Australalasian Journal of Philosophy

ISSN: 0004-8402 (Print) 1471-6828 (Online) Journal homepage: https://www.tandfonline.com/loi/rajp20

How To Count 212 Oranges

Eric Snyder & Jefferson Barlew

To cite this article: Eric Snyder & Jefferson Barlew (2019): How To Count 212 Oranges, Australasian Journal of Philosophy, DOI: 10.1080/00048402.2018.1542738

To link to this article: https://doi.org/10.1080/00048402.2018.1542738

Published online: 11 Feb 2019.

AUSTRALASIAN JOURNAL OF PHILOSOPHY
https://doi.org/10.1080/00048402.2018.1542738

ARTICLE

How To Count $2\frac{1}{2}$ Oranges

Eric Snyder and Jefferson Barlew

Ohio State University

ABSTRACT

We address a puzzle about the meanings of fraction words, due to Nathan Salmon. Counting $2\frac{1}{2}$ oranges seemingly requires enumerating a collection of objects with a non-whole cardinal number, which is incoherent by the lights of traditional analyses of cardinality expressions. Thus, some argue that counting $2\frac{1}{2}$ oranges really amounts to measuring quantities of orange-stuff. However, we claim that '$2\frac{1}{2}$ oranges' is ambiguous between counting- and measuring-interpretations; and recognizing this affords a novel solution to Salmon's puzzle that is consistent with traditional analyses.

ARTICLE HISTORY Received 9 June 2017; Revised 23 August 2018

KEYWORDS number expressions; fractions; counting; measuring; cardinality

1. The Counting Oranges Puzzle

Here's a puzzle due to Salmon (1997). Suppose there are three oranges on the table. I take one, cut it in half, eat one of the halves, and set the remaining half on the table. Now consider the Question:

(Question) How many oranges are on the table?
(Answer) There are $2\frac{1}{2}$ oranges on the table.

In this scenario, the intuitively correct answer is the Answer, '$2\frac{1}{2}$ oranges'. But there's a problem. Either half an orange is an orange or it isn't. If it is, then there are three oranges on the table. If it isn't, then there are only two. Either way, the answer to the Question is not the intuitively correct Answer. Call this *the Counting Oranges Puzzle* (COP).

COP is primarily a puzzle about the meanings of cardinality expressions—expressions whose semantic function is to count. According to Frege's [1884] celebrated analysis, these denote second-order concepts, and cardinal numbers are identified with classes of equinumerous concepts. For example, the cardinal number two is the class of two-membered concepts. Accordingly, (1) is true if the class of oranges on the table is two-membered.

(1) There are two oranges on the table.

Now, classes are extensional: they are identified by the objects belonging to them. It follows that cardinal numbers must be whole, on Frege's view. There cannot be a concept ϕ such that more than two but fewer than three objects fall under ϕ, for instance. But this is what the Answer apparently requires. We seem to be counting the number of oranges on the table in terms of a fractional number. Consequently, stating coherent Fregean truth-conditions for the Answer appears impossible.

According to Salmon, the problem with Frege's analysis is its extensionality. Thus, the lesson of COP is that counting is inherently *intensional*. Cardinal numbers aren't classes of equinumerous concepts; rather, they are measures of pluralities relative to some counting property. Depending on that property, we get different answers to 'how many'-questions. For example, relative to being a whole orange, the answer to the Question will be 'two'. Relative to being some orange-stuff, however, the answer will be the Answer. That's because we assign different 'weights' to members of the plurality in question: we assign 'weight' 1 to the whole oranges, 'weight' .5 to the half-orange, and, adding these together, we get the Answer '2.5 oranges'.

Frege's and Salmon's analyses have analogues in linguistic semantics. Like Frege, Barwise and Cooper [1981], Landman [2004], Scontras [2014], and Rothstein [2017] analyze count nouns as denoting sets of individuated, and thus countable, objects, where cardinalities represent how many such objects constitute a collection.[1] As a result, cardinalities are necessarily whole.

The alternative is due to Krifka [1989]. Basing his analysis on classifier languages, he analyses all English nouns as basically mass, denoting quantities of stuff, so that all numerical modification corresponds to measuring amounts of stuff.[2] On Krifka's view, the primary difference between, for example, 'orange' and 'water' is that count nouns have a certain classifier-like element built into their meanings, called a natural unit (NU). Krifka's analysis of 'orange' is given in (2), where 'NU' is a contextually determined function mapping quantities of orange-stuff to real numbers representing a ratio of that stuff to a natural unit of orange.

