Philos Stud (2021) 178:209–234
https://doi.org/10.1007/s11098-020-01428-9

# Double-counting and the problem of the many

**David Liebesman**[1]

Published online: 13 February 2020


**Abstract** There is a defeasible constraint against double counting. When I count colours, for instance, I can't freely count both a color and its shades. Once we properly grasp this constraint, we can solve the problem of the many. Unlike other solutions, this solution requires us to reject neither our counting judgments, nor the metaphysical principles that seemingly conflict with them. The key is recognizing that the judgments and principles are compatible due to the targeted effects of the defeasible constraint.

**Keywords** Problem of the many · Quantifier domain restriction · Nominal restriction · Double-counting · Semantics of counting · Pragmatics of counting

## 1 The problem

It's a beautiful day for a hike: there's just one cloud in the sky. So you thought. Some attractive metaphysical principles seem to show there are many more. All it takes for there to be a cloud is enough water molecules to be close enough together. There are many such groups in the sky: pick an arbitrary group, subtract one molecule at the edge, and you have another. Different groups constitute different clouds. So the multitude of groups entails a multitude of clouds. Either your judgment is mistaken, or one of the metaphysical principles is. At least that's the

✉ David Liebesman
david.liebesman@ucalgary.ca

[1] Department of Philosophy, University of Calgary, 2500 University Dr. NW, Calgary, AB T2N 1N4, Canada

conventional wisdom about the problem of the many. The following six claims seem jointly inconsistent and independently plausible[1]:

(i) Exactly one cloud is in the sky.
(ii) There are a multitude of cloud candidates (collections of water molecules), such that if any constitutes a cloud then they all constitute clouds.
(iii) If x and y are non-identical cloud candidates, then if they constitute clouds, the clouds they constitute are not identical.
(iv) If there is a cloud in the sky it is constituted by one of the candidates
(v) If x is a cloud in the sky and y is a cloud in the sky and x is not identical to y then at least two clouds are in the sky.
(vi) If at least two clouds are in the sky then it is not the case that there is exactly one cloud in the sky.

Familiar solutions to the problem of the many reject one of these claims.[2] Whether these rejections are plausible is irrelevant to solving the problem, or so I'll argue. How can this be? On my view, (i)–(vi) are consistent. The claims contain context-sensitive terms—common nouns like 'cloud'—and we have independent reason to think that the context-sensitivity is resolved differently in different claims. The result is consistency.[3]

In particular, I will argue that the occurrence of 'cloud' in (i) is governed by a pragmatic principle that rules out double-counting. The result is that if there are a multitude of overlapping clouds, the extension of 'cloud' in (i) is restricted. This principle, however, does not govern 'cloud' in (ii). Since 'cloud' has different extensions in (i) and (ii), we cannot derive a contradiction from (i)–(vi).

The pragmatic principle is not ad hoc: it is independently motivated by a variety of constructions. In Sect. 2 I'll motivate and articulate the principle. In Sect. 3 I'll return to the problem of the many and explain how, given the principle, (i)–(vi) are consistent. My solution is compatible with familiar solutions to the problem. Each denies one or the other of (i)–(vi); I don't deny or affirm any of (i)–(vi), I simply affirm their compatibility. In Sect. 4 I'll respond to objections, and in Sect. 5 I'll mention other potential implications of the pragmatic principle. In an "Appendix", I discuss functional overlap, which had a role in the preceding discussion.

[1] This presentation is adapted from Weatherson (2016).

[2] Unger (1980) denies (i), and rejects the existence of clouds. López de Sa (2014) and Williams and Robert (2006) reject (i) because they think there are a multitude of clouds. McGee and McLaughlin (2000) reject (ii) because they think that it is false under all precisifications, though each candidate constitutes a cloud relative to some precisification. Korman (2015) rejects (ii) because he thinks that just one of the candidates constitutes a cloud, though it is metaphysically indeterminate which. Woods (forthcoming) rejects (ii) because he thinks that there is one maximal candidate. Jones (2015) rejects (iii) because he thinks that a single cloud can be constituted by multiple different collections. Chisolm (1976) and Noonan (1993) reject (v) because they think we count by a relation weaker than identity.

[3] Sattig (2010) also develops a view that takes (i)–(vi) to be consistent. Space precludes a comprehensive discussion of his view, but, as he recognizes, his view is based on a semantic theory that lacks empirical support. (Though he doesn't take this to be a problem.) Given that I take there to be substantial empirical support for my view, it has at least that advantage over Sattig's.

Lewis (1993) inspired the view I'll defend. Like me, he thinks that when entities massively overlap we often disregard their multiplicity. However, Lewis' views about counting compel him to deny either (i) or (vi) (depending on how we interpret him). By contrast, I don't deny either.

## 2 The double-counting constraint

For most of this discussion, I will follow Stanley and Szabó (2001) and Stanley (2002) in taking contextual domain restriction to derive from variables associated with nouns. 'Cloud', as a matter of its context-independent meaning, has all clouds—no matter their relevance—in its extension. 'Cloud' also occurs in sentences with a free domain variable that is then either bound or contextually saturated. In the latter case, the extension of the occurrence is determined by intersecting the contextually specified domain with the context-independent extension.[4] What goes for 'cloud' goes for nouns more generally. For instance, if I truly utter 'Every book is on the shelf', 'book' co-occurs with a free domain variable that is saturated in context—e.g. by the domain of all things in my office. The context-independent extension of 'book' is then intersected with the contextually provided domain, and the sentence is true just in case every object in the resultant set—the set of books in my office—is on the shelf. Though this view of domain restriction is plausible and well-motivated, it is inessential to my argument. In Sect. 4.3 I'll discuss other possible views of domain restriction and demonstrate how to generalize my argument.

Assume, then, that nouns co-occur with free variables for contextually-provided domains. Here's a difficult question: what determines which domain saturates a domain variable in an arbitrary context? I won't pretend to fully answer that question, but, in the remainder of this section, I will identify a constraint.

### 2.1 Determinates and determinables

You are painting using two different cans: one filled with maroon paint and the other with crimson. You dip your brush in the crimson can and paint a swatch on the previously blank canvas, then you do the same with paint from the maroon can. There is a true reading of (1). (This case is adapted from Liebesman and Magidor (2017).)

(1) Two colours are on the canvas.

There's also a true reading of (2). To see this just imagine following (2) with 'red'.

(2) One colour is on the canvas.

Crimson and maroon are colours, and they witness the truth of (1). Red is also a colour, and it witnesses the truth of (2). What's surprising is that it is extremely

[4] I'm simplifying the view and omitting the functional variable that Stanley and Szabó posit to explain binding readings.

difficult (though perhaps not impossible) to get a true reading of (3), with red, crimson and maroon as the intended witnesses.

(3) Three colours are on the canvas.

(3) is difficult to read as true because it is usually illegitimate to count red as well as maroon or crimson, the latter of which are shades of the former. Counting a colour and its shades simultaneously is an illegitimate form of double-counting. Shades are determinates of colour determinables, and just as it is illegitimate to simultaneously count a colour and one of its shades, it is usually illegitimate to simultaneously count any determinable along with its determinates. The observation generalizes. Here is another case:

Imagine that I paint a square onto the canvas. In that case, there's a true reading of (4), as witnessed by the shape *square*, but there's no (readily available) true reading of (5) as witnessed by the shapes *square and rectangle*.

(4) One shape is on the canvas.
(5) Two shapes are on the canvas.

Once we recognize the restriction on simultaneously counting a determinable and its determinates, we can generate similar cases at will. Our next step is to articulate a principle that captures this constraint on counting. Before doing that, though, we'll need to make some observations about these cases.[5]

## 2.2 Four observations and a principle

**Observation 1:** the mere fact that some entities are in the extension of a term does not entail that we can easily count all of those entities at once.

The fact that red, maroon, and crimson are all in the extension of 'colour' does not entail that there is an available true reading of (3), at least without substantial contextual background. This is important because it divorces theses about whether something is a P, where P is an arbitrary property, from our judgments about how many Ps there are. As we'll see, this divorce has number of implications.

