of an event.¹⁴ For example, consider Donald Davidson’s claim that while no mental event *type* is identical to any physical or functional event type, every mental event *token* is identical to a physical event token; that is to say, every token of a mental event type is also a token of a physical event type.¹⁵ But there is nothing in the principles by which events are individuated that could make any identification of a mental event token with a physical event token determinately true, given that no mental event type is identical to any physical or functional event type. It’s given, let’s suppose, that at time *t* there occurs in you a certain twinge and an *X*-fiber firing. How might you determine whether there is one event token in you at *t* that is both the twinge and the *X*-fiber firing, as opposed to one event that is the twinge and a distinct co-temporal event that is the *X*-fiber firing? It should be clear that there is no procedure you could in principle carry out to resolve the question. This isn’t because the answer can’t be known, or can be known only on the basis of some abstruse methodological principle of simplicity or the unity of science, but because there is nothing in our concept of an event to determine the issue. It is, at best, indeterminate.

The same is true for properties and the claim that mental properties are identical to physical or functional properties, where the correct identity claim is one that is supposed to be knowable only *a posteriori*, if knowable at all. For what procedure might you follow to determine whether, say, the property of being a pain = the property of being a *c*-fiber stimulation? We can, to simplify more than a little, know *a posteriori* that water = H₂O because our concept of water is the concept of whatever it is that we drink, swim

¹⁴ Here (and elsewhere in this book) I flout the Quinean principle that in order for there to be *F*s there must be a non-trivial criterion of identity for *F*s. After all, there are pains, but where is the non-trivial criterion of identity for pains? As implied in the text, very often the lack of a non-trivial criterion for *F*s will render the proposition that *x* is the same *F* as *y* indeterminate, but not always, as is obvious from the fact that I can know, say, that the Φ is the same *F* as the Ψ without a non-trivial criterion just by virtue of knowing of some given *F* that it’s both uniquely Φ and uniquely Ψ .

¹⁵ Davidson 1980b.

in, etc., and we then discover that that stuff is composed of H₂O molecules. But we don't think of the property of being a pain under a concept that identifies it in terms of one of its contingent properties, and there is nothing in the practices constitutive of our concept of a property, or of the property of being a pain, that leaves a conceptual space for it to be determinately true that the property of being a pain is identical to any physical property.

It might be helpful to think of a contrasting view. Consider the property of being a dog. On my view, there isn't a lot more to this property than can be culled from the something-from-nothing transformation that allows us to move back and forth between

(5) x is a dog

and its pleonastic equivalent

(6) x has the property of being a dog.

Now, suppose (what is probably not the case), that, necessarily, a thing is a dog just in case it belongs to such-and-such genotype. Can we then say that the property of being a dog = the property of belonging to such-and-such genotype? I don't see that we can, for there is nothing in the something-from-nothing practice or in any ancillary practice to establish that either

The property of being a dog = the property of belonging to such-and-such genotype

or

The property of being a dog and the property of belonging to such-and-such genotype are distinct metaphysically equivalent properties.

On the contrasting view, 'doghood' is simply a proper name of a thing enjoying the same conceptual and ontological independence from our hypostatizing practices as a volcano. Perhaps the link between 'doghood' and doghood by virtue of which the former names the latter is a causal link. In any case, once we grasp the bearer of 'doghood', it's then a matter for some sort of further, *a posteriori* investigation to determine the nature of, and hence the individuating conditions for, the property that bears that name. For all we yet know, we may discover that this property is identical to this, that, or the other thing.

Consider the difference between volcanoes and fictional characters. We can first discover volcanoes and then explain the introduction of 'volcano' in terms of the

discovered volcanoes, but we can't first discover fictional characters and then explain the introduction of 'fictional character' in terms of the discovered fictional characters. This is because of the way, already made a big deal of, fictional entities enter our ontology only *qua* fictional entities, via the something-from-nothing transformations that give us our knowledge of them. The problem with the contrasting view is that it makes properties out to be like volcanoes when they're really like fictional entities. There can be no intrinsically identifying the property of being a dog apart from identifying it as the property of being a dog. If something falls under the concept of doghood, we can't look at what it is apart from its falling under that concept, for there is nothing there apart from its falling under that concept. There can be nothing more to the identity and individuation of a particular property than is determined by the canonical concept for that property. This is why the property of being a dog has no hidden nature for empirical investigation to unearth; it's a shadow of the word 'dog'.

I turn now to certain critical questions whose answers will further clarify the notion of a pleonastic entity.