(2) $[[\text{orange}]] = \lambda n.\lambda x.orange(x) \wedge \text{NU}(orange)(x) = n$

As Kennedy and Stanley [2009: 619n16] note, Krifka's semantics affords a solution to COP that is strikingly similar to Salmon's:

> This proposal is very similar to the one suggested by Nathan Salmon (1997, p. 10), who suggests that '… numbers are not merely properties of pluralities simpliciter, but relativized properties'. An advantage of this analysis is that it provides a semantics for nominals with fractional number terms … Assuming that the sorted NU function is not constrained to return whole numbers as values, a phrase like '2.5 oranges' denotes a property that is true of orange-stuff whose measure equals 2.5 orange-units.[3]

The solution on offer is this. In the context of COP, we are measuring orange-stuff relative to a typical whole orange. Since this measurement is determined by a ratio, and since that ratio is 2.5 to 1, it's little wonder that the intuitively correct

[1] For example, Landman, Scontras, and Rothstein analyze cardinal 'two' as (ia), where'' maps pluralities to numbers representing their atomic parts [Link 1983]. Because atoms have no proper parts, these numbers are necessarily isomorphic to the naturals.

(i) a. $[[\text{two}]] = \lambda P.\lambda x.(x) = 4 \wedge P(x)$
b. $[[\text{two oranges}]] = \lambda x.(x) = 2 \wedge oranges(x)$

[2] See Rothstein [2017] for discussion.

[3] See Liebesman (2016) for a similar view.

Answer involves a non-whole number. Call this the Krifka-Inspired Solution [+Salmon] (KISS).

We believe that KISS predicts an interpretation of the Answer that is genuinely available, one that involves measuring orange-stuff in terms of a non-whole number. However, we also think that the Answer has an interpretation that is not predicted by KISS, one that involves counting via discrete numbers. In fact, we believe that the Answer is ambiguous in a way resembling container phrases such as 'glass of water'. Consider (3a):

(3) a. There are four glasses of water in the soup.

b. There's a plurality of glasses x such that x consists of four individual glasses, each of which is filled with water and is in the soup.

c. There's a quantity of water x such that x measures four glasses worth and x is in the soup.

Rothstein (2017) shows that (3a) is ambiguous between an individuating interpretation (II) paraphrased in (3b) and a measure interpretation (MI) paraphrased in (3c).[4] For the II, suppose that Mary has a strange way of heating water: she pours it into glasses and places them into boiling soup. For the MI, imagine instead that Mary wants to make soup, and the recipe calls for four glasses worth of water. She fills a certain glass four times with water, pouring the contents each time into the soup, or perhaps she fills a jug with an equivalent amount and pours that directly into the soup.

One of our primary empirical contentions is that I/M-ambiguities are not limited to container phrases. In fact, (4a) is similarly ambiguous:

(4) a. There are four oranges in the punch.

b. There's a plurality of oranges x such that x consists of four individual oranges, each of which is in the punch.

c. There's a quantity of orange x such that x measures four oranges worth and x is in the punch.

For the II in (4b), suppose that Mary has already made the punch and wants to decorate it for the party. She thinks that adding some fruit would work, so she drops several apples and four oranges into the punch. For the MI in (4c), imagine instead that Mary wants to make punch for the party, and the recipe calls for some orange pulp. Accordingly, she pours approximately four oranges worth of prepackaged orange pulp directly into the punch.[5]

The theoretical significance of I/M-ambiguities is that, generally speaking, IIs involve *counting* individuated objects such as glasses or oranges, while MIs involve *measuring* an amount of a substance such as water or orange,[6] relative to a

[4] See also Landman [2004], Partee and Borschev [2012], and Scontras [2014].

[5] A reviewer observes that the peel may or may not figure into MIs. Whereas the MI is naturally understood as measuring quantities of orange without the peel, (i) is naturally interpreted as measuring quantities of chopped orange with the peel.

(i) There are four candied oranges in the fruitcake.

[6] Following, e.g., Chierchia [1998], by 'substance' we mean the sorts of things denoted by plural and mass nouns, which we neutrally call 'pluralities' and 'quantities'. Pluralities are collections (sums or sets) of 'atoms' [Link 1983], or individuated, countable objