

Copredication and Property Inheritance

David Liebesman and Ofra Magidor

1. Copredication and its Lessons

Books are physical things found on shelves. Dave's copy of *Naming and Necessity* was produced on a particular day, has a particular weight, and has (embarrassing) marginalia from his undergraduate days. Ofra's copy differs in each of these ways, and therefore Ofra has a different book on her shelf. Books are also informational things that are not tethered to shelves. Given that we have both studied *Naming and Necessity*, we have studied one and the same book.

A natural way to account for this is to take “book” to be ambiguous between a physical and an informational sense. This would explain why (1) has both a true reading and a false reading when there are three copies of *Naming and Necessity* on an otherwise empty shelf.

(1) Three books are on the shelf.

An ambiguity view faces an immediate challenge: certain sentences seem to make use of both the physical and informational senses of “book”, despite containing only a single occurrence of the word. Consider (2):

(2) Three red books are informative.

When we use “red” to characterize books, we tend to focus on physical books. When we use “informative”, we tend to focus on informational books.¹ If we assume that physical books can't be informative, and informational books can't be red, then it seems we're forced to disambiguate the single occurrence of “book” in (2) in two different ways. It is hard to see how this can be done.

Nouns like “book” and sentences like (2) are common. The following are all examples of copredication, where—intuitively—we seem to use a single noun to designate two different types of entities (soon we will reject this intuitive characterization):

¹ Throughout the paper we will use the term “informational book” for the kind of book that generates the false reading of (1). However, the reader should be careful to avoid one misleading implication of this terminology – namely, that informational books are individuated merely by the information, or list of words, that they contain. We take it that informational books can be individuated by more than just the information they contain: on the one hand, two versions of *War and Peace* differing by one or two words might still count as the same informational book, and on the other hand, if casually isolated aliens happen to produce a book that's word-to-word identical to *War and Peace* we might still want to say that they have produced a different informational book.

- (3) Lunch was delicious, but lasted hours.
- (4) The bank was vandalized after calling in Bob's debt.
- (5) Fred mastered three heavy books.

The lessons drawn from co-predication are dramatic. Chomsky, for instance, argues on the basis of copredication that we should abandon referential semantics altogether:²

Contemporary philosophy of language...asks to what a word refers, giving various answers. But the question has no clear meaning. The example of "book" is typical. It makes little sense to ask to what *thing* the expression "Tolstoy's *War and Peace*" refers, when Peter and John take identical copies out of the library... In general, a word, even of the simplest kind, does not pick out an entity of the world, or of our "belief space". Conventional assumptions about these matters seem to me very dubious. (Chomsky (2000): 16-17).

Even among those who aren't compelled to entirely abandon referential semantics, there's consensus that accounting for copredication requires radical semantic and/or metaphysical commitments. Semantically, some have taken copredication to show that types are extremely fine-grained (Pustejovsky (1995), Luo (2012), Asher (2011)), or that predicates encode criteria of individuation to which quantifiers are sensitive (Gotham (2016)). Metaphysically, it has been argued on the basis of copredication that there are complex objects with both physical and abstract aspects to which predication and quantification are sensitive (Gotham (2016)), and moreover these require us to accept the doctrine of relative identity (Asher (2011)).

By contrast, we argue that accounting for copredication requires no revisionary semantics or metaphysics, and that copredication is perfectly compatible with standard referential semantics. Our view is that previous discussions of copredication are driven by two false views: one linguistic and the other metaphysical. The false linguistic view is that nouns like "book" are ambiguous (or at least polysemous). Against this, we'll argue that "book" has a single sense and it designates both informational and physical books (section II). The false metaphysical view is that there are categorical constraints on property instantiation that ensure that, e.g. informational books can't be on shelves and physical books can't be informative. Against this, we'll argue that instantiation is not so selective and objects like informational books can instantiate properties like being on a shelf, even if this instantiation is derivative (section III).

After we've defended the components of our view, we will examine a variety of examples of copredication (section IV), arguing that our view provides the correct truth-conditions in all these cases. We will also discuss (section V), three potentially difficult cases for our view, but show that we can handle these as well. Finally, we'll compare our view to alternative views (section VI).

II. Univocity

² Pietroski (2005) and Collins (2009; *forthcoming*) concur with Chomsky.

The primary evidence for the polysemy of nouns like “book”, “lunch”, and “bank” is that quantificational sentences containing those nouns admit multiple readings. Reconsider (1):

(1) Three books are on the shelf.

In the situation envisioned, there are three copies of *Naming and Necessity* on an otherwise empty shelf. (1) has both a true reading and a false reading. The true reading can be forced by focusing on book factories. Imagine that Dave works at the factory and his task is to print 10,000 copies of *Naming and Necessity* by the end of the day. Unfortunately, the equipment is malfunctioning and he was only able to produce three copies by noon. At noon, in the depths of despair, he utters “What a complete disaster! I’ve been working all day and look what we have to show for it: three books are on the shelf.” Dave’s utterance of (1) is true.

Imagine, however, that Ofra manages the catalogue for the book factory and she wishes to know how many books from the catalogue have been produced and put on the shelf that day. If she calls the factory manager and asks how many books are on the shelf, and they respond with (1)—knowing full well that she’s asking about the number of books in the catalogue—then their utterance of (1) is false.

The standard view is that (1) has two readings due to the polysemy of “book”: it has one sense that designates informational books, and a second sense that designates physical books. By contrast, our view is that “book” has a single sense and, used with this sense, it designates both physical and informational books. However, like all nouns, its domain can be contextually restricted (cf. Stanley and Szabó (2001)). The two natural readings of (1) come from two different natural ways to restrict the extension of “book”: to informational or to physical books. To argue for this, we’ll first deflate the main consideration in favour of the ambiguity approach, then we’ll give two independent reasons to favour univocity.

The main consideration in favour of the polysemy of “book” is that counting sentences like (1) have multiple readings. However, there are analogous sentences with multiple readings, where we shouldn’t treat the relevant noun as ambiguous.

(6) Three colours are on the canvas.

(7) Exactly one wooden thing is in the room.

Imagine that the canvas has been painted with three shades of red. It seems that (6) has both a true reading and a false reading. If we focus on highly determinable colours like red and green, then it is false. If we focus on more determinate colours like crimson and maroon, it is true.

Imagine that the only wooden object in the room is a table with a top and four legs. Focusing on entire pieces of furniture, (7) is true. However, if we are particularly interested in the production of wood tables, it is easy to read (7) as false, due to the existence of more than one wooden table leg.

We should not treat “colour” and “wooden thing” as polysemous because of the multiple readings of (6) and (7). Treating “colour” as polysemous would yield a huge number of senses because there are an arbitrary number of levels of colour specificity on which we can focus. Just as we can treat different shades of red as a single colour, we can treat different shades of maroon as a single colour, while treating maroon as distinct from crimson. There simply aren’t two natural “levels” of colour specificity such that “colour” is naturally ambiguous between them. The point may be even clearer with “wooden thing”. We may treat entire pieces of furniture as single wooden things, or their salient parts, or their smallest parts, or their smallest perceptible parts, or any other “level” which happens to suit our interests. Given the infinite number of possibilities, an ambiguity view would result in an infinite number of senses for “wooden thing”.

Furthermore, there is an independent explanation of the multiple readings of (6) and (7) which requires no such proliferation of senses: that “colour” and “wooden thing” are contextually restricted.³ We can naturally restrict each noun to a subset of entities in their extensions. One lesson of (6) and (7) is that multiple readings of counting sentences may be due to contextual domain restriction on the restrictor noun, rather than its ambiguity.

There is another important lesson we can draw from “colour” and “wooden thing”. Consider a natural objection to the view that “book” has a univocal sense: if “book” has a univocal sense, we should be able to use it to count all of the things it designates at once. This, however, leads to false predictions about which readings are possible. It leads us to expect, for instance, that (8) has a true reading in the scenario in which three copies of *Naming and Necessity* are on the shelf (and nothing else is).

(8) Four books are on the shelf.

After all, if each physical copy is in the extension of “book”, as well as the informational book itself, then there are four distinct entities in the extension of “book”. (Is it of any help to insist that the informational book is not on the shelf? No: both the ambiguity view and the context-sensitivity view must hold that it is on the shelf, or else neither would be able to explain how “(Exactly) one book is on the shelf” has a true reading.)

To defuse this objection, consider an attempt to use “colour” or “wooden thing” in analogous unrestricted senses.

(9) Four colours are on the canvas.

(10) Six wooden things are in the room.

Imagine, again, that we’ve painted with three shades of red. Given that each shade of red is a colour and red is itself a colour, using “colour” in a completely unrestricted sense would lead us to predict that (9) has a true reading. However, it seems that (9) has no true reading. Or, more

³ In fact, Chierchia (2010) gives a context-sensitive account of “thing”.

cautiously, that we cannot achieve a true reading without a substantial amount of contextual supplementation.⁴ Similar remarks apply to (10). Intuitively, it is clear that counting both a determinable and its determinate is ruled out due to some restriction on double counting, as is counting both an object and its parts. In other words, there's an independent consideration ruling-out counting with "colour" or "wooden thing" in their most unrestricted senses.

An interesting and difficult project is determining the nature of the constraint on double counting. In particular, it is not clear from these examples either what the content of the constraint is—which cases count as double counting—or what the strength of the constraint is—how it can be overridden.⁵ For our purposes, however, the exact nature of the constraint is unimportant. What's relevant is that the fact that, for some nouns, it is extremely difficult to count with them in their most unrestricted sense. Our hypothesis is that this explains why we cannot easily generate a true reading of (8).