**Observation 2:** merely avoiding double-counting is not enough to make a domain salient.

To see this, consider (6).

(6) Four colours are on the canvas.

Imagine that our paint exhibits minor colour variation. Given this variation, we could individuate colours more finely than maroon, e.g. light and dark maroon, and light and dark crimson. These colours do not overlap in any sense—none is a

[5] One worry is that these cases exemplify ambiguity in the restricting noun, rather than restriction on its domain. To see that an ambiguity view is misguided note two things: (1) such a view would lead to an infinite number of senses for terms like 'colour', given that there are an infinite number of levels of colour-individuation, (2) there are occurrences of 'colour' in which it must have both red and maroon in its extension (e.g. 'Red and maroon are both colours'.) Cf. Liebesman and Magidor (2017).

determinable or shade of another. Nonetheless, without contextual supplementation, we cannot get a true reading of (6) in the envisioned context.

The problem is that there are a large number of non-double-counting domains for ‘colour’. {red, green, blue, ...} is one (the ellipses stands for the completion of the list of colours that are at the same level of individuation as red, green, and blue), as is {crimson, maroon, ...}, but so is {light crimson, dark crimson, light maroon, dark maroon, ...}. And, of course, there are an infinitude of additional non-double-counting domains corresponding to the infinitude of possible levels of colour individuation. What is it, then, that makes a non-double-counting domain eligible to saturate a free domain variable? Using the term as a placeholder, we’ll call this additional feature ‘salience’.[6]

What goes for domain variables goes for variables more generally. The mere fact that the Pope is a man does not give rise to a reading on which ‘he’ refers to the Pope, for any arbitrary sentence containing ‘he’ in any arbitrary context. To generate that reading, we must somehow distinguish the Pope from the multitude of possible referents.[7]

What determines salience? There are a multitude of factors, and salience is the result of complex interaction between them. One factor is the general social environment, including linguistic environment. English contains oft-used lexical items that designate red, blue, green, etc. and this in and of itself may be enough to make {red, blue, green, ...} a salient domain.[8] In addition to general environmental features, specific contextual features will affect salience. If I explicitly refer to light and dark maroon, this may be enough to make them salient. And, of course, non-linguistic environment may raise domains to salience. If we clearly separate light and dark maroon on the canvas in a way that is obvious to any ordinary observer, we can expect them to compose a salient domain.

I won’t heroically tackle a theory of salience. The only crucial point for my purposes is that not double counting does not suffice.

> **Observation 3:** sufficient contextual supplementation may be able to override the prohibition on double-counting.

Consider the following scenario from Liebesman and Magidor (2017: 141):

> 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;

[6] This technical notion of salience differs from our intuitive notion. As Kratzer (2005) stresses, there are intuitively salient domains that are nonetheless not eligible to be the values of domain variables.

[7] I’ll speak of domains as themselves double-counting. This is shorthand for the claim that, a domain d, relative to a context c and noun n is such that if n’s domain variable were saturated with d in c, that would suffice for double-counting.

[8] Other languages contain other colour words and, given that, other domains will be salient. Russian, for instance, contains distinct lexical items for light and dark blue. Some have argued that this gives rise to processing differences: see Winawer et al. (2007).

Anger; Happiness'. In this context, it seems acceptable to say 'Four colours are on the list: red, crimson, scarlet, and maroon'.

Though some find it strained, most can get a true reading of the counting sentence, where 'colours' ranges over both red and some of its determinates. This shows that we shouldn't build a double-counting restriction directly into our semantics, but we should rather find a defeasible pragmatic constraint on the saturation of domain variables.[9]

If we include linguistic environment as part of context, then we'd predict that some linguistic environments may themselves override the prohibition on double-counting. Consider (7):

(7) Somtimes when you have two colours, one is a shade of the other.

(7) has a true reading. To force it, imagine following the sentence with 'e.g. red and maroon'. That shows that 'colours', even when it restricts number words, must sometimes be allowed to range over both determinates and their determinables. There are likely a variety of different sorts of predicates that enable us to override the prohibition on double-counting. I won't generalize about the exceptions here. One reason is that there may not be true generalizations about exceptions. Another is that my reasoning about the problem of the many will only require focusing on particular predicates.

**Observation 4:** the constraint on double-counting does not affect all occurrences of nouns.

To illustrate this, consider several different occurrences of 'colour':

(8) Red and maroon are both colours.
(9) Red is a colour, as is maroon. (In fact, the latter is a shade of the former.)
(10) Every shade of a colour is also a colour.
(11) Every combination of hue, saturation, and brightness determines a colour. (Some are more specific versions of others.)
(12) I picked a colour, as did my friend. (In fact, I picked red, while he—with his more specific tastes—picked maroon.)

(8) and (9) both have salient true readings. If 'colour' couldn't designate both determinates and their determinables in these sentences, then they would lack those readings. The truth of (8) would be straightforwardly ruled out. One may try to explain this by arguing that the constraint on double-counting is overridden by explicit mention of both red and maroon before the predication. However, this explanation does not apply to (9), which contains an elided occurrence of 'is a colour'. The first occurrence of 'is a colour' in (9) must designate both red and

9 Baxter (1988) and Cotnoir (2013), by contrast, build a double-counting constraint directly into their semantics.

maroon, or when it is copied it wouldn't then guarantee the truth of the second clause.[10]

Similarly, the most salient reading of (10) is true. However, if double-counting were prohibited in (10), then the second occurrence of 'colour' would have to range over a domain that didn't include a colour as well as any of its shades, and (10) would have only a false reading. The upshot of (8), (9), and (10), is that at least some occurrences of 'colour' in predicate position (either in singular predications, or as the matrix predicate of a quantificational sentence) designate both determinables and their determinates.

In (11) and (12), 'a colour' appears as the object term of a transitive verb. Both sentences have salient true readings that require 'colour' to range over both determinates and their determinables. In both cases, I've added a parenthetical remark to force these readings. If the domain of 'colour' in (11) couldn't contain a determinable and its determinate then 'colour' couldn't range over two colours such that one is a more specific version of (i.e. a determinate of) the other. If the domain of 'colour' in (12) couldn't contain a determinable and its determinate, then the sentence wouldn't be true in a scenario in which I picked red and my friend picked maroon. The upshot of (11) and (12) is that at least some occurrences of 'colour' as part of an indefinite that is the object of a transitive verb designates both a determinable and one of its determinates.

Given that the double-counting prohibition seems to most clearly arise when nouns restrict quantifiers, our pragmatic principle will generalize only about such occurrences. This is not to claim that a double-counting prohibition never affects other sorts of occurrences. Rather, we merely aim to provide a sufficient condition for the prohibition, and, as such, we'll focus on the cases in which it uncontroversially arises.

Taking these observations into account, we can articulate a prima facie plausible principle about domain restriction:

> **NDC**: when saturating a free domain variable of a noun that restricts a quantifier, speakers determine a salient domain that contains no double-counting.

NDC should be read as a exception-tolerating generalization, akin to a ceteris paribus law. Compare it to a similar principle we may formulate about 'he': when saturating a free occurrence of 'he', speakers determine the most salient male. This is a true exception-tolerating generalization. It captures a fact about the determination of values for 'he' while being compatible with contexts in which we use 'he' to refer to inanimate objects. Of course, ultimately we'd like to refine NDC as well as incorporate it into a comprehensive pragmatic theory. For our purposes, though, this version will suffice.

What licences us to move from our aforementioned cases to a general principle? The fact that, as already stated, the phenomenon is productive. Once we note that, as a (defeasible) rule, we can't count a colour and its shades, we can see that this applies to any determinable/determinate pair. That NDC is a general principle

[10] A rejoinder: when the predicate is copied, it is the free variable that is copied, not its saturation. A response: on this view, the occurrence of 'as is' in (9) would be very odd, as we'd be ascribing different properties to red and maroon.

simply reflects the fact that the phenomenon that motivates it is generalizable. One hedge: ultimately, NDC will likely have to be refined in order to account for the full diversity of counting data as well as to be incorporated into more general pragmatic and metasemantic theories. As I'll emphasize in Sect. 3.1, the fact that NDC is merely a plausible first pass will not undermine my reasoning.