The way properties enter our conceptual scheme via something-from-nothing transformations and the way they take to the metaphor of being shadows of predicates may suggest that in some sense they are creations of our hypostatizing linguistic or conceptual practices. I believe there is something to this. Conceptualism about properties is the view that properties are creations of our conceptual or linguistic practices. I don't see how this can be literally true, since properties exist in every possible world, and thus in possible worlds in which there are neither thinkers nor speakers. Since it's a metaphysically necessary truth that properties exist, they can hardly be created by anything we do. Indeed, their being independent of us in this way is a consequence of the pleonastic conception of them. For example, from

Necessarily, there are dogs or there are not dogs

we may infer

Necessarily, there are things that have the property of being a dog

or there are not things that have the property of being a dog

which entails

Necessarily, the property of being a dog is or is not instantiated

which entails

Necessarily, the property of being a dog exists.

The property of being a dog, in other words, exists in every metaphysically possible world (although, to be sure, it's only in some of these possible worlds that it's instantiated).

So, conceptualism about properties is false: nothing we do creates properties. At the same time, I believe that the theory of pleonastic properties is a conceptualist *manqué* theory: it's a theory that should relieve the need to be a conceptualist, a theory the conceptualist would have accepted had she thought of it. In this sense, the doctrine of pleonastic properties, and of pleonastic entities generally, is a deflationary theory, but that quasi-evaluative term has in recent years been used in so many diverse ways that it is now almost entirely devoid of descriptive content.

There are three questions about something-from-nothing transformations that must be addressed. The questions arise for any kind of pleonastic entity, but I'll address them, for the most part, with respect to properties, although what I say is intended to apply, *mutatis mutandis*, to other kinds of pleonastic entities.

The first question is about the validity of the inference schema

(7) x is F

So, x has the property of being F

For virtually all substitution instances of ' F ', e. g.,

Fido is a dog

So, Fido has the property of being a dog

the schema (7) yields a valid inference; but, notoriously, we can't safely conclude that (7) is valid for all substitution instances, as is shown by the at first sight seemingly innocuous inference

(8) Doghood is a property that doesn't instantiate itself.

So, doghood has the property of being a property that doesn't instantiate itself.¹⁶

¹⁶ Another exception to the unqualified validity of the schema (7) is provided by a sentence such as 'Cell phones are ubiquitous,' from whose truth we shouldn't want to infer that every cell phone has the property of being ubiquitous.

For what are we now to say about *the property of being a property that doesn't instantiate itself*? Does it instantiate itself? If it does, it doesn't; if it doesn't, it does. We are landed with a contradiction. Let me cautiously say that it can't be said that (8) is determinately valid. My reason for putting it in this coy way will be apparent in chapter 5, when I present my own views on the 'determinately' operator. For now, however, it doesn't really matter whether we say either that there can be no such property as the property of being a property that doesn't instantiate itself or that it's indeterminate whether there can be such a property. The important thing for now is that any claims about the validity of the schema (7) must be qualified, and the question of moment, for me, is how to qualify (7) without undermining the pleonastic conception of properties.

Before directly addressing this question I want to say something about the nature of philosophical paradoxes and the bearing of that on the nature of our quotidian concepts. A philosophical paradox is a set of apparently mutually incompatible propositions each one of which enjoys some significant degree of plausibility when viewed on its own. The problem of free will is a paradigm of such a paradox. Here there are three evidently mutually incompatible propositions: (i) that we have free will, i. e., that we sometimes act freely; (ii) that everything we do is such that we were caused to do it by events that occurred even before we were born and, therefore, over which we had no control; and (iii) that the foregoing two propositions can't both be true. Now, a *happy-face solution* to a philosophical paradox would do two things. First, it would identify the odd-guy-out: it would tell us that the apparently incompatible propositions were not really incompatible, or that a certain one of the plausible propositions wasn't the truth it appeared to be; and second, it would, in unmasking the odd-guy-out, do so in a way that removed its patina of plausibility, so that we would never again be taken in by it. Most of the traditional solutions to the classical philosophical paradoxes aspire to be happy-face solutions. For example, the problem of free will admits in principle of three potential happy-face solutions, each with its school of proponents: "compatibilism," which accepts (i) and (ii) but denies (iii); "hard determinism," which accepts (ii) and (iii) but denies (i); and "libertarianism," which accepts (i) and (iii) but denies (ii).

A nice example of a genuinely happy-face solution is the solution to the "paradox" of the barber who shaves all and only those who don't shave themselves.

Does he, or doesn't he, shave himself? Here the solution is that it's logically impossible for there to be such a barber, and once this is appreciated there is nothing whatever paradoxical or puzzling about the barber. But what classical philosophical problem enjoys a happy-face solution? The trouble in each case is that while any given "solution" offers an identification of the odd-guy-out, none succeeds in removing from it its patina of plausibility. This is illustrated by each position in the space of potential happy-face solutions to the problem of free will. Consider compatibilism, no doubt the most widely held solution. If to say that an act is free is just to say that the actor would have acted otherwise if she had decided to (or, anyway, something in that direction), then how did the problem ever come to be a problem, and why has no one succeeded in eliminating the worry that we can't be free in any sense that matters if everything we do is caused by factors over which we had no control? When one thinks of the classical philosophical paradoxes—skepticism about the external world, the problem of induction, the sorites, the liar, etc.—it seems clear that none of them truly enjoys a happy-face solution. There are competing solutions to each paradox, and none succeeds in getting us to see through the alleged spurious plausibility of its chosen odd-guy-out.