The upshot is that counting sentences don't give us strong reason to favour the ambiguity view of "book". The exact same phenomena that occur in counting sentences with "book" also occur in counting sentences where the relevant nouns are restricted in context rather than ambiguous. We'll now turn from arguing that there's no good reason to treat "book" as ambiguous, to arguing that there are good reasons to treat it as unambiguous.

One reason we gave to treat "colour" and "wooden thing" as restricted by context rather than ambiguous is that treating them as ambiguous would lead to an implausibly arbitrary number of senses. In fact, this applies to "book" as well.

Consider a bookshelf with two copies of the French translation of *Naming and Necessity*, and one copy of the Hebrew translation. Now consider the following sentences:

- (11) One book is on the shelf.
- (12) Two books are on the shelf.
- (13) Three books are on the shelf.

We contend that each has a true reading, as well as some false readings. To access the true readings, we can expand the sentences as follows:

- (11*) One book is on the shelf: *Naming and Necessity*.

⁴ Here is one special context which might achieve such an unrestricted reading of "colour": an art teacher gives the students a list of items, where each student needs to pick an item to constitute the theme of their next painting. Suppose the list reads as follows: "Red; Crimson; Scarlet; Maroon; Square; Triangle; Anger; Happiness;". In this context, it seems acceptable to say "Four colours are on the list: red, crimson, scarlet, and maroon". However, note that exactly the same kind of contextual supplementation would allow for an unrestricted sense of "book" (cf. our "history of the book" example below).

⁵ Some work has been done investigating double-counting, though nobody has advanced a theory that would answer basic questions about its source and force. See Kratzer (2016) and Liebesman (2016) for such discussions.

(12*) Two books are on the shelf: the Hebrew translation of Naming and Necessity and the French translation of Naming and Necessity.

(13*) Three books are on the shelf: the one on the left, the one in the middle, and the one on the right.

On our univocal proposal, these readings are accounted for, as above, by domain restriction. However, the view that accounts for different readings of counting sentences by positing ambiguity would now have to posit at least three meanings of “book” (in addition to the physical sense, there will be at least two kinds of informational senses – one according to which the translation of the same book into two different languages counts as the same book and one according to which they do not). Reflection on similar issues concerning the individuation conditions for books (Do two editions of the same book count as the same book? Does a second physical copy with one word misprinted count as the same book as the original?) suggests that, on an ambiguity view, “book” would have to have a large number of senses.

A second reason to treat “book” as univocal is that there are sentences in which “book” must be true of both informational and physical copies.

Consider a history of the book course in which the students learn general facts about the evolution of books from a textbook and then utilize this information to examine particular copies. On the syllabus, the professor may write the following:

(14) Each week, you will be assigned a single book, on which you will be required to submit a report: on odd weeks, your book will be a theoretical book on the history of book making, while on even weeks, our special collections librarian will assign you a volume from our historical collection to examine.

“Book” in (14) ranges over both physical and informational books: the books the students are mastering on odd weeks are informational books, while the ones they are mastering on even weeks are physical books.⁶

A final, perhaps more flat-footed consideration in favour of univocity is this: suppose you are pointing at a copy of War and Peace and say ‘This isn’t a book!’. If ‘book’ is ambiguous between a meaning that is only true of physical books, and a meaning that is only true of informational

⁶ In fact, (14) is similar to actual sentences you find in descriptions of history of the book courses. Here, for example, are two taken from Harvard’s description:

(14’) Drawing on the rich collections of Harvard’s libraries and museums, learners are invited to explore the book not simply as a container of content, but as a meaningful physical object that has shaped the way we understand the world around us.

(14’’) The first three modules (sic.) focus on the advent of printing in Europe, the role of books as physical and intellectual objects during the Enlightenment, and the rise of the reading public.

books, the sentence should have a true reading.⁷ But it really does not seem to have a true reading.

Our goal in this section has been to defend the view that nouns like “book” aren’t polysemous. While that is our considered view, it is worth stressing that it is not central to the rest of our discussion. We will next turn to defending the metaphysical view that in true copredication sentences, there is a single object that instantiates both predicated properties, e.g. an object that is both heavy and informative (a book), or that is both delicious and long (a lunch). That view is independent of the (non)ambiguity of the nouns that designate those objects. Furthermore, if the primary desideratum on a view of copredication is to explain how sentences like (2)-(5) have true readings, then our metaphysical view will suffice. That said, the linguistic and metaphysical components of our view do cohere particularly well, for a reason we’ll return to.

III. Inheritance

The metaphysical component of our view concerns the relationship between the different sorts of entities designated by “book”. A central supposition underlying typical discussions of copredication concerns which properties can be instantiated by physical books, and which can be instantiated by informational books. In particular, the key challenge raised by copredication arises from the supposition that the properties ascribed in copredicational sentences cannot be instantiated by both. For instance, physical books can be heavy but not informative and informational books can be informative but not heavy.

Our view is that this is bad metaphysics. Informational books are distinct from physical books, but there are many properties that both can instantiate. This may seem particularly surprising when it comes to properties that are most familiarly instantiated by one or the other: e.g. being informative or being on a shelf. However, in many such cases one object instantiates the property only derivatively, inheriting it from the other.⁸ For instance, a physical book may be informative, but in typical cases, when it is informative, that is in virtue of the fact that it is an instance of an informative informational book.⁹

⁷ Note that it doesn’t matter whether ‘This’ here picks out a physical or an informational book. Either way ‘book’ could receive the reading that applies only to the other kind of books, rendering the sentence true.

⁸ Liebesman (2011) defends the view that kinds can inherit properties from their instances, alongside the view that all generics should be analysed as kind-predications with the truth-conditional variance explained by property inheritance. The view we defend in this paper coheres straightforwardly with that view of generics, though nothing in our approach to copredication *requires* endorsing the view of generics given in Liebesman (2011).

⁹ We say “in typical cases”, because sometimes the fact that a physical book is informative has nothing to do with the informational book it instantiates. For example, if the detectives find a physical book with a bloodstain on it on the scene, the book might be informative in the sense that it is providing the detectives with pertinent information.

On the most ham-fisted version of this view one accepts the following:

- (i) For every property P , if there exists a physical book which has P , then the corresponding informational book has P .
- (ii) For every property P , if an informational book has P , then every physical instance of it has P .

This version of the view clearly fails. Some physical copies of *Naming and Necessity* have maroon covers, while others have orange covers. The copies with orange covers have the property of being non-maroon. If the informational book inherited this property, and it inherited the property of being maroon (from the maroon copies), then a single object would be both maroon and non-maroon. Moreover, since on the proposed view each physical book inherits every property that the corresponding informational book has, each physical copy will be both maroon and non-maroon as well. That's impossible, so property inheritance must be more discriminating. Our task in this section is to advance a more discriminating view.

One way to give a more discriminating view would be to endorse a different quantificational schema regarding how many physical books must instantiate P in order for the corresponding informational book to instantiate P .¹⁰ We reject any such straightforward account: such accounts are false because inheritance, and property-instantiation more generally, is not so systematic. An analogy makes the point vivid. Complex objects can inherit properties from their proper parts. A table may be green in virtue of all of its proper parts being green; a symphony may be beautiful in virtue of each of its movements being beautiful. However, there is no simple quantificational account of when a complex object inherits properties from its proper parts.¹¹ For a table to have the property of touching a wall it suffices for just some of its parts to touch the wall. By contrast, a person may be animate even though each of their proper parts is inanimate. Whether a complex instantiates a property depends not just on how many of its parts instantiate that property, but also on (at least) the nature of the property, the nature of the parts, and the nature of the complex. Given the variety of properties, parts and complexes, we shouldn't expect a single straightforward quantificational account. This lesson carries over to informational objects.

III.i The Properties of Informational Books

Our strategy for discerning the properties of informational books is simple: we identify uncontroversially true sentences in which the subject term designates an informational book i and the predicate ascribes a property P , and we conclude that i instantiates P . Stated abstractly, the test is uncontroversial. What's more controversial, of course, is to assess whether a given sentence has this proposed structure.

¹⁰ For example, one might suggest replacing (i) above with the claim that if *every* physical instance has P , then so does the corresponding informational book. It is easy to see that this fails, though, when ' P ' is taken to be: "is a physical book".

¹¹ These examples were first given in Liebesman (2011).

Suppose we are in a bookstore, which contains many copies of *War and Peace*, all located on the top shelf. Consider an utterance of the following sentence:

(15) *War and Peace* is on the top shelf.

Given that in this context there is no particular physical object for “*War and Peace*” to designate (there are many copies, none of which are particularly salient), we claim that the term designates the informational book.¹² More precisely, we hold the following three theses about (15):

(T1): (15) has the simple subject/predicate form of atomic sentences, in which “*War and Peace*” is the subject and “is on the top shelf” is the predicate.

(T2): “*War and Peace*” in (15), in the envisioned context, designates an informational book.

(T3): “is on the top shelf” in (15), in the envisioned context, designates the same property as when it is used to ascribe a property to a physical entity.

T1-T3 are motivated by considering a number of linguistic phenomena. Though each of these is, in principle, compatible with the denial of T1-T3, they jointly provide the premises of a strong abductive argument.

The first phenomenon, which motivates T1, is anaphora support. “*War and Peace*”, as it occurs in (15) supports subsequent singular anaphoric pronouns, as in (16):

(16) *War and Peace* is on the top shelf. Reading it takes forever.

The best explanation for this anaphora support is that “*War and Peace*” is a singular referential term. To see this, consider an alternate hypothesis on which (15) contains implicit quantification over physical copies, so “*War and Peace*” isn’t a term referring to an informational book, rather it is a predicate designating physical books. On such a view, it would be wholly mysterious that the predicate supports a singular anaphoric pronoun. Consider, for instance, (17) and (18):

(17) # Every dog is barking. It is disobedient.

(18) # Some dogs are barking. It is disobedient.

“Every dog” and “some dogs” can support anaphora, but only plural anaphora. By contrast, singular existential quantifier phrases can support singular anaphoric pronouns, as in (19)

(19) Some dog is barking. It is disobedient.

¹² In fact, we are tempted by the view that the name “*War and Peace*” never designates a physical book, always an informational book, but we will not argue for this here.

In order, then, for a quantificational view of (15) to make sense of anaphora-support as illustrated in (16) we would have to hypothesize that the implicit quantifier is a singular existential.¹³ However, while it is arguably necessary and sufficient for the truth of (15) that some physical copy of *War and Peace* is on the top shelf, maintaining that this follows from the general structure of the sentence (rather than, as we maintain, from the metaphysics of what it takes for an informational book to be on the top shelf), leads to truth-conditional disasters. For on such a hypothesis, (20) should have a reading that is truth-conditionally equivalent to (21), and (22), a reading that is equivalent to (23):

- (20) *War and Peace* is read by many people.
- (21) Some physical copy of *War and Peace* is read by many people.
- (22) *War and Peace* was written by Tolstoy
- (23) Some physical copy of *War and Peace* was written by Tolstoy.

But (20) and (22) simply do not have readings with these truth-conditions: note for example that either sentence can be true even if every single physical copy of the book has been destroyed, and at least the latter can be true even if no physical copy has ever existed, e.g. because Tolstoy had never bothered transcribing the book.

The second phenomenon, which supports both T1 and T2, is that of true quantificational sentences. As we've already noted, (1) (repeated here as (24)) has a true reading even when there are three copies of *War and Peace* on the top shelf:

- (24) One book is on the top shelf.

For (24) to be true it must have a witness: a single entity that is a book on the shelf. We can easily name this witness, as in (25):

- (25) One book is on the top shelf: *War and Peace*.

"War and Peace" in (25) must refer to a single entity, or it would not name something that witnesses the truth of (24). Given that there are three physical books on the top shelf, the single entity it names cannot be a physical book.

Indeed, such quantificational uses provide the primary test for discerning properties of informational books. Thus note, for example, the contrast between the following two utterances:

- (26) One book is on the top shelf: *War and Peace*, and it is 587,287 words long.

¹³ We could hypothesize that the quantifier is a definite article (so the sentence is read as saying that "The *War and Peace* is on the top shelf"), but this will not work in our current context where there is no salient copy of the book.

(27) #One book is on the top shelf: *War and Peace*, and it is worn out.

(26) seems perfectly acceptable in context, suggesting the informational book has the property of being 587,287 words long, while (27) sounds infelicitous (even assuming that all copies of *War and Peace* on the top shelf are worn out), suggesting the informational book doesn't have the property of being worn out.

The third phenomenon, which supports T1, is that of coordination. Notice that quantificational NPs do not usually conjoin with names to form complex phrases¹⁴:

(28) # Somebody and Emma are barking.

However, we can freely coordinate within categories:

(29) Oli and Emma are barking.

"War and Peace" can freely conjoin with proper names, which at least puts pressure on the view that it acts in this context as a quantificational expression or bound predicate.

(30) *War and Peace* and Emma are both on the rug.

The fourth phenomenon, which supports T3, is that of predicate elision.

(31) *War and Peace* is on the top shelf, as is Ofra's blue pencil.

In (31) we attribute the very same property to *War and Peace* as we do to Ofra's blue pencil. Since the latter is a physical entity, then the property we attribute to *War and Peace* is the very same property we can attribute to physical entities.¹⁵

Each of these four phenomena is relatively straightforward and unsurprising. Together, they support T1-T3. We doubt that our opponents would deny any of this. Their rejection of T1-T3 comes not from a denial that they are supported by linguistic phenomena, but, rather, from a belief that T1-T3 are otherwise unacceptable. Arguments for the unacceptability of T1-T3 have a familiar form: (i) If T1-T3 are all true then some informational book is *P* (ii) No informational book is *P*; Therefore, (iii) At least one of T1-T3 is false.

¹⁴ Arguably, they can be conjoined with names in special cases such as 'Emma and someone else are talking', though even this strikes us as at best borderline acceptable. (Thanks to Paul Elbourne for raising this point.)

¹⁵ Can "Ofra's blue pencil" refer here to a type of pencil rather than to a physical object? To fend off this worry, imagine that Ofra owns two qualitatively identical blue pencils, and replace (31) with "*War and Peace* is on the top shelf, as are both of Ofra's blue pencils". The fact that we are assuming that there are *two* blue pencils (rather than one) indicates that we must be referring to physical objects.

The most familiar instances of this form arise in the literature on the ontology of art. Focusing on sentences similar to (15), some conclude that abstract fictional works (e.g. novels) exist and are distinct from their instances. Objectors give arguments in the familiar form for a variety of different properties. For instance, one may object that if *War and Peace* (the informational book) were to exist it would be a contingent non-eternal abstract object, but no such abstracta exist, therefore *War and Peace* does not exist. Walters (2013) contains an illuminating discussion of three such objections: that no abstract objects (i) can be created, (ii) can have non-essential intrinsic properties, and (iii) can have causal powers. Rather than repeating any of Walters' considerations, to which we are sympathetic, we'll focus on another version of the objection which we take to be stronger. The stronger objection shifts focus from relatively high-level properties like entering into causal relations to more familiar properties. Kleinschmidt and Ross (2012) give a particularly vivid instance of the objection focusing on the temporal properties of musical works. We'll follow them in moving from *War and Peace* to *The Moonlight Sonata*.

(32) The Moonlight Sonata is 15 minutes long.

The analogues of T1-T3 for (32) are as follows:

(T1'): (32) has the simple subject/predicate form of atomic sentences, in which "The Moonlight Sonata" is the subject and "is 15 minutes long" is the predicate.

(T2'): "The Moonlight Sonata" in (32), in the envisioned context, designates an informational musical work rather than a particular performance.

(T3'): "is 15 minutes long" in (32), in the envisioned context, designates the same property as when it is used to ascribe a property to a physical entity.

Usually when entities are 15-minutes long, we can ask when they happen. This question, however, seems to admit no obvious answer when it comes to an entity like *The Moonlight Sonata*. One may react to the oddity as follows: *The Moonlight Sonata* just isn't the sort of thing that can instantiate the property of being 15 minutes long. T1'-T3' are false because they commit us to the view that it does.

In particular, Kleinschmidt and Ross argue that T1'-T3' commit us to holding one of four unacceptable options about the temporal location of *The Moonlight Sonata*: (i) *The Moonlight Sonata* is temporally located where one of its "privileged" instances is located, (ii) *The Moonlight Sonata* is located in the fusion of all of the intervals at which its instances are located, (iii) *The Moonlight Sonata* is wholly present wherever each of its instances is, and (iv) *The Moonlight Sonata* has no temporal location.

Our view is that both (iii) and (iv) are coherent options (though we ultimately favour (iii)), and that Kleinschmidt and Ross' objections to them are unconvincing. Before explaining why, we will explain some key components of our metaphysics of informational musical works (and, more generally, objects that can have instances).

On our view, *The Moonlight Sonata* instantiates the property of being 15 minutes long, the very same property instantiated by token events, e.g. Ofra's drinking of coffee this morning. For token events with duration it is perfectly reasonable to ask when they occur. Ofra's coffee drinking, for instance, occurred between 7:32 and 7:47 AM. Generalizing from morning coffee, one may hold a principle like the following: if an entity has a temporal duration then it has a temporal location. It would then be perfectly reasonable to ask when *The Moonlight Sonata* is located. Our view is that the aforementioned principle is false: simply because something has a temporal duration doesn't entail that it has a temporal location.

How can it be that something with a temporal duration fails to have a temporal location? The key is in recognizing that there are different ways to instantiate durational properties. When it comes to event tokens, they have durational properties only when they have temporal locations as well. However, this principle does not generalize to durational properties of types. The principle that duration requires location fails, then, because it overgeneralizes from a particular way of instantiating durational properties.

This view may sound dramatic, but it is no more dramatic than the universally accepted view that properties can be instantiated in different ways. Consider determinables like redness. Objects can be red in lots of different ways: by being different shades. Consider properties that can be instantiated by both events and non-events, like being tolerable. It may be that certain sorts of condescending explanations are intolerable because one is miserable when listening to them. It also may be that a person is intolerable in virtue of their love of giving such condescending explanations. One and the same property—being intolerable—can be instantiated by an event in virtue of very different sorts of facts than those that ensure it is instantiated in a person. These are just two of many possible analogies to drive home a simple point: properties can be instantiated by different objects in different ways. Our thesis merely applies this to the property of being 15 minutes long as it is instantiated by token events and informational musical works. This is the first key component of our metaphysical picture.