NDC raises an important question: what counts as double-counting? We've already seen that a domain that contains a determinable and its determinate contains double counting. In the next two subsections I'll consider two other ways to double count. First, though, I will consider an objection.

The objection is that NDC is too strong. To use an example from Lewis (1993), imagine that you draw two diagonals in a square. There may be a true reading of (13), which is in prima facie conflict with NDP, as it requires us to count overlapping triangles.

(13) Exactly eight triangles are in the figure.

Recall, first, that NDC is presented as a defeasible pragmatic principle. So, the most obvious thing to say about such counterexamples is that, somehow, NDC is being overridden. In this particular example, in fact, it seems much harder to get a true reading of (13), than it does to get a true reading of (14), an observation which supports the existence of a defeasible principle like NDC.

(14) Exactly four triangles are in the figure.

## 2.3 Entities and their parts

Kratzer (2012: 112) provides us with *Dialogue with a lunatic* (her title):

> Lunatic: What did you do yesterday evening?
> Paula: The only thing I did yesterday evening was paint this still life over there.
> Lunatic: That's not true. You also painted these apples, and you also painted these bananas. Hence painting this still life was not the only thing you did yesterday evening.

The source of lunacy is the fact that Lunatic takes 'things' to range over both the event of painting, and the events of painting some particular apples, and some particular bananas, where the latter two events are parts of the former.[11]

What goes for events goes for objects more generally: any domain with an object and its parts will violate the double-counting constraint, and not just any way to arbitrarily satisfy the double-counting constraint will yield an acceptable reading.

[11] Kratzer contrasts the lunatic with the pedant. The latter responds to an restricted quantificational claim by denying it and asserting a less restricted counterpart. The contrast is that the pedant seems annoying, but within his rights as a speaker, while the lunatic is not.

Consider a room with one wooden table and no other objects.[12] If we focus on the table itself, we can get a true reading of (15). If we focus on the top and legs of the table separately, we can get a true reading of (16). However, we cannot get a true reading of (17), as witnessed by the table itself, and its five components—at least not without substantial contextual supplementation. The upshot is that another way to violate the double-counting constraint is to count both objects and their parts. As in the determinate/determinable case, this phenomenon generalizes.

(15) One wooden thing is in the room.
(16) Five wooden things are in the room.
(17) Six wooden things are in the room.

That this is an instance of the same phenomenon exhibited by the colours on the canvas case can be seen by revisiting our observations. Our first observation about the colour case was that the mere fact that some entities are in the extension of a noun doesn't entail that we can count them at once–at least not without contextual supplementation. This is true of the wooden thing example as well. The mere fact that the table and its parts are all wooden things doesn't mean we can (easily) count them all at once.

Our second observation was that merely avoiding double-counting does not make a domain eligible to saturate a free domain variable. Imagine that I arbitrarily divide the table into eight non-overlapping parts. If this domain were eligible without contextual supplementation, we should be able to get a true reading of (18), however, we cannot. However, if we explicitly raise the eight non-overlapping parts to salience, then the domain they constitute will itself be salient and there will be an available true reading of (18).

(18) Eight wooden things are in the room.[13]

Our third observation was that contextual supplementation may override the restriction on double-counting. Imagine that we explicitly mention the table as well as its parts, and make it clear that we wish to discuss them all at once. In that case, we can generate readings of the counting sentences on which 'wooden thing' ranges over both.

Our fourth observation was that the restriction on double-counting doesn't extend to all occurrences of the term. To see this, we can consider other occurrences of

[12] This example comes from Liebesman and Magidor (2017). Korman (2015) uses an almost identical example.

[13] There is a tempting metaphysical rejoinder to this examples: to argue that in the wooden thing case the arbitrary parts are not wooden things because they don't exist, perhaps because arbitrary undetached parts do not exist more generally. (See van Inwagen (1981) for an early influential defense.) Tempting as it is, this rejoinder is unconvincing. Imagine that each of the table's five parts were constructed from two pieces of wood glued together, though the gluing was seamless. In that case, it is uncontroversial that the ten parts exist, though there is still not a true reading of (19), at least without additional contextual supplementation. Even if we cannot arbitrarily decompose objects/events, some decompositions will still exist. We can then generate examples using these decompositions.

(19) Ten wooden things are in the room.

‘wooden thing’. (20) and (21) have intuitively true readings, and can’t have such readings without ‘wooden things’ containing both the table and (at least some of) its parts.

(20) The table and its top are both wooden things.
(21) If something is a wooden thing, then any large enough part of it is as well.

## 2.4 Overlapping entities

Our charity is delivering clothes to the needy. One box contains fifty identical shirts, and the other contains fifty identical pairs of pants. (22) is true, but (23) doesn’t have a true reading, at least not without enriching the context.[14]

(22) Fifty outfits are being delivered.
(23) Two thousand five hundred outfits are being delivered.

Without appeal to a restriction on double-counting, the lack of a true reading of (23) is odd. After all, there are two thousand five hundred distinct pairs of shirts and pants, given that we have fifty different shirts and fifty different pairs of pants. These outfits will overlap massively–any single shirt or pair of paints will figure in myriad different pairs. If, however, we restrict to non-overlapping pairs, then the result is fifty.[15]

As Krifka (2009) emphasizes, this sort of example can be generalized, just as in our other two examples of double-counting. To give just one more example, if we have produced three copies each of each volume of a two-volume book, then we’ve produced three books, not nine.

Should we conclude from the outfit and multi-volume book examples that overlap is another way to violate the double-counting constraint? Matters are not so simple, especially if we take overlap to consist merely in part-sharing. After all, imagine that my closet contains two very different pairs of pants and two very different shirts. Furthermore, I’ve carefully chosen my shirts and pants such that each shirt can pair with each pair of paints, creating a drastically different, but nonetheless stylish, outfit (this example is adapted from Krifka (2009)). In this scenario, (24) has a clearly true reading.

[14] This case, and the two-volume book case, are adapted from cases discussed extensively by Krifka (2009).

[15] There’s an important feature of this case that doesn’t apply to the others: there isn’t a unique set of shirt/pants pairs that witnesses the truth of (22). Rather, it seems that any set of non-overlapping pairs will work just as well as any other. We can make sense of this linguistically by taking the domain variable associated with ‘outfits’ to be indeterminate between any domain that doesn’t double-count. Any acceptable resolution of the indeterminacy ensures the truth of (22).

(24) Four outfits are in my closet.

Importantly, I can count each shirt/pants pair, even though each pair will overlap with another. This suggests that it is not mere part-sharing that leads to a violation of the double-counting constraint in (23)

Lewis (1993) argues that we can group multiple overlapping entities together in a count. For reasons similar to those I just gave, Sutton (2015) argues that we shouldn't understand Lewis' notion of overlap in terms of part-sharing. Here's one of her examples (2015: 47): imagine a pair of houses that share a wall and were built by a wall-obsessive, who makes the shared wall 95% of the total space of each house. Even though these houses massively overlap, we still count them as two. Sutton proposes that we group multiple objects together on the basis of functional overlap. The idea is that the function of a house is to shelter, and two houses perform different sheltering functions. Therefore, we must count them as two, in spite of the fact that they share 95% of their total space. This idea applies to both outfits cases. In the charity case, the function of outfits is to *clothe*, and the 2500 pairs cannot clothe 2500 people. In the closet case, however, we've shifted the function to *providing wardrobe diversity*, and the multiple outfits can perform this function, even if they share components.

Functional overlap, then, is another way to violate the double-counting constraint. Functional overlap, however, brings with it some mysteries. To see this, note that we can't simply understand functional overlap in terms of parts of functions. After all, it is hardly a forgone conclusion that functions have parts at all. In the appendix, I attempt to explicate the notion of functional overlap. However, the purposes of our central argument, we'll take functional overlap to be whatever notion explains our judgments in Krifka's and Sutton's cases.