It's my view that most, if not all, of these paradoxes don't have happy-face solutions; they have *unhappy-face* solutions.¹⁷ An unhappy-face solution to a paradox would do two things. First, it would tell us that there can be no determinately correct complete identification of the odd-guy(s)-out; and second, it would tell us what it is about the concepts involved that explains this. In each case, the explanation will find a glitch in the concept or concepts involved, a tension in the underived conceptual roles that individuate those concepts. For example, as regards the problem of free will, our concept of free will has one component that inclines us to ascribe free will in certain conditions ("paradigm cases" of free acts, if there are any), whereas another component disinclines us to ascribe free will when we know that the agent's decision was caused by factors over which he or she had no control, and there is nothing else in the concept or elsewhere to adjudicate the issue. The troublesomeness of the paradox-inducing conceptual glitch may vary from case to case. In some cases it may be disastrous, as it arguably is for our concepts of free will and moral responsibility, but in other cases it may be benign. One

¹⁷ Schiffer 1995/6.

way in which it may not be disastrous is nicely illustrated by an example made up by Charles Chihara.¹⁸ A group of secretaries of clubs they are not allowed to join themselves form a club whose condition of membership is that one be a secretary of a club one isn't allowed to join. All works well until one day the club itself hires a secretary. Is this person eligible to join the club? If yes, then the condition of membership is violated, and so, no. But if no, then yes, for then the condition of membership is satisfied. Chihara suggests that something like this lies behind the semantic paradoxes, and I'm inclined to agree.

I would put the point this way as regards the version of Grelling's paradox represented by "the property of being a property that doesn't instantiate itself." Our concept of a property is governed by an underived conceptual role that disposes us to accept instances of (7); it's *as if* the notion of a property had been explicitly introduced without qualification by the schema. All goes well until we hit upon the paradoxical instance (8), at which point it's not clear what to say, and for which there may well be no determinately correct thing to say (here I look ahead to chapter 5). What is crucial to the present discussion is that we must not suppose that the paradox generated by "the property of being a property that doesn't instantiate itself" has a happy-face solution. Our concept of a property *is* largely determined by our disposition to accept instances of the property-hypostatizing something-from-nothing transformation (7) (that's why the paradox is a paradox); it's just that the concept of a property, like the concepts of truth, knowledge, free will, and more, has a glitch, one revealed in the present version of Grelling's paradox. The existence of the paradox neither shows that the uninfected instances of (7) are invalid, nor that our concept of a property isn't based on the underived conceptual role that leads us to take uninfected instances of something-from-nothing property-entailment claims as conceptual truths. Where we go from here is another matter, but I will defer further discussion of what can be done when confronted with an unhappy-face solution to chapter 5.

A second point of clarification concerns the scope of something-from-nothing entailments. It's important to appreciate that it doesn't follow from the fact that '*x* has the property of being *F*' is a pleonastic equivalent of '*x* is *F*' that *every* statement

¹⁸ Chihara 1979.

containing a singular term whose referent is a property is pleonastically (or in *any* way) equivalent to a statement containing no such singular term. The statement that so-and-so is humble is pleonastically equivalent to the statement that so-and-so has humility, the property of being humble, but I doubt that the statement that humility is a virtue is pleonastically equivalent to any statement not containing a singular term referring to humility.¹⁹ When we get to propositions in the next section, it will be especially important to appreciate that what is stated by a belief report of the form ‘*A* believes that *S*’ can’t be stated without reference to the proposition to which the belief report’s that-clause refers.

Finally, what about inexpressible properties? A property isn’t inexpressible simply by virtue of not having a single *word* to express it: from the fact that *x* is a figure having fourteen sides and a fuzzy texture we may infer that *x* has the property of being a figure that has fourteen sides and a fuzzy texture. Still, there might be properties that are ineffable, or perhaps effable only in the languages of vastly more intelligent creatures, or of creatures with much different sensory abilities. My view here is that we can make sense of such properties just by virtue of our ability to make sense of there being a language—an enrichment of our own language or a completely different language—in which such properties are expressible.

2.4 Pleonastic Propositions

Propositions, the things to which that-clauses refer, are also pleonastic entities. They have their something-from-nothing transformations, such as the one that takes us from

Lassie is a dog,

whose only singular term continues to be ‘Lassie’, to another of its pleonastic equivalents,

That Lassie is a dog is true

(more colloquially, ‘It’s true that Lassie is a dog’), which contains the singular term ‘that Lassie is a dog’, whose referent is the proposition that Lassie is a dog. Consequently, the concept of a proposition—that is to say, of a that-clause referent—satisfies CE, the

¹⁹ Cf. Armstrong 1989.