Earlier, we stated that our univocal view of "book" coheres nicely with our metaphysical views. Now we can explain why. The mere fact that some occurrences of "book" designate informational objects and others designate physical objects does not, on our view, support the ambiguity of "book". Rather, we can draw a more mundane lesson: there are different ways to be a book. This puts "book" perfectly on par with other predicates and the property of being a book perfectly on par with other properties. To hold fast to the ambiguity of "book" while independently stressing the fact that there are many ways to instantiate ordinary properties and exemplify ordinary predicates would run the risk of arbitrariness. So, even if the combination of our metaphysical view with an ambiguity view of "book" isn't incoherent, it is unmotivated.

The second key component of our underlying metaphysical view, which we've already outlined, is that there is no simple quantificational formula for determining when a type has a property as based on when its instances have that property.

With these two components on the table we can return to Kleinschmidt and Ross' objection. They argue that none of their four options (i)-(iv) is plausible. We agree that neither (i) nor (ii) is particularly promising. As we stated, however, our view is that both (iii) and (iv) are *prima facie* reasonable options. We'll begin with (iv).

Kleinschmidt and Ross argue that (iv) is false because it leads to paradox. In particular, they think the proponent of (iv) must hold the following principle: if, necessarily, each instance of The Moonlight Sonata is P, then The Moonlight Sonata is P. That the principle leads to a contradiction is straightforward: necessarily each instance of The Moonlight Sonata is a non-kind. By the principle, this entails that The Moonlight Sonata is a non-kind. However, this contradicts the supposition that The Moonlight Sonata is a kind.

Our response is predictable: it is false that if necessarily each instance of a kind K is P then K is P. No such quantificational claim about property inheritance holds. Kleinschmidt and Ross anticipate this response and argue that this view is unacceptable because it would leave us without a unified view of truth conditions of predications like (11). We think that the lack of a unified view is desirable. On our view, (11) belong to a much more general class: atomic predications. There is no reason to think we could give a reductive account of the truth conditions of atomic predications.¹⁶

There's a sense in which pointing out that (iv) is not paradoxical is our most important response to the objection. Quite generally, there is no good reason to think that kinds will have a temporal location even if they have a duration.

However, we do think that attention to the data suggests that (iii), namely the claim that The Moonlight Sonata is wholly located where each of its instances is located, is the correct view to take here. Indeed, this claim about temporal location is suggested by truths such as the following:

- (33) The Moonlight Sonata will start in 15 minutes.
- (34) The Moonlight Sonata started at 1:30 and finished at 1:45 every day this week.

(33) is true in a scenario where we attend a concert. (34) is true when speaking about the practice regimen for that concert. To determine whether the truth of (33) and (34) entails that The Moonlight Sonata has locational properties, we need to determine their correct analyses.

In fact, the tests we've already considered, as well as the implausibility of alternate analyses, leads to the conclusion that (33) and (34) are best analysed as atomic predications that ascribe

¹⁶ Lewis (1983) makes this point in a critical discussion of Armstrong (1978). Note, also, that claiming that there is no reductive account of the truth conditions of atomic predications doesn't commit us to the view that nothing at all informative and general can be said about predication. Liebesman (2015) gives such a non-reductive account.

familiar properties to an informational object. As the following examples demonstrate, (33) and (34) support anaphora, true quantificational sentences, coordination and predicate elision:

(35) The Moonlight Sonata started at 1:30 every day this week. It was composed by Beethoven.

(36) One (and the same) piece started at 1:30 every day this week: The Moonlight Sonata.

(37) The Moonlight Sonata and Professor Star's colloquium both started at 1:30 today.

(38) The Moonlight Sonata started at 1:30 today, as did Professor Star's colloquium.

We will not rehearse the argumentation again, but for similar reasons as above, we think this data strongly suggests that (33) and (34) do ascribe temporal location to The Moonlight Sonata. Ultimately, then, we hold Kleinschmidt and Ross' option (iii): The Moonlight Sonata is located where its instances are. Importantly, though, we do not hold this view because of some general connection between temporal duration and temporal location. Rather, we hold the view because a specific analysis of truths about The Moonlight Sonata supports it.

Before concluding our discussion of the properties of informational objects, we want to consider one more objection to our favoured position that The Moonlight Sonata has both a temporal location and duration. The objection is raised by Kleinschmidt and Ross, who start out by assuming the following two principles (2012: 132):

Locative Claim 1: each part of The Moonlight Sonata is located where and only where it is performed.

Locative Claim 2: if some but not all of a wholly material object, *o*, is present within some region, *r1* then there must be some other, disjoint region, *r2*, such that some of *o* is present in *r2*, and the parts of *o* present in *r1* and the parts of *o* present in *r2* together make up *o*.

Here is why they take Locative Claim 1 and Locative Claim 2 to make trouble for our view. Imagine that there are exactly two performances of The Moonlight Sonata, and the later performance omits the third movement. From Locative Claim 1, it follows that the first two movements are located at the later time. Then, from Locative Claim 2, they infer that there is another part of The Moonlight Sonata, located at the earlier time such that this part together with the parts located at the later time make up The Sonata. However, according to them, this leads to the rejection of the following claim:

(39) The Moonlight Sonata is such that its third movement always comes after its first and second movement.

Finally, they take this rejection to be absurd.

There are several problems with this objection. To begin with, Locative Claim 2 is not clearly applicable to this case – after all, it concerns wholly material objects and it is far from clear that The Moonlight Sonata is a wholly material object (for one thing, it presumably exists even if no

tokens of it exist). More importantly, even if it were applicable, Locative Claim 2 is false: imagine that there is only ever one performance of The Moonlight Sonata and it omits the third movement. In that case, there is obviously no disjoint region that contains an object that combines with the incomplete performance to compose The Moonlight Sonata. This is exactly what we'd expect given that there can be informational objects that are only ever incompletely exemplified, as well as informational objects that have no instances whatsoever.

Even if we assumed Locative Claim 2 were true and applicable, it is hard to understand what Kleinschmidt and Ross's complaint amounts to. We are not sure what exactly they mean by saying that the third movement 'always comes after' the first two movements. If what they mean is that tokens of the third movement always come after tokens of the first two movements, that's clearly false – and for reasons that have nothing to do with incomplete performances. If there are 100 complete performances of The Moonlight Sonata, then 99 token performances of the third movement occur before the final token performance of the first and second movements. Perhaps a better way to phrase their worry would be to complain instead that we must reject the following:

- (40) The first and second movement of The Moonlight Sonata precede its third movement.

Imagine now that there are only ever two incomplete performances of The Sonata: one consisting of just the third movement, and a later performance of just the first two movements. Since we accept that each movement is temporally located where its performances are, it follows from our view that the (only) temporal location of the third movement is earlier than the (only) temporal location of the first and second movements. And one might imagine that this is sufficient to lead to a rejection of (40).

This objection, however, relies on an overly simplistic view of the *precedence* relation. In the case of movement tokens, it is plausible that precedence is determined solely by relative temporal location. However, this is implausible for movement types. For movement types, precedence is plausibly determined by the properties The Moonlight Sonata itself. One reason to hold this is that there can be precedence relations between the movements even if there are no performances of The Sonata whatsoever. We thus conclude that incomplete performances give us no reason to reject (40).¹⁷

III.ii The Properties of Physical Books

¹⁷ We also note in passing that Kleinschmidt and Ross's own positive view concerning these cases presumably does require them to reject (40). Roughly, their view is that (40) has the following truth-conditions: Generically, things that are members of the plurality of token performances of the first and second movements, precede things that are members of the plurality of token performances of the third movement. But as the case of 100 complete performances demonstrates, it is entirely typical and mundane for token performances of the third movement to come before token performances of the first two, so it is hard to see why this generic claim would be true.

Our strategy for discerning the properties of informational books was straightforward: find a true subject/predicate sentence where the subject designates an informational book, and determine which property is ascribed by the predicate. Our strategy for discerning the properties of physical books is the same.

According to Wikipedia, The Leningrad Codex is the “oldest complete manuscript of the Hebrew bible in Hebrew.”^{18 19},

(41) The Leningrad Codex is informative.

(41) is true. Our contention is that its subject term refers to a physical book and its predicate ascribes the very same property that is instantiated by informational books. Here are analogues of (T1)-(T3) for (41):

(T1’): (41) has the simple subject/predicate form of atomic sentences, in which “The Leningrad Codex” is the subject and “is informative” is the predicate.

(T2’): “The Leningrad Codex” in (41), in the envisioned context, designates a physical book.

(T3’): “is informative” in (41), in the envisioned context, designates the same property as when it is used to ascribe a property to an informational entity.

(T1’) and (T2’) can be supported by similar linguistic phenomena to those that supported (T1) and (T2), but we suspect that in this case they will be taken as uncontroversial.²⁰ (T3’) is supported by the acceptability of the following:

(42) The Leningrad Codex is informative, as is Biblia Hebraica.

Note that “Biblia Hebraica” does not refer to a single physical book. Rather, it refers to three editions of the Bible edited by Rudolf Kittel. Though Biblia Hebraica has many physical instances, it is not itself a single physical book. (42), then, uses predicate elision to illustrate that a physical book (The Leningrad Codex) can instantiate the property of being informative—the very same property instantiated by a book type (Biblia Hebraica).

¹⁸ https://en.wikipedia.org/wiki/Leningrad_Codex

¹⁹ Note that there are several reasons for which The Leningrad Codex can be informative: it can be informative in virtue of the informational book it exhibits or it can be informative in virtue of its physical attributes (for example, the fact that it was written on a particular kind of parchment might be informative for historians of the period). To focus on a case where it is even controversial that The Codex has the property, imagine that one is uttering the sentence in a context where the focus is on the informational book the Codex exhibits.

²⁰ One could propose an unusual alternative account according to which “The Leningrad Codex” denotes a very specific informational book, one that essentially has exactly one physical instance. However, this hypothesis does not explain the acceptability of the following:

(41’) There is exactly one oldest intact copy of the Hebrew Bible: The Leningrad Codex. It is informative.

Note that “Leningrad Codex” here is the witness for the predicate “*copy of the Hebrew Bible*” and hence denotes a physical book, one which is then said to be informative.

As was the case with book types, we don't think there is any simple recipe for determining when a physical book instantiates a particular property. Rather, we must test on a case-by-case basis. The fact that physical books can be informative, however, shows that there aren't the strong sortal restrictions on instantiation sometimes supposed in the copredication literature.²¹

The issue of sortal restrictions has been, perhaps, one reason that a theory of the sort we develop here has not been defended in the past. The worry is that attributing the property of being informative to a physical book, or alternatively attributing the property of being on the shelf to an informational book, would result in a *category mistake* – a sentence that is (in some sense or other) inappropriate in a similar manner to “The number two is green” or “Saturday is having breakfast”.²²

The worry is, as stated, odd: after all, while there are a range of different competing accounts of the infelicity of category mistakes, the one aspect that these accounts share is the explanandum – namely the claim that category mistakes *are infelicitous*. But sentences such as “The Codex Bible is informative” or “War and Peace is on the top shelf” are *not* infelicitous, and thus the worry that they give rise to a category mistake seems unmotivated.

Nevertheless, there is a more nuanced version of the worry that runs by considering sentences such as the following:

- (43) * The number two is on the top shelf.
- (44) * The property of being green is on the top shelf.
- (45) * The tree is informative.
- (46) * My house is informative.

All of these sentences seem to be clear cases of category mistakes.²³ Moreover, a tempting theory of what accounts for the fact that these sentences are category mistakes is that the predicate “is on the top shelf” is restricted so that it can only legitimately apply to physical objects and thus predicating it of an abstract object such as the number two or the property of being green results in a category mistake. Similarly, the predicate “is informative” is restricted so that it can only legitimately apply to informational (or at perhaps abstract) objects and thus

²¹ For example, Collins (*forthcoming*) says “perfectly quotidian nominals, such as *London* or *book*, may occur copredicatively as a single argument of categorically mismatched predicates, which prima facie preclude a coherent uniform construal of the nominal argument” (p. 1-2).

²² See e.g. Asher (2011), Luo (2012), or Gotham (2014). Note that Gotham and Luo allow that it is always sortally acceptable to predicate “informative” and “heavy” of books, but this is because they take “book” to denote a special type of object which is a kind of “product” of an informational book and a physical book, and hence counts as satisfying both of the sortal restrictions.

²³ At least in the absence of substantial contextual supplementation. ((45) seems acceptable in a context where the tree has some important information written on it, or alternatively if the very presence of particular physical properties is informative in context.)

predicating it to a physical object such as a tree or a house results in a category mistake. (For current purposes we can remain neutral on what the source of this “restriction” is – i.e. whether it is syntactic, semantic, or pragmatic.)

The worry is then that these restrictions predict that the following sentences, if they are to be analysed along the line we propose – should also constitute category mistakes:

- (47) *War and Peace* is on the top shelf.
- (48) The Leningrad Codex is informative.

After all, according to our analysis, in (47), “is on the top shelf” is predicated of an informational book, and in (48), “is informative” is predicated of a physical book. But this gives rise to the objection: since (47) and (48) are perfectly felicitous, they are obviously not category mistakes. The prediction, the objection goes, is thus false and our account fails.

Our response to this objection is that it relies on an incorrect account of category mistakes. Magidor argues that the kinds of requirements that account for category mistakes are specific adjectival properties rather than type (or “is-a”) restrictions.²⁴ We will not repeat the argument in full here, but briefly, consider a paradigmatic category mistake such as the following:

- (49) *The number two is green.

While it might be tempting to think that “is green” imposes on its subject term the constraint that it is a concrete object, that hypothesis is at odds with the fact that the following speech is perfectly felicitous (even if false):

- (50) *Physicist*: you know, even though photons aren’t concrete objects we discovered they do have colours. This photon, for example, is green.

Rather, the predicate “is green” requires that the subject be *coloured*.²⁵

Similarly, we maintain that the infelicity of (43)-(46) can be explained by the claim that the predicate “is on the top shelf” gives rise to the restriction that its subject *is located*, and the predicate “is informative” gives rise to the restriction that its subject *is either informative or uninformative*.²⁶ This kind of account allows both for an explanation of why (43)-(46) are category mistakes (because it’s clear in context that the number two is not located and that my house is not informative or uninformative) at the same time as allowing that (47) and (48) can

²⁴ Magidor (2013), Ch. 5 §4 and Magidor (2016): 562-63.

²⁵ Magidor (2013) argues that the nature of the requirement is that it generates a pragmatic presupposition, but as noted above the current discussion is largely orthogonal to the precise source of sortal requirements.

²⁶ The precise formulation of the sortal restrictions in question can be debated (and for some such requirements there may not be a simple way to phrase them in English), but the crucial point is that the requirement is not a type one (see Magidor (2013), ch.5 §4 for discussion.)

perfectly felicitous. According to our account type-books can be located, and physical books can be informative, so no restriction is violated.

IV. Our view in action

With all of the components of our view explicated and defended, we can now revisit our initial examples of copredication and explain how our view handles them.

(2) Three red books are informative.

Theorists have been puzzled by sentences such as (2) because they think that only physical books can be red, while only informational books can be informative. We reject this. Physical books, on our view, can be both red and informative. The fact that physical books can be informative is explained by our analysis of (41), which entails that The Leningrad Codex—a physical book—is informative. So, on a true reading of (2), context determines that “books” designates physical books, and these are informative in virtue of the information they express.

Another possible treatment of (2) takes informational books to be red. We don’t rule out this treatment for any principled reason. Rather, we simply find that attention to the data in this case yields the result that informational books are not, at least in typical cases – red.²⁷ Consider, for example, a bookshelf with exactly three physical books – three copies of *War and Peace*, two with red covers and one with a blue cover. An utterance of “One book on the shelf is red” seems false in this context, contrary to the prediction we would get if it were true of the informational book that it is red.

Indeed, similar attention to data suggests that while it is easy to ascribe to physical books the property of being informative, physical books don’t typically instantiate other “informational” properties. Thus, for example, suppose a bookshelf contains two copies of book *A*, two copies of book *B*, and two copies of book *C*. Assume also, that only *A* and *B* (but not *C*) were printed in over 10,000 copies. In this context, it’s hard to get a true reading of the following:

(51) Four books on the shelf were printed in more than 10,000 copies.

This suggests that “was printed in more than 10,000 copies” does not apply to physical books.²⁸

²⁷ Some non-typical cases might be ones where all the copies of an informational book stably have a red cover (a good example might be *Quotations from Chairman Mao Tse-tung*). Another relevant case would be an informational book where the very information the book contains requires it to have a red cover (consider, e.g., a book on the nature of colours which mentions in the text the colour of the cover). In that case, even if there are currently no physical copies of the informational book we might want to classify it as red.

²⁸ We note in passing, that there are special ways to set up the context so that we get acceptable or borderline-acceptable readings of (51) here. Our view is that such readings are generated by a deferred reading of the predicate, but the details of the case would take us too far afield.

Next consider:

(3) Lunch was delicious, but lasted hours.

The worry about (3) is that only events can last hours but only non-events can be delicious. Our view is that non-events can also last hours, though they have this property in virtue of some related event. Consider, for instance, (52):

(52) That dog bone lasted two minutes.

When I give Emma a bone that only takes her two minutes to consume, (52) has a true reading. On the true reading, “that dog bone” designates a non-event object that has the property of lasting two minutes. It has this property in virtue of the fact that the event of Emma consuming the bone was two minutes long.

Similarly, some food served midday may have the property of lasting hours in virtue of the associated event. This is most obvious when the associated event of, e.g. physically consuming the food, is hours-long. However, we need not be so constrained in associating an event with an object. A bone may last two hours when Emma is amused by it for two hours, and some food may last hours when the associated event of sitting in a dining hall lasts hours—even if the food is itself consumed before the end of that event.

Again, another possible option is to argue that an event may be delicious.²⁹ We have no principled objection to that option, we simply find that the data suggests that it is not a promising analysis of this case. This is so because it seems hard to find true readings of sentences that explicitly refer to an event as delicious, even when we describe the event in a way that make its association with food very salient. Thus, for example, the following two sentences sound entirely infelicitous:

(53) The lunch meeting was delicious, but lasted hours.

(54) The meeting, at which lots of food was served, was delicious but lasted hours.

Finally consider:

(4) The bank was vandalized after calling in Bob's debt.

On our analysis, institutions can be both vandalized and call in debt. They can be vandalized in virtue of one of their physical buildings being vandalized. (They call in debt in virtue of one of their employees calling in debt, though, for whatever reason, it is not controversial in the literature that institutions really do call in debt.) That institutions can be vandalized is supported by sentences like (55):

²⁹ In fact, this is suggested (but rejected) by Cooper (2007).

(55) The Royal Bank of Canada, which was vandalised earlier this year, had to be moved following the incident.

(55) can clearly be true in a scenario where there is no single building that was vandalised and then moved. “The Royal Bank of Canada” in that case, must refer to an institution, which was itself vandalised.

V. Problem cases

Having defended the core of our account, in this section we consider a range of examples that might be thought to pose some special difficulty, and discuss them in turn.

V.I Destroying physical instances

Asher provides the following objection against theories that postulate two readings of “book” (an informational and a physical one), and allow for a sense in which informational books can be burned or physical books mastered.³⁰

(56) John burned a book and then mastered it.

Asher’s worry is that “the account predicts that intuitively *a priori* false or odd sentences can be true” (p. 137). Asher’s objection seems to be the following. Consider a situation where John burns one copy of *War and Peace*, and then masters the book, using a different copy. An account such as ours, the worry goes, would predict that on every reading, (56) is true relative to this situation. This is so because (as the objector assumes the theory goes) either “book” is interpreted as an informational book and informational books can be burned in virtue of one of their instances being burned, or “book” interpreted as physical book and physical books can be mastered in virtue of mastering the information in that physical book (however it is instantiated).

The problem is that Asher seems to be implicitly assuming here the most ham-fisted version of property inheritance (one where if a single physical instance has *P* so does the informational object, and where if the informational object has *P* so do all its physical instances). But as we have stressed in section III, we do not accept any such general principles of property inheritance, and certainly not this particular schema (which trivially leads to inconsistency, see section III).

Taking our more nuanced approach to property inheritance avoids Asher’s problem. Consider first the interpretation of (56) where “book” is restricted to physical books. We maintain that physical books *can be mastered*, but in order to master a physical book one must causally engage with that very copy. To see that this is so, consider a case where John’s shelf contains

³⁰ Asher (2011): 136-7. The theory Asher is considering here is somewhat different than ours, but the worry he raises applies just as much to our theory.

exactly three physical books: three copies of *War and Peace* that John has never touched, though he has extensively studied (another copy of) *War and Peace*. Our judgement is that in this situation it would be false to say “John mastered three books on his shelf”. This judgement, however, is reversed where John has extensively engaged with the particular details of each of the three volumes. Precisely what is required to master is a physical book is a difficult question, but what is clear is that no one can satisfy that requirement with respect to a physical book that has been burned.³¹ With that in mind, it is clear that (56) at least has a reading - arguably the most salient reading - where “book” is restricted to physical books, and where the sentence would come out as false relative to every situation. We maintain that it is this reading that accounts for Asher’s judgement that (56) is “intuitively a priori false”.

Does the sentence also have another reading, according to which it is true relative to some situations? That depends on the question of whether type-books can be burned. As should be clear by now, there is no straightforward way to answer this question. However, even if it turns out that type-books *can* be burned, at most it would show that (56) has a true readings (as well as a salient false one), and moreover, that this reading would have the kind of truth-conditions which make it completely unsurprising that it is true.³²

V.II Referring to type-books indirectly

Consider the following case: last week, Jane picked up a blue-covered copy of *War and Peace* from a certain library. Today, she is standing in Jill’s office. Looking at a pile of books on the desk, she spots a red-covered copy of *War and Peace* and says:

(57) Hey, I picked up that red book from the library last week!

The utterance seems felicitous and true in the situation described. But it is not trivial to account for this fact. If “that red book” refers to the physical book in Jill’s office then the utterance would be false (Jane didn’t pick up that very physical copy from the library last week). On the other hand, if “that red book” refers to the informational book *War and Peace* it seems that we must maintain that the type-book is red. The problem is that, as noted above, we think that – except in special cases – informational books do not inherit colour properties in this way. After all, if *War and Peace* is sometimes printed with blue covers and sometimes with red covers, there is no non-arbitrary way to determine whether it would be red or blue. And while one could perhaps accept that the type books is both red and blue (just as we accepted that it can

³¹ A trivial reason for this is that the physical book no longer exists after it has been burnt. But we can change the example to “John blacked out all the words in the book and then mastered it” – in that case the physical book would still exist, but would not be available for mastering.

³² Another case discussed by Asher (Asher (2011):133), which might be thought to cause trouble for our view, is the contrast between “The student mastered every book in the library” and “The student carried off every book in the library”. The worry arises from the claim that even if informational books can be “carried off”, the latter sentence does not have a reading where the quantification ranges over informational books (i.e. it does not have a true reading when the student carried off just one copy from each informational book). In a brief, our response is the sentence does in fact have the relevant reading – see Magidor (2016): 567-8 for a context that makes this reading salient.

be multiply located), that would not help with (57) – as the colour in the demonstrative would do little to pick out *War and Peace* in particular (many other informational books on Jill’s desk would count as red too). Even worse, variants of (57) would require the type book to be both red and not red. (Imagine that Jill’s desk contains a bunch of red books, and just a single book – a copy of *War and Peace* – that is clearly not red, but also does not have a very definite colour. In that case Jane might naturally say “I picked up that book – the one that isn’t red – from the library last week!”). We conclude that allowing for type-books to be red is not the correct way to account for the truth of (57).

Rather, we maintain that this example involves a case of deferred reference. A classic example of deferred reference is Nunberg’s “ham sandwich” case: a waiter says “The ham sandwich left without paying” to communicate the proposition that the costumer who was eating a ham sandwich left without paying.³³ The exact mechanism for deferred reference (as well as the question of whether the resulting proposition is semantically asserted or merely pragmatically communicated) need not be settled here. Our suggestion is that the demonstrative “that red book” in (57) is interpreted in the first instance as picking out the physical copy of *War and Peace*, but then its reference is deferred to the informational-book.

At least one piece of data to support this interpretation is the following. Note that when definite descriptions (rather than complex demonstratives) are used for deferred reference, the standard requirement of uniqueness (in context) are waved with respect the original referent. Thus, for example, if one costumer is in the process of eating five ham sandwiches at once (and is the only costumer eating ham sandwiches) an onlooker can felicitously say “The ham sandwich is sure eating a lot!” – even though there isn’t a single ham sandwich for the description to refer to.³⁴ And interestingly, an analogous case to (57) passes this test: if Jill’s desk contains many red copies of *War and Peace* (and lots of other books with covers in other colours), Jane can felicitously utter:

(58) Hey, I picked up a copy of the red book from the library last week!

V.III Individuating books both informationally and physically?

An example presented in Gotham (2016), one which plays a key role in his discussion, is the following:

(59) Three informative books are heavy.

³³ See Nunberg (1995)

³⁴ One explanation is that the mechanism for deferred reference in this case is that the deference happens already at the level of the predicate embedded within the description: if “ham sandwich” is interpreted as *costumer eating a ham sandwich* and then the definite article is applied to the latter property, and uniqueness will not be violated.

According to Gotham, the truth of this sentence requires there to be “*three books individuated both physically and informationally*” (p.2). Consider the following three situations:

S1: There are three physical volumes: one copy of *War and Peace*, one copy of *Anna Karenina*, and one copy of *The Kreutzer Sonata*.

S2: There are three physical volumes: three copies of *War and Peace*

S3: There is one physical volume which binds together the text of *War and Peace*, *Anna Karenina*, and *The Kreutzer Sonata*.

According to Gotham, (59) is true relative to S1, but false relative to both S2 and S3.

Can our theory account for this claim? To do so, we would have to maintain that the noun “book” applies not only to objects individuated physically and to objects individuated informationally, but also to an additional kind of object that is individuated both informationally and physically.³⁵ Absent any compelling evidence, we are reluctant to accept the existence of such objects.

Rather, we wish to contest that (59) really has the truth-conditions that Gotham attributes to it. (Interestingly, even though accounting for these alleged truth-conditions plays a central role in the development of Gotham’s account of copredication, he says very little to justify the claim that the sentence really has these truth-conditions).

For a start, note that the sentence clearly has at least *some readings* on which (contra Gotham) it is true relative to S2. Imagine that a certain library has a project of dusting each of its (physical) books. Moreover, since the informative books are much more popular among readers than the uninformative ones, they decide to start the project by dusting all the informative (physical) books in the library. Now, in that context, one can appropriately utter the following sentence, in reference to the three copies of *War and Peace*:

(60) Three informative books are heavy, so take extra care as you’re pulling them down from the shelf for dusting.

Are there also readings of the sentence that are true relative to S3? One complication is that Gotham assumes in his discussion that this scenario does not involve three physical books – but that is not obvious (on one perspective, this scenario involves three physical books that are attached to each other). Even if we assume this scenario has only one physical book, there remains the question of whether informational books can be heavy. Arguably they *can* be heavy in cases where all or at least typical instances are stably heavy. Thus if *War and Peace* is a very long book that is typically printed on heavy physical volumes it is arguably heavy (even if it was

³⁵ It is worth noting that we do already have many more than two kinds of individuation conditions for books (we allow for one kind of book where the Hebrew translation and the English translation of *War and Peace* count as the same book and another where it doesn’t). However, such flexibility seems much less objectionable where “type”-like objects are concerned, than in the case of physical books.

once printed in microscopic font in a tiny non-heavy volume). Assuming the three informational books in question are all heavy, (59) can also have a reading which is true relative to S3.³⁶

Thus far, we have argued that (59) has at least *a* reading which (contra Gotham) is true relative to S2 and perhaps also *a* reading which (contra Gotham) is true relative to S3. The question remains, however, whether the sentence also has *an additional* reading which has precisely the truth-conditions that Gotham suggests. The problem is that it is not trivial to test for the presence of such readings. Again, recall that Gotham points to one situation involving informationally and physical distinct books (S1) relative to which the sentence has a true reading; a second situation (S2) involving three physically distinct but informationally identical books relative to which the sentence has a false reading; and a third situation (S3) involving three informationally distinct but physically identical books relative to which the sentence has a false reading. We do not deny that the sentence has readings with precisely these profiles: we can easily get these readings by interpreting “book” as physical in the first and third case, and as informational in the second. The problem, however, is that to get the truth-conditions Gotham postulates we need an additional argument that all of these judgements involve *the very same reading*, and so far nothing has been said to justify this claim.