## 3 The problem of the many

On the readings intended, claims (i)–(vi) are compatible. In particular, the occurrence of 'cloud' in (i) ranges over a domain containing just a single cloud. 'Cloud' in (ii) is not so-restricted. The result is that the occurrences of 'cloud' range over different entities. As I'll demonstrate in Sect. 3.3, this blocks the derivation of a contradiction from (i)–(vi).

### 3.1 'Cloud' in (i) is restricted by NDC

i. Exactly one cloud is in the sky.

In (i) 'Cloud' restricts 'exactly one', a quantifier, so NDC, as stated, applies. It follows that the domain variable that co-occurs with 'cloud' in (i) must be saturated by a salient domain that does not double-count. This isn't surprising. (i) has almost the exact same linguistic structure as sentences like (1). As I emphasized, though, NDC can be overridden by linguistic context. I'll argue that it isn't.

Recall from Sect. 2.2 that some predicates can give rise to linguistic contexts in which NDC is overridden. To hold that 'cloud' is restricted in (i), we must first establish that ‘is in the sky’ is not such a predicate. To see that it isn’t, imagine that our outfits for charity are on a plane to their destination. (25) is then true, while (26) is false.

(25) Fifty outfits are in the sky.
(26) minsky 2,500 outfits are in the sky.

Similarly, imagine that I take our wooden desk and toss it in the sky. (27) and (28) have true readings, while (29) doesn’t have a true reading where it is witnessed by the entire desk and its five parts.

(27) One wooden thing is in the sky.
(28) Five wooden things are in the sky.
(29) Six wooden things are in the sky.

The fact that we can so-modify our earlier examples of the double-counting prohibition demonstrates that the predicate ‘in the sky’ does not override the NDC.

In sum, NDC, as stated, applies to ‘cloud’ in (i). This is further bolstered by the fact that (i) is perfectly analogous the sentences that motivate it, and that the predicates in (i) do not override NDC. Of course, the fact that NDC applies to ‘cloud’ in (i) does not mean that NDC cannot be overridden with sufficient effort. Quite the contrary: I emphasized in Sect. 2.2 that NDC can be overriden. I’ll return to this in Sect. 4.1.

### 3.2 ‘Cloud’ in (ii) is not restricted by NDC

The fact that ‘cloud’ in (i) is restricted in accordance with NDC does nothing to demonstrate the compatibility of (i)–(vi) unless ‘cloud’ is not restricted in the same way in at least one of the other premises.[16] I’ll now turn to arguing that, at least in (ii), ‘cloud’ is not restricted in the same way. To stress again, I will not argue that (ii) is true. Rather, I’ll merely argue that (an occurrence of) ‘cloud’ in it is not restricted in accordance with NDC.

> (ii) There are a multitude of cloud candidates (collections of water molecules), such that if any constitutes a cloud then they all constitute clouds.

‘Cloud’ has three occurrences in (ii): as part of the compound ‘cloud candidates’, as part of the indefinite ‘a cloud’, and in its plural form ‘clouds’. The first and second occurrences are irrelevant to the issue at hand: the first could be easily replaced with another term, and the second need only designate a single cloud in order to play its argumentative role. I’ll focus on the third.

One of the observations in Sect. 2.2 was that not all occurrences of nouns are governed by a constraint against double-counting. This is why the articulation of NDC covers only nouns as they occur in quantificational restrictor position. ‘Clouds’ in (ii) occurs not as a quantificational restrictor, but as a bare plural complement of the transitive verb ‘constitute’. We have independent reason to think

[16] Korman (2015: 221) makes this point forcefully.

that nouns in such contexts aren't naturally read as restricted in accordance with NDC. Consider (30) and (31).

(30) All shirt/pants pairs constitute outfits.
(31) All collections of four legs and a top constitute tables.

Return to our original outfit scenario in which we're providing 50 outfits to charity by providing 50 pairs of pants and 50 shirts. The occurrence of 'outfits' in (30) designates all shirt/pants pairs (or, rather, what they constitute), even when they overlap. Why? The salient reading of (30) is true. However, it wouldn't be if 'outfits' were restricted in accordance with NDC. In that case 'outfits' would, on any acceptable restriction, designate a set of non-overlapping outfits and many of the shirt/pants pairs would not be in that set. Note that appeal to indeterminacy is of no help here. Even if 'outfits' were indeterminate between a number of NDC-satisfying domains, it would be (at best) indeterminate whether some arbitrarily chosen pairs constitute outfits and (30) would be (at best) indeterminate, rather than true. All of this reasoning carries over to (31), where we imagine that we are in a furniture factory with table-parts scattered about.

These examples show that the cases we've discussed, bare plural complements of 'constitute' are not restricted in accordance with NDC. This is hardly an isolated phenomenon. In general, it appears that nouns in their occurrences as complements of transitive verbs are not so-restricted. We've already, in effect, illustrated this with (11) and (12).

(11) Every combination of hue, saturation, and brightness determines a colour. (Some are more specific versions of others.)
(12) I picked a colour, as did my friend. (In fact, I picked red, while he—with his more specific tastes–picked maroon.)

None of this should be surprising. NDC, as its name implies, is a principle that governs counting. As such, it is predictable that it applies to nouns as they occur as quantificational restrictors, and not in other sorts of occurrences. Indeed, this is what our investigations have revealed. It is not just transitive verb complement occurrences of nouns that aren't so-restricted, but we also earlier established that predicative occurrences of nouns aren't so restricted.

So, a strong reason to think that 'clouds' in (ii) is not restricted in accordance with NDC is that, in general, transitive verb complement occurrences of nouns aren't so-restricted. There is another more specific reason. Consider this close analog to (ii):

(32) There are a multitude of outfit candidates (pairs of shirts and pants), such that if any constitutes an outfit then they all constitute outfits.

(32) strikes many as clearly true. In fact, amounts to little more of the conjunction of (30) with the claim that there are, in fact, many shirt/pants pairs. However, insofar as it is taken to be true, then 'outfits' as it occurs in (32) cannot be restricted. After all, if it is so restricted then the conditional in (32) will be false (this follows from the same reasoning that established that 'outfits' cannot be restricted in (30)).

The argument, then, is that (32) is so similar to (ii) that if ‘outfits’ in (32) is not restricted in accordance with NDC, then neither is ‘clouds’ in (ii).

In sum, I’ve given two reasons that ‘clouds’ in (ii) is not restricted in accordance with NDC. The first is that it is not the sort of occurrence of a noun that is usually restricted, and the second is that in an extremely similar sentence the corresponding noun is not restricted.

Recall, also, that I have regarded NDC as a plausible first-pass principle but I have allowed that it will likely ultimately have to be refined. We can now see why that doesn’t weaken my reasoning. The ultimate evidence for a double-counting restriction on ‘cloud’ in (i), and a lack thereof in (ii) comes not from NDC itself but from the data that motivated NDC. In particular, I’ve drawn analagies and disanalogies between (i) and (ii) and other counting sentences in order to make the case that (i) and (ii) have the readings I allege.

### 3.3 Blocking the contradiction

That NDC applies to ‘cloud’ in (i), and doesn’t apply to ‘cloud’ in (ii), doesn’t yet show us much. After all, NDC may apply but have no effect. To block the contradiction we need to show that ‘cloud’ has different extensions in the different claims. Given that I am merely aiming to establish that (i)–(vi) are compatible, rather than true, I’ll assume that (ii)–(vi) are true and show how they are compabtible with (i).

Begin with the occurrence of ‘cloud’ in (ii). If (ii)–(vi) are true it follows that the extension of ‘cloud’ contains a multitude of potentially overlapping entities in *i*. After all, *i* contains a multitude of cloud candidates, each of which constitutes a numerically distinct cloud. Given that the occurrence of ‘cloud’ in (ii) is not restricted by NDC, we should expect that ‘cloud’ has in its extension all of the same things in the context-independent extension of ‘cloud’.

What about ‘cloud’ in (i)? NDC applies to it, so, if the context-independent extension of ‘cloud’ double-counts, then ‘cloud’ in (i) must be restricted. In fact, the context-independent extension of ‘cloud’ does double-count (assuming (ii)–(vi) are true) in at least two ways: by containing entities and their parts, and by functional overlap.

Consider some cloud-candidate C that contains molecules $\{m_1...m_n\}$, and a candidate C′ that contains molecules $\{m_1...m_{n-1}\}$. By (ii)–(vi), both C and C′ constitute clouds. On the plausible assumption, that C′ is a part of C, then the domain double-counts by containing an entity and its parts. In fact, the plausible assumption is likely inessential. Reconsider our *wooden thing* case. Even if, in a mereological surprise, the left half the the table is not a part of it, it would remain illegitimate to count both.

What about functional overlap? We’re being purposely non-committal about this notion (though see the “Appendix”). However, no matter our understanding of functional overlap, it is highly plausible that the many entities in the extension of ‘cloud’ (assuming (ii)–(vi)) functionally overlap. We can see this in two ways. First, any functional I can conjure for clouds ensures that they overlap. Consider, for example, the function of *producing rain*. The entities in the extension of ‘cloud’ overlap enormously! The total amount of rain they can produce is certainly not the sum total of the rain they can produce individually. After all, they share many water molecules. Reconsidering our earlier cases makes this point vivid. We couldn't count overlapping outfits twice, because overlapping outfits can't clothe more people. We could count the overlapping houses twice, because, despite sharing a massive wall, the houses have independent sheltering functions. Second, note that it is ubiquitous in discussions of the problem of the many to stipulate that the many clouds (should they exit) exhibit only minute differences. A natural way to understand minute differences is precisely in terms of the functions of the clouds: C and C′ differ only minutely because they are almost exactly the same in terms of all of the cloud functions we normally care about.

With our observation that NDC applies to 'cloud' in (i), and this application is non-vacous, we can now block derivation of a contradiction from (i)–(vi).Here is how one would try to elicit a contradiction:

1. Exactly one cloud is in the sky. ((i))
2. $C_1$ is a cloud in the sky. ((i), existential instantiation)
3. $C_1$ is constituted by $CC_1$ ((iv), existential instantiation)
4. If $CC_1$ constitutes a cloud, then all candidates do. ((ii), universal instantiation)
5. $CC_2$ constitues a cloud, $C_2$. ((ii), 2, existential instantiation)
6. $C_1$ is not identical to $C_2$. (iii, 5, 3, non-identity of $CC_1$ and $CC_2$)
7. There are two clouds in the sky (6, v)
8. It is not the case that there is exactly one cloud in the sky (7, vi).

Given that 'cloud' in (i) is restricted in accordance with NDC, and 'clouds' in (ii) is not, then we cannot elicit a contradiction in this way. Either (8) does not contradict (1), or this reasoning is flawed along the way. Why would (8) not contradict (1)? If 'cloud' is restricted in (1) but not in (8), then there is no contradiction: after all, there may be exactly one entity in the restricted extension of 'cloud' even if there is more than one entity in the unrestricted extension. Insofar as 8 is derived from earlier steps, then it is plausible to think that 'cloud' in (8) is meant to be unrestricted. If cloud in (8) is restricted in accordance with NDC, then it won't follow from earlier steps. At least the occurrence of 'clouds' in (4) must be unrestricted (as it is supported by (ii)). If that's right then any shift along the way to a restricted reading of 'clouds' will block the reasoning. For instance, if 'clouds' is restricted in 7, it won't follow from 6 and v. However you slice it, there is no contradiction: either (8) simply doesn't contradict (1) or the reasoning deriving the contradiction stumbles when it illicitly moves from an unrestricted reading of 'clouds' to a restricted one, or vice-versa.

# 4 Objections

## 4.1 Unrestricting

On the view defended, ‘cloud’ in (i) is restricted by NDC, a defeasible pragmatic principle. An objector my seize on this defeasibility. In principle, the objector insists, we should be able to defeat NDC and generate a reading on which (i) is false. However, the objector continues, try as they might they cannot modify the context such that they intuit that (i) is false. To put it more bluntly: it seems obvious that there just is one cloud in the sky and no context hocus-pocus changes that.

First, a preliminary point. As we learned in the canvas case, it is often very hard to override NDC. So, the mere fact that it is difficult to generate a case on which (i) is intuitively false does not show that ‘cloud’ in (i) is not governed by NDC.

Second, a more substantive point. I do not predict the following: *if NDC were overridden, we would intuit (i) to be false*. After all, it could be that another of (ii)–(vi) is false for whatever reason, and there is just one entity in the extension of ‘cloud’. In such a case, even if NDC is lifted it would make no difference to the truth-value of (i). Rather, I make the more subtle prediction: *If NDC were overridden, and (ii)–(vi) were true, then we will intuit that (i) is false*.

This more subtle prediction is hard to test. Many theorists will hold that one of (ii)–(vi) is impossible. If that’s correct, then we would have to bring in our favoured view of counterpossible conditionals to evaluate this prediction.[17]

On the other hand, if (ii)–(vi) are compossible, then we need merely consider the closest possible world in which they are. In this world, the extension of ‘cloud’ has more than member. Insofar as I can process this counterfactual, it seems true. So, either way, the more subtle prediction seems accurate. Or, more cautiously, we have no good reason to think that it is inaccurate.

Though nobody has defended my particular approach to the problem of the many, Lewis (1993) proposes a similar approach. On the relevant interpretation of Lewis,[18] he takes (i) to be strictly speaking/literally false, but thinks we can use it to express a truth. In many ordinary contexts, Lewis contends, mere non-identity of x and y does not suffice for them to count as two: they must also not substantially overlap.

This solution contrasts with the one I’ve advanced.[19] I do not hold that (i) is strictly speaking/literally false. In fact, I hold no view whatsoever as to its truth value. My only contention is that its free domain variable is governed by NDC. (How could (i) be false given that the variable is governed by NDC? Easy: if there’s a salient non-singleton domain that satisfies NDC, or if Unger (1980) is correct, and there are no clouds.) Nonetheless, there is a key similarity: in both cases, the plausibility of (i) is sourced to contextual effects.

[17] In fact, Korman (2015: 202) claims that relevantly similar counterpossibles are true.

[18] This is not uncontroversial, see Weatherson (2016) and López de Sa (2014) for competing interpretations of Lewis (1993).

[19] Another contrast: Lewis and I motivate our counting claims with different cases and, as I argue in Liebesman (2015), I don’t think his cases work.

Hudson (2001) is skeptical about this reliance on context. He argues that while we can, of course, use (i) to express non-discriminating counts *a la* Lewis, we also know how to count strictly. Furthermore, the problem of the many arises when we count strictly as well, so invoking non-strict counting is of no help. We can attempt to generalize Hudson's objection to my view. The Hudson-inspired objection, then, is that the problem of the many arises even when NDC is overidden.

Let's get clearer on the objection. The idea is that the problem of the many arises even when NDC is overridden. The problem of the many arises if claims (i)–(vi) are jointly incompatible, as well as independently plausible. So, the idea behind the Hudson-inspired objection is that overriding NDC ensures that the claims are incompatible, while retaining their independent appeal.

Our discussion thus far already gives us the tools to undermine this objection. If NDC is overridden, we have no reason to think that (i)–(vi) are compatible. The question then becomes whether (i)–(vi) are independently plausible when NDC is overridden. This, in fact, just reduces to the question of whether the following is true: if NDC were overridden and (ii)–(vi) were independently plausible, then (i) would be implausible. I am committed to endorsing this conditional while the Hudson-inspired objector must reject it. This conditional should sound familiar: it is just our earlier more subtle prediction with truth/falsity swapped for independent plausibility/implausibility. For all of the same reasons rehearsed above, we should think that this conditional is true, so the objection is rebutted.

### 4.2 Generalizing

I've focussed solely on clouds. The problem of the many is bigger than that. One can think of the problem of the many as schematic, where (i)–(vi) are just one instance and other instances are given by replacing 'cloud' with another count noun. The generality of the problem gives rise to a worry: even if my solution is plausible for this particular instance of the problem, why should we think that it will generalize to other instances?

To see that my solution generalizes, we merely need to reconsider my argument that NDC is non-vacuous for 'cloud' in (i), assuming the truth of (ii)–(vi). The general idea was that if (ii)–(vi) are true, the context-independent extension of 'cloud' contains many entities that overlap both mereologically and functionally. Importantly, none of this reasoning had anything to do with the specific features of clouds. In fact, the exact same reasoning will generalize to all material object kinds, which are just the kinds that give rise to instances of the problem of the many. Why will the reasoning generalize? Consider, first, functional overlap. Whenever two material objects don't exhibit functional overlap then we will no longer take (i) to be plausible (cf. Sutton's wall-sharing houses). In fact, most discussed instances of the problem of the many are explicitly designed such that the would-be instances exhibit functional overlap. Second, consider the double-counting due to a thing and its parts. Again, all discussed cases of the problem of the many are such that many would-be instances are parts of others. NDC rules out domains containing both.

### 4.3 Other views of domain restriction

Throughout this discussion, I've been assuming the view of domain restriction defended in Stanley and Szabó (2000): that all nouns co-occur with free domain variable and the extension of the complex (noun+variable) is the intersection of their interpretations. A natural question, then, is whether my arguments are compatible with other views of domain restriction. In this section I'll argue that they are by considering three alternatives.

The first alternative is defended by von Fintel (1994). On his view, domain restriction occurs via a free domain variable associated with quantifiers, rather than restricting nouns. There is no problem modifying my considerations to fit this view: NDC will be re-cast as a constraint on the quantifier-associated domain variables, and all other considerations will go through unmodified.

The second alternative is defended by a variety of theorists, e.g. Recanati (2004). On this view, there is no domain variable associated with either nouns or quantifiers. Rather, the semantic interpretation of such terms in a context is identical to their context-independent interpretation. However, sentences are often used to communicate far more than their semantic contents. This sort of view comes in a variety of flavours. For instance, one could hold that the semantic value of a sentence is a proposition, or something not fully propositional, and one can hold a variety of views about the determination and nature of communicated content. The important point for our purposes is that the considerations motivating NDC will now allow us to state a constraint on communicated content. Since our judgments, on the view being considered, track communicated rather than semantic content, NDC will explain our judgments about the problem of the many. More bluntly: the judgments motivating NDC are fully compatible with either semantic or pragmatic explanations, and the arguments will carry over straightforwardly.

The third alternative is defended by Kratzer (2005). On this view, there is also no domain variable associated with nouns or quantifiers. However, there is a tacit situational variable and restricting the range of relevant situations is how we achieve domain restriction (which is perhaps a misnomer on this view). Again, the considerations that motivate NDC will carry over straightforwardly, but placing a constraint on the saturation of the relevant domain variable.

In all of these cases, the important point is that the basic data driving NDC comes from our judgments about the communicated content of sentences in contexts. That communicated content must obey some constraint on double-counting. Whether that constraint is semantic or pragmatic, and whether it occurs on the noun, quantifier, or as a separate situation variable, is independent of the conclusion that there must be some constraint. However such a constraint is manifested, it will generate the compatibility of (i)–(vi) by divorcing the context-independent extension of 'cloud' from what particular counts using 'cloud' communicate.

In this subsection, I've shown that my argument is compatible with all extant views of domain restriction. However, another worry looms. Perhaps the data also supports a different view that doesn't cohere with my reasoning. In fact, there is a salient option: the data may be accounted for in terms of a restriction on *how* we count, rather than *what* we count. To make this idea more precise, consider a view on which the meanings of number-words require that we count by some contextually salient relation that may be less discriminating than identity. On this view, the truth conditions of a sentence of the form ⌜N Ps are Q⌝ is given as follows, where C-distinctness is a contextually determined distinctness relation (and for a set to be C-distinct is for each of its members to be C-distinct from one another):

> ⌜N Ps are Q⌝ iff the maximal C-distinct subset of $\{P \cap Q\}$ has the cardinality N.

Importantly, this view can account for the data that motivated NDC. The observation that we usually don't double count can now be recast as a generalization about C-distinctness relations in various contexts. For instance, it will be a true generalization that sets containing determinables and determinates are not C-distinct (mutatis mutandis for other ways of double counting). Since counting is context sensitive we can override these constraints. Finally, since the view is a view about number words only, it will not be surprising that other occurrences of nouns are not affected.

There are potential empirical advantages and disadvantages to such a view. In its favour, it targets number words specifically, so it is need not differentiate between occurrences of nouns which may have looked ad hoc. (Though note that this is not an advantage over vin Fintel's version of the domain restriction view.) Also in its favour, note that it may avoid indeterminacies that plague other views. Consider, again, the outfits case. To generate the true reading of '50 outfits are being delivered', the domain variable must be restricted to a non-overlapping domain. Since there are a multitude of these, it will be indeterminate which is chosen. By contrast, on the alternative view, the only contextual paramater is the C-distinctness relation and that seems to be satisfied by a unique relation.[20] Against it, note that it posits a fairly radical revision to well-established semantics for number words which is not forced by the data, since the independently motivated mechanism of domain-restriction can account for the data equally well. A full comparison of the views is beyond the scope of this discussion.

There is one way in which the truth of the C-distinctness account would not undermine the forgoing discussion. My key claim was that (i)–(vi) are compatible. I demonstrated this by assuming (ii)–(vi), and showing that NDC would be operative in (i) thereby securing its compatibility. Adopting the C-distinctness view does nothing to undermine this reasoning. Assume (ii)–(vi). (i) may nonetheless be true as long as the C-distinctness relation is more discriminating than non-identity. Given that all of our generalizations about domain restrictions will equally well motivate generalizations about C-distinctness, we will be able to conclude that, in the salient reading of (i), C-distinctness is more discriminating than identity. However, assuming that (v) is true, the C-distinctness relation relevant for interpreting (v) must be the identity relation. So, we have an equivocation: different

[20] The force of this prima facie advantage will depend on a variety of issues involving indeterminacy and interpretation.

C-distinctness relations are salient in different claims, and no contradiction can be derived.

However, the fact that we can generalize my reasoning to the C-distinctness view masks a deeper worry. The worry is that on the C-distinctness view, the most salient reading of (v) will be false. After all, any C-distinctness relation besides identity will yield the falsity of (v).

There are several responses to this deeper worry. The first is that it is not obviously a problem for the view defended in this paper if (v) ends up as false. After all, the main lesson is that (i)–(vi) are consistent, and that is compatible with the falsity of (v). The second is that there is some reason to think that (v) will not be interpreted as false. After all, it contains explicit reference to non-identity, making that a contextually salient distinctness relation.

A full evaluation of this alternative would require a fully developed view and there isn't one in the literature. However, adopting such a view would not undermine the main thesis in this paper: that on their salient readings, (i)–(vi) are compatible, and, at worst, would support an additional thesis: that (v) is false. This would merely allow us to solve the problem of the many in two ways!

### 4.4 Is the solution irrelevant?

Finally, I can imagine some theorists reacting to this discussion with a shrug. Such theorists are independently convinced that one of (i)–(vi) is false, and that this solves the problem of the many. So, they may reply, let the context fall where it may, the problem is already solved. The problem with this reaction is that, at least for many such theorists, it is dialectically unstable.

Consider, first, theorists that reject one of (ii)–(vi). If the following inference underlies their position, they are in trouble: (i)–(vi) are inconsistent, x is the least plausible, therefore, x is false (where x is their favourite premise to reject). I hope the problem is clear: the underlying motivation for their position derives from a desire to square our ordinary counting judgments with plausible metaphysical principles. Unfortunately for them, this is unnecessary: (i)–(vi) are compatible.

Now consider theorists who reject (i). They may do this either because they think that 'cloud' has an empty extension, or because they think that its extension has more members than a singleton. The latter, it should now be clear, is no reason to reject (i): NDC may guarantee that the contextually restricted extension of 'cloud' is a singleton even when its context-independent extension isn't. The former is also likely unstable. Why would one reject the existence of clouds? Like Unger (1980) one may hold that *if clouds exist, they are governed by (ii)–(vi), but those principles are incompatible with our everyday judgments*; the result of this inconsistency is that there are no clouds. Again, this reasoning falters where it alleges inconsistency.

A final reason that metaphysicians may be skeptical about the importance of the forgoing discussion is that they think that the mere possibility of domain restriction is irrelevant to metaphysicians, who intend their claims to be taken unrestrictedly. After all, if we interpret (i)–(vi) unrestrictedly (i.e. with no domain restriction) then they are incompatible.

We can reconstruct this reasoning as follows: (1) if read unrestrictedly, (i)–(vi) are incompatible, (2) metaphysicians read (i)–(iv) unrestrictedly, therefore, (3) as metaphysicians read (i)–(vi), they are incompatible.

There are two (compatible) responses to this objection. The first is that (2) is false, and the second is that even if the argument is good it won't deflate the importance of the forgoing discussion.

What reason do we have to think (2) is false? If one holds that absolutely unrestricted readings are in general impossible, then they will reject (2). Even if one holds that absolutely unrestricted readings are possible, (2) may still be suspect. Contextual domain restriction is ubiquitous in ordinary discourse, and we often restrict our domains without being aware that we are, and, in fact, in some cases it is plausible that we are not even in a position to know that we are. When engaging in metaphysics, we may often aim to interpret claims unrestrictedly (though I doubt most metaphysicians make such intentions explicit or keep them in mind) but the pervasiveness of domain restriction makes it implausible that we will always succeed. Furthermore, in doing metaphysics we often discuss the judgments of ordinary speakers in ordinary contexts and, in such cases, there is no reason to think that the content of those judgments is unrestricted. In fact, if we think of the problem of the many as bringing out a conflict between ordinary counting judgments like (i) and plausible metaphysical principle like (ii)–(vi), then there is pressure to interpret (i) as it is ordinarily interpreted which, as I've argued, accords with NDC.

Setting aside worries about (2), consider what happens if the objection succeeds. In that case, we can conclude that when interpreted as metaphysicians intend, (i)–(vi) are incompatible. What does that tell us about the problem of the many? We can now distinguish two versions of the problem of the many. On this first version, (i)–(vi) are interpreted in the ways I've been considering and they are compatible. On the second version, (i)–(vi) are interpreted as metaphysicians intend and aren't. The forgoing discussion clearly offers a solution to the first version. But it also offers insight into the second. If, as the forgoing discussion alleges, NDC is a ubiquitous and hard-to-suspend principle governing the interpretation of counting claims, then we should lose confidence in our ability to judge counting claims unrestrictedly. After all, such unrestricted interpretations are exceptional and hard to achieve. So, we should be suspicious of the independent plausibility of counting claims on the second version. This most naturally leads to a view on which we solve the second version of the problem by rejecting (i) and we explain the counter-intuitiveness as generated by the difference between the unrestricted reading of (i) and the NDC-conforming reading, which is hard to override. Of course, nothing in the forgoing discussion forces us to reject the unrestricted reading of (i), but it is a natural additional position to hold if one is convinced by my reasoning.

In addition to providing motivation for one potential solution to the unrestricted problem, our discussion may also make us skeptical about its importance. The problem of the many only arises insofar as (i)–(iv) are independently plausible; mere inconsistency of claims does not suffice for a philosophical problem. If the independently plausibility of (i)–(vi) rests on taking them to have their most salient readings, and (i) is restricted according to NDC on its most salient reading, then we're left without evidence that (i)–(vi) are independently plausible. This would deflate the unrestricted problem.[21]

## 5 Further applications

NDC has philosophical applications beyond the problem of the many. I'll conclude by mentioning three.

First, consider debates about the possibility of co-located material objects. Many intuit the impossibility of 'Two distinct objects are in the exact same place at the exact same time'. However, familiar statue/clay arguments seem to establish just such co-location. The principle may be compatible with the examples if NDC governs 'distinct objects' in the relevant contexts.

Second, consider debates about overdetermination. Some think that it is implausible that physical actions are systematically overdetermined. If overdetermination is glossed as 'Two causes give rise to the same effect', and then it may be false that physical actions are systematically overdetermined, even if every physical action has both physical and mental causal antecedents. This would be due to the fact that 'causes' must obey NDC, and no domain with both a physical and corresponding mental antecedent will.

Third, consider debates about the persistence of material objects. Ship S will undergo fission into $S_1$ and $S_2$ at *t*. How many ships are there at pre-fission time *t*-? According to Lewis' (1986) account of material objects as four-dimensional sums of temporal parts (worms), $S_1$ and $S_2$ are both extant at *t*-. Sider (1996) objects: he thinks this gives rise to the counterintuitive counting judgment that there are two ships at *t*-. To account for this judgment (among other things), he adopts stage theory, on which ordinary count nouns like 'ship' designate momentary stages. However, if NDC governs 'ship' in the relevant contexts, there may be no conflict between Lewis' theory and our ordinary counting judgments. Both $S_1$ and $S_2$ are in the context-independent extension of 'ship' at *t*-, but NDC ensures that the contextually-restricted extension is a singleton (though which singleton it is may be indeterminate).

**Acknowledgements** Thanks to Bernhard Nickel, Jonathan Payton, Matti Eklund, Jeremy Fantl, Nicholas Jones, Ofra Magidor, Alexis Wellwood, Barry Schein, Jeremy Goodman, Mark Richard, and Dan Korman, two anonymous referees as well as audiences at USC and The Society of Exact Philosophy. Support for this paper was provided by an Insight Grant from The Social Sciences and Humanities Research Council of Canada.

[21] In general, one may reasonably think that different versions of the problem of the many can be solved in different manners. See Simon (2018) for a defense of the view that the version of the problem of the many that arises for experiences is particularly difficult, and that the best solution to it may be to adopt property dualism.

## Appendix: Functional overlap

What is functional overlap? There are several challenges in understanding the notion. First, there's a worry that we don't know what it is for functions to overlap. After all, overlap is standardly understood in terms of part-sharing, and it is not clear that functions have parts, let alone share them. Second, there's a worry that many objects aren't associated with functions. I'll address these in turn.

Consider our outfits for charity from Sect. 2.4. An arbitrary shirt/pants pair has the power to clothe a single person. We can understand this power dispositionally. If we think of clothing somebody as protecting them from the environment, then a shirt/pair pants has the disposition to clothe somebody when worn. The idea, then, is that the function of a shirt is given by some of its dispositional properties.

We can think of the function of an object as given by a subset of its total dispositions.[22] Just which subset? This is partly determined by what we are counting, and partly determined by context. If we are counting houses, then the relevant function (in many contexts) is that of sheltering, which is what we associate with the kind *house*. When we are counting artworks, there will likely be a different associated function. What if something is both a house and an artwork? Well, then the function that is relevant to double-counting will depend on whether we are restricting our quantifier with 'house' or 'artwork'.

Taking the function of a kind of object to be a contextually-selected set of dispositions, we can now consider our other question: what is it for functions of two objects to overlap? I won't give a fully general answer here—one that applies to any functions whatsoever—but I will give a sufficient condition that covers some functions. Return to our outfits: in the charity scenario we associate the kind *outfit* with the dispositional property of *clothing* (understood as the dispositional property of *protecting from the environment when worn*). Particular outfits have this general property, but they also have a more specific variant: a single pants/shirt pair has the disposition *to clothe one person*. A wholly distinct pair (one that differs in both components) also has the disposition to clothe one person. Those pairs, taken together, have the disposition to clothe two people. This, however, is not true of overlapping pairs. Consider the pairs p1/s1 and p1/s2 that differ in their shirt-components but not in their pants-components. Taken separately, each has the disposition to clothe one person, but taken together they do not have the disposition to clothe two—after all two people can't wear a single pair of paints at single time. These pairs functionally overlap. We can make this a bit more precise by noting that some dispositional properties correspond to measure functions.

Measure functions are functions from entities to values on a conventional scale. For instance, the measure function volume-in-liters is a function from entities to numbers, where the numbers measure the volume (in liters) of those entities. Some measure functions are dispositional, in the sense that they measure powers/abilities. For instance there is a measure function people-can-clothe (call it PCC) that maps

[22] More carefully, for the purposes of understanding the notion of functional overlap, we can understand the function of an object as given by a subset of its set of dispositions. This account may be unsuitable for analyzing the notion of a function more generally, or for other purposes.

objects to the number of people they can clothe, such that, e.g. p1/s1 is mapped to one. For dispositional properties like the property of *clothing one person* there will be a corresponding measure function/measure pair, in this case (PCC, 1). Necessarily, if something instantiates the property of *clothing one person when worn*, PCC maps it to 1. This is the sense in which the measure function/measure pair corresponds to the dispositional property. More general properties like *being able to clothe people* correspond to the measure functions themselves, in the following sense: necessarily, if an object has the property of being able to clothe people, then it maps PCC to some positive number. Measure functions like PCC may abstract away from some irrelevant physical properties of the objects they are measuring. Even if s1 has a few more molecules than s2, PCC may map both s1/p1 and s2/p1 to 1: after all those extra molecules on s2 do not affect the number of people it can clothe.

For our purposes, measure functions are functions from both individual objects, and pluralities of objects. A measure function M is non-additive relative to some objects o1 and o2 just in case, M(o1)+M(o2) does not equal M(o1,o2), where the latter signifies that the plurality of o1 and o2 is the argument of M. We've already seen how this can happen. If o1 and o2 share a component, then the total measure may reflect this. Finally, we can understand the relevant notion of functional overlap: function f overlaps for objects o1 and o2 just in case f (understood as a disposition) corresponds to a measure function M and M is non-additive relative to o1 and o2. Again, the outfit case illustrates this. The function of clothing people that is shared by p1/s1 and p1/s2 overlaps for them because that function corresponds to PCC and PCC is non-additive for p1/s1 and p1/s2.

We can now articulate a sufficient condition for a domain to double-count relative to a context c, based on this understanding of functional overlap.

> A candidate domain d for a noun N double-counts relative to c if the function F associated with N in c corresponds to a measure function M, and M(d) is identical to M(d′), where d′ is some proper subset of d such that M(d-d′) > 0.

In articulating this condition, I slid from a measure function applying to a plurality, to that function applying to a domain, this was merely shorthand: when I wrote of a measure-function applying to a domain I intend that to be understood as the measure function applying to the plurality of the members of the domain. The idea behind this sufficient condition is that some object in the domain must add *nothing whatsoever* to the measure, and that suffices for double-counting. Again, this constraint makes perfect sense of the outfit case. When we go beyond fifty shirt/pants pairs, we will add nothing whatsoever to the measure given by PCC.[23]

Consider, for contrast, Sutton's two houses that share a massive wall. The relevant F is *sheltering*, and the measure function M is a function from amount of sheltered space they provide. Despite their shared wall, M maps the two houses to

[23] The condition that M(d-d′) > 0 guarantees that there is genine overlap, not just irrelevance. After all, if we lift this requirement then we could guarantee that a domain overlaps relative to an M solely by adding some entities that M maps to 0, e.g. a set containing five shirts and one dog would functionally overlap relative to PCC because dogs can't clothe people.

twice the number that it maps each to individually, so the sufficient condition for functional overlap is not met, despite the massive mereological overlap of the houses.

Another worry about invoking functional overlap is that in for many kinds K and contexts c, there just is no function associated with K in c. One response is that this simply doesn't undermine our sufficient condition for double-counting: when there's no associated function then we can't double-count via functional overlap. There needn't always be the possibility of functional overlap for it to sometimes be actual. Another response is that our notion of functional overlap is a proprietary one, and is not directly connected with independent investigations into teleology. Rather, we should think of it as whatever notion allows us to make sense of the outfit and house cases. These cases show us that some notion of overlap is relevant to NDC, but it is not simple part-sharing. Our notion of functional overlap is merely a first attempt to capture the notion of overlap that makes sense of our judgments in these cases. We have no reason to expect that it directly connects to indpendent discussions of the nature of functions.

## References


Baxter, D. (1988). Many-one identity. *Philosophical Papers*, *17*, 93–216.

Chisolm, R. (1976). *Person and object*. London: George Allen and Unwin Ltd.

Cotnoir, A. (2013). Composition as general identity. In K. Bennett & D. Zimmerman (Eds.), *Oxford studies in metaphysics* (Vol. 8, pp. 295–322). Oxford: Oxford University Press.

Hudson, H. (2001). *A materialist metaphysics of the human person*. Ithaca: Cornell University Press.

Jones, N. (2015). Multiple constitution. In K. Bennett & D. Zimmerman (Eds.), *Oxford studies in metaphysics* (Vol. 6). Oxford: Oxford University Press.

Korman, D. (2015). *Objects: Nothing out of the ordinary*. Oxford: Oxford University Press.

Kratzer, A. (2005). Covert quantifier restrictions in natural language, talk given at EA Ling 2005. Slides accessed online at https://semanticsarchive.net/Archive/mIzMGUyZ/Covert%20Quantifier%20Domain%20Restrictions.pdf.

Kratzer, A. (2012). *Modality and conditionals*. Oxford: Oxford University Press.

Krifka, M. (2009). Counting configurations. In A. Riester & T. Solstad (Eds.), *Proceedings of Sinn und Bedeutung* (Vol. 13, pp. 309–324). U Stuttgart: SinSpeC.

Lewis, D. (1986). *On the plurality of worlds*. Oxford: Blackwell.

Lewis, D. (1993). Many, but almost one. In J. Bacon (Ed.), *Ontology, causality and mind: Essays in honour of D M Armstrong*. New York: Cambridge University Press.

Liebesman, D. (2015). We do not count by identity. *Australasian Journal of Philosophy*, *93*(1), 121–42.

Liebesman, D., & Magidor, O. (2017). Copredication and property inheritance. *Philosophical Issues*, *27*, 131–166.

López de Sa, D. (2014). Lewis vs Lewis on the problem of the many. *Synthese*, *191*, 1105–1117.

McGee, V., & McLaughlin, B. (2000). The lessons of the many. *Philosophical Topics*, *28*, 129–51.

Noonan, H. (1993). Constitution is identity. *Mind*, *102*, 133–46.

Recanati, F. (2004). *Literal meaning*. Cambridge: Cambridge University Press.

Sattig, T. (2010). Many as one. *Oxford Studies in Metaphysics*, *5*, 145–78.

Sider, T. (1996). All the world's a stage. *Australasian Journal of Philosophy*, *74*, 433–53.

Simon, J. (2018). The hard problem of the many. *Philosophical Perspectives*, *31*, 449–68.

Sutton, C. (2015). Almost one, overlap and function. *Analysis*, *75*, 45–52.

Unger, P. (1980). The problem of the many. *Midwest Studies in Philosophy*, *5*, 411–67.

van Inwagen, P. (1981). The doctrine of arbitrary undetached parts. *Pacific Philosophical Quarterly*, *62*, 123–37.

von Fintel, K. (1994). *Restrictions on quantifier domains*, Ph.D. thesis, UMass, Amherst
Weatherson, B. (2016). The problem of the many. *The Stanford encyclopedia of philosophy (Winter 2016 Edition)*. E. N. Zalta (ed.). https://plato.stanford.edu/archives/win2016/entries/problem-of-many/.
Williams, J., & Robert, G. (2006). An argument for the many. *Proceedings of the Aristotelian Society*, *106*, 411–419.
Winawer, J., Witthoft, N., Frank, M. C., Wulund, L., Wade, A., & Boroditsky, L. (2007). The Russian Blues reveal effects of language on color discrimination. *Proceedings of the National Academy of Sciences of the United States of America*, *104*, 7780–5.
Woods, E. (forthcoming). Many, but one. *Synthese*.