Perhaps a more sensitive test would be the following. Suppose that Jane’s library contains exactly three physical books – three heavy copies of *War and Peace* and John’s library contains exactly one physical book – a volume of three distinct heavy informational books. We might then test whether the following sentence has a true reading:³⁷

(61) Neither Jane nor John owns three informative books that are heavy.

The sentence and the situation in question are quite convoluted, so it is hard to get any clear judgment here. But for what it’s worth it is far from clear to us that it has a true reading in this situation. We conclude that we do not have sufficient evidence for thinking that (59) has any reading with truth-conditions that require books individuated both informationally and physically.

VI. Comparisons

³⁶ It is also important to note that if no informational books are heavy that will make little difference to our argument. In that case, (59) will have two readings: one where “book” is restricted to physical books and one where the noun is restricted to informational books. Both readings will be false relative to S3, but the sentence would still not have any reading with the truth-conditions that Gotham attributes to it (the physical book reading will sometimes be true relative to S2, and the informational book reading will be false relative to all three situations).

³⁷ Note that this test assumes that informational books can be heavy. If they cannot, then the presence of a true reading of (61) would not support Gotham’s theory, as that reading can be obtained simply by restricting “book” to informational books.

Gotham (2014), (2016) and Asher (2011) develop sophisticated views that contrast with our own view that copredication requires no revisionary semantics. Our final task is to compare their views to ours.

VI.1 Asher (2011)

Asher's account of co-predication proceeds against a background type theory. According to Asher, each lexical item is assigned a (very fine-grained) type. In addition, predicative items place type presuppositions on their arguments. For example, "heavy" presupposes that its argument is a physical object, while "informative" presupposes that its argument is an informational object. When a predicate is applied to an argument, the presuppositions "must either be satisfied by the types of the arguments or accommodated if the predication is to be semantically well-formed" (Asher (2011): 97).

Consider a typical case of co-predication:

(62) The book is heavy but informative.

What type should "book" receive? The two predicates ("heavy" and "informative") place conflicting type presuppositions on their arguments. If "book" belongs to the type informational object ('I') the first predication will crash, and if "book" belongs to the type physical object ('P') the second predication will crash. Asher postulates that the term belongs to *neither* of these types. Rather, the term unambiguously belongs to a special kind of type ("dot type") which we can think of as a kind of product of the two types: $P \cdot I$ (read "P dot I"). Now this in itself does not solve the problem with (62). On the contrary: on this analysis "book" satisfies *neither* the presupposition of "heavy" nor of "informative", so the sentence should suffer from a "double crash".

To address this, Asher postulates a complex semantic and syntactic mechanism, which has the following effect: (62) is interpreted as having a structure whereby "heavy" and "informative" do not apply directly to the subject, but rather to a modulated version of the subject. This modulation is achieved via a functor which takes the dot object (an object of type $P \cdot I$) to one of its "physical aspects" in case of "heavy", and one its "informational aspects" in the case of "informative".

We will not dwell on the syntactic and semantic complexities of this system, but let us pause for a moment to consider the metaphysical commitments involved. It is far from clear what the underlying metaphysics of aspects is. Asher's brief remarks on the topic do little to clarify this issue. With respect to "book" he writes: "Books, for instance, appear to have both a "physical" aspect and an "informational aspect"... Books have a "dual" nature, – two conceptualizations, if you will, that are equally "true of" or "faithful to" the object" (Asher (2011): 130-131). Even more strikingly, he argues that almost every noun has both a "kind aspect" and an "individual aspect" despite claiming that "kinds and individuals are incompatible aspects, in the sense that they have different, even incompatible individuation and identity conditions" (Asher (2011):

131). The upshot seems to be that different aspects are equally “true of” the object, despite being incompatible – a view that seems to lead to a straightforward contradiction.

There is, admittedly, a tradition in metaphysics of reconciling apparently contradictory properties that a single object has, by arguing that it is different *parts* of the object that are the bearers of the relevant properties (see e.g. the temporal parts solution to the problem of change), but Asher is clear he does not want to construe aspects in this way. Here, for example, is what he says with respect to the case of “lunch” (which he takes to similarly belong to a dot-type combining an event type with a food type, a manoeuvre which is intended to account for cases such as (3)): “A lunch object is *wholly* an event (under one aspect) and wholly food (under another aspect). When we speak or think of lunches as food, there is no “other part” of the lunch itself that’s left out and that is an event” (Asher (2011): 150, emphasis in the original).

Things become even worse when we consider sentences such as (1) (“Three books are on the shelf”). As we noted at the outset, in a scenario containing three copies of *Naming and Necessity* the sentence has both a true reading and a false reading. However, since Asher takes “book” to univocally pick out an object of type P-I it is not clear how he can get these readings. To address this, Asher says the following: “We cannot coherently count objects of • types except with respect to one of the constituent type’s counting criterion...So we mustn’t count the physical aspect of a particular book and the informational aspect of a book as two different books. But we can count books as either informational or physical (though not coherently both)...Does this mean that we have in effect two models with two different sets of inhabitants of *book*, depending on which criterion of individuation is chosen? Not necessarily. But it does force us to the position that counting requires a criterion of identity and individuation, which is part of the doctrine of relative identity” (Asher (2011): 157).

In contrast to Asher’s view, our own view postulates perfectly ordinary syntax and semantics for sentences involving co-predication ((62) on our view has exactly the same syntactic structure as “The book is heavy and blue”). Our view also avoids metaphysical commitment to relative identity, as well as an ontology of aspects. So, both linguistically and metaphysically, our view is more conservative than Asher’s, while accounting for all of the relevant data. Of course, our view does take on the commitment that there are no clear categorical constraints on property instantiation, but we have argued at length that this is well-motivated for reasons wholly independent of copredication.

VI.II Gotham (2016)

Gotham’s theory of copredication shares some key elements with ours. Like us, Gotham holds that “book” is unambiguous, and that it designates objects that can instantiate both properties ascribed in a copredication sentence like (2).

(2) Three red books are informative.

However, Gotham’s development of this view differs from ours in a number of important ways.

The differences are explained in part by the fact that Gotham's primary aim is to compositionally derive truth-conditions for (2) on which the books are both physically and informationally distinct. We've already argued that (2) does not have these truth-conditions, so we reject the data on which Gotham's theory is founded.

To generate his desired truth-conditions, Gotham adopts the following theses. First, nouns like "book" designate mereological sums of informational and physical books. Consider, for instance, the copy of *Naming and Necessity* on the shelf. This is not in the extension of "book" for Gotham. Rather, the fusion of that copy, along with the informational book it expresses, is the in the extension of "book". However, Gotham maintains that such mereological sums inherit properties from their parts, so that the sum of the physical book and the informational book inherits the property of being on the shelf, as well as the property of being informative.³⁸ Second, each predicate encodes a criterion of individuation. "Heavy" for instance encodes a physical criterion, while "informative" encodes an informational criterion. Nouns like "book" which designate sums of physical and informational entities encode composite criteria of individuation: physical and informational. Third, counting, and quantification more generally, are sensitive to criteria of individuation.³⁹ We ordinarily think of a counting sentence like "Three Ps Q" as requiring mere non-identity between the Ps; Gotham thinks that they can encode more specific distinctness criteria, e.g. being *physically distinct*. Importantly, two numerically distinct objects may fail to be physically distinct. This happens when two objects have identical physical components but distinct non-physical components.

The reason that Gotham takes quantification to be sensitive to criteria of individuation is that without it, he would generate incorrect truth-conditions. Consider a scenario involving three copies of *Naming and Necessity*. In that scenario there is a *false* reading of (63):⁴⁰

(63) There are three informative books.

For Gotham, "book" designates fusions of physical copies and informational books. There are three of those, and each of them inherits the property of being informative. How, then, can he generate the false reading of (63)? His solution is that quantification, in this case, is sensitive to an *informational* criterion of individuation. Thus the sentence is read as expressing the claim that there is plurality of three books, *that are informationally distinct from each other*, and such

³⁸ We note in passing that Gotham maintains that such mereological sums have *all* the properties of each of their parts (Gotham (2014): 34). But this cannot be correct: the physical book has the property of being purely physical or of *not having an informational part* – properties that the sum clearly does not inherit.

³⁹ Unlike Asher, Gotham does not adopt the view of relative identity. Rather, he builds criteria-sensitivity into the semantics of quantifiers.

⁴⁰ In fact, Gotham maintains that the sentence *only* has a false reading, because he thinks that the use of the predicate "informative" forces us to individuate books informationally. According to us, the sentence has both a true and a false reading, but that still leaves the question of how to generate the false reading.

that each of them is informative. Since in our scenario there is no plurality containing three informationally distinct books, the sentence is false.

To make quantification sensitive to criteria of identity, Gotham needs to substantially complicate his semantics. He associates each predicate with a criteria of individuation, along with its ordinary semantic value, and then revises quantifier semantic-values so that they can compose with these criteria. Transitive verbs must be associated with multiple criteria, one for each argument place. As is clear in his work, pursuing these innovations in a compositional setting requires an enormous amount of ingenuity and complication. By contrast, our view requires none. Given that our view adequately accounts for copredication without such complexity, we should prefer it.

In addition to objecting to its complexity and the data that drives it, there are a number of more specific worries about Gotham's view. Recall our scenario in which there are three copies of *Naming and Necessity* on the shelf (and no other books). As we have argued, there is a true reading of (11):

(11) One book is on the shelf.

The problem, however, is that Gotham's view does not generate this true reading. The reason is that according to Gotham "on the shelf" requires a *physical* criteria of individuation, entailing that whenever we count "books on the shelf" we count physical distinct copies separately. Perhaps Gotham could maintain that "on the shelf" is ambiguous, between one predicate encoding a physical criterion of individuation and one encoding an informational criterion. But apart from generating multiple ambiguities, this solution would run into further problems. Consider the envisioned reading on which (11) is true. In such cases it should have a single witness to its truth, one that we can go on and name or singularly refer to. Thus consider for example the following:

(64) One book is on the shelf. The book is *Naming and Necessity*, and it has been in print continuously since 1980.

For Gotham, the first quantification does not involve a single witness. Rather, its truth involves *three* equally good witnesses (each a sum of one of the three physical copies together with the informational book). Which book, then, are we referring to then, in the second part of (64)?⁴¹ It is hard to see, then, how Gotham can account for such cases.

⁴¹ Furthermore, if the witness is a mereological sum of a physical and informational book, and if (as Gotham's discussion suggests) such sums inherit all of the properties of their physical book parts, then it should be okay to say "One book is on the shelf: it is furthest to the left and has a red cover". After all, there's just one mereological sum of the physical book on the left and the informational book. However, this sentence seems extremely odd.

Another more specific problem is that Gotham cannot easily handle the fact that there is more than one way of individuating books informationally. Recall the scenario on which there are three copies of *Naming and Necessity* on the shelf: two of the French translation and one of the Hebrew translations. As we have argued, in that scenario each of the following has a true reading:

- (65) One book is on the shelf (namely, *Naming and Necessity*)
- (66) Two books are on the shelf (namely, the Hebrew version of *Naming and Necessity* and the French Version).

Can Gotham account for this data? One way for him to try and do so, would be to revise his ontology of books. Thus he could propose that there are three distinct informational books in play in this scenario - *Naming and Necessity* (NN), The Hebrew version of *Naming and Necessity* (HNN), and the French version of the book (FNN). Call the three physical volumes in play V_1 - V_3 . The scenario will then involve the following six composite objects: V_1 +NN, V_2 +NN, V_3 +NN, V_1 +FNN, V_2 +FNN, and V_3 +HNN. Since V_2 +FNN and V_3 +HNN are informationally distinct, then a plurality containing both of them can, on Gotham's theory, be the witness for "At least two books are on the shelf" (assuming an informational criterion of individuation). This, however, already bars (65) (read as an exact counting statement) from being true. Moreover, on this way of expanding his theory, it will also be true (given an informational criterion of individuation) that "At least *three* books are on the shelf", since V_1 +NN, V_2 +FNN, V_3 +HNN are all informationally distinct, so that (66) will be false as well.⁴² Expanding the ontology of books in this way, then, would not allow Gotham to get the right truth-conditions in this case.

An alternative option for Gotham is to acknowledge only the fine-grained informational books HNN and FNN, and then add to his compositional theory *multiple informational criteria* of individuation (one according to which V_1 +HNN and V_3 +FNN count as informationally distinct and one according to which they do not). And of course, the point generalizes beyond

⁴² Can Gotham maintain that one never quantifies over pluralities containing both V_1 +NN and V_2 +FNN because these two objects are too similar to each other? (Cf. our discussion of restriction on counting both red and crimson in the same context in §II). We think he cannot. For recall that Gotham allows us to quantify over pluralities containing both V_1 +NN and V_2 +NN – if pluralities containing objects with *identical* informational components are acceptable, then surely a plurality containing the merely similar informational components (FNN and NN) should be acceptable.

translations to editions, printings, etc. entailing a massive ambiguity in the interpretation of quantifiers.⁴³, ⁴⁴

References

- Armstrong, D.M. 1978. *Universals and Scientific Realism*. Cambridge: Cambridge University Press.
- Asher, N. 2011. *Lexical Meaning in Context*. Cambridge: Cambridge University Press.
- Chierchia, G. 2010. "Mass Nouns, Vagueness, and Semantic Variation." *Synthese* 174: 99-149.
- Chomsky, N. 2000. *New Horizons in the Study of Language and Mind*. Cambridge: Cambridge University Press.
- Collins, J. 2009. "Methodology, Not Metaphysics: Against Semantic Externalism." *Aristotelian Society Supplementary Volume* 83: 53-69.
- Collins, J. *forthcoming*. "The copredication Argument", *forthcoming in Inquiry*.
- Cooper, R. 2007. "Copredication, Dynamic Generalized Quantification and Lexical Innovation by Coercion." In P. Bouillon, L. Danlos, and K. Kanzaki (eds.), *Fourth International Workshop on Generative Approaches to the Lexicon*. Paris, France.
- Gotham, M. 2014. *Copredication, Quantification, and Individuation*. Ph.D. Thesis, University College London.
- Gotham, M. 2016. "Composing Criteria of Individuation in Copredication." *Journal of Semantics* 1-39.
- Kleinschmidt, S. and J. Ross 2012. "Repeatable Artwork Sentences and Generics." Mag Uidhir, C., (ed.), *Art and Abstract Objects*. Oxford: Oxford University Press.
- Kratzer, A. 2016. "Situations in Natural Language Semantics," *The Stanford Encyclopedia of Philosophy* (Spring 2016 Edition), Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/spr2016/entries/situations-semantics/>.
- Lewis, D. 1983. "New Work for a Theory of Universals." *Australasian Journal of Philosophy* 61: 343-77.
- Liebesman, D. 2011. "Simple Generics." *Noûs* 405: 409-442.

⁴³ A further account of co-predication that has some features in common with each of Gotham and Asher is that of Luo (see e.g. Luo (2010), Luo (2012)). Like Asher, Luo postulates a fine-grained system of types, and takes books to belong to a kind of product type: P•I. However, unlike Asher, he takes P•I to be a *sub-type* of each of the types P and I (note that this does not entail that there is a *sub-set* relation holding between the set of individual members falling under these types). Thus even though he takes "informative" to be sortally restricted to objects of type I and "heavy" to be sortally restricted to objects of type P, he can allow both predicates to be true of the same object of type P•I. However, while this view allows for a relatively straightforward account of co-predication it has nothing at all to offer on the issue of counting and individuation - namely, questions such as how to generate both true and false readings of sentences such as "Three books are on the shelf" in situations where three copies of Naming and Necessity are on the shelf. (Cf. Gotham (2014): 122 who shares this criticism.)

⁴⁴ Thanks to Cian Dorr, Paul Elbourne, Jeremy Fantl, John Hawthorne, Kerstin Hoge, Ali Kazmi, Dan Korman, Beau Madison Mount, Milan Ney, Bernhard Nickel, Jeff Pelletier, Bryan Pickel, Byeong-uk Yi, audiences at University of British Columbia, The Society of Exact Philosophy (at the University of Calgary), The University of Edinburgh, and Oxford University, as well as financial support from the Social Sciences and Humanities Research Council of Canada (SSHRC) and the Leverhulme Trust.

- Liebesman, D. 2015. "Predication as Ascription." *Mind* 124: 517-69.
- Liebesman, D. 2016. "Counting as a Type of Measuring." *Philosophers' Imprint* 16(12): 1-25.
- Luo, Z. 2010. "Type-theoretical semantics with coercive subtyping." In Nan Li and David Lutz (eds.), *Semantics and Linguistic Theory (SALT)*. Vol. 20. CLC Publications. Ithaca, NY. 38-56.
- Luo, Z. 2012. "Formal Semantics in Modern Type Theories with Coercive Subtyping." *Linguistics and Philosophy* 35.6: 491-513.
- Magidor, O. 2013. *Category Mistakes*. Oxford: Oxford University Press.
- Magidor, O. 2016. "Response to Abrusán, Shaw, and Elbourne." *Inquiry* 59: 559-586.
- Nunberg, G. 1995. "Transfers of meaning." *Journal of Semantics* 12 (2): 109–132
- Pietroski, P. 2005. "Meaning Before Truth." in *Contextualism in Philosophy*. Preyer and Peters, (eds.). Oxford: Oxford University Press, 2005.
- Pustejovsky, J. 1995. *The Generative Lexicon*. Cambridge, MA: MIT Press.
- Stanley, J. and Szabó, Z. 2000. "On Quantifier Domain Restriction." *Mind and Language* 15: 219-61.
- Walters, L. 2013. "Repeatable Artworks as Created Types." *British Journal of Aesthetics* 53 (4): 461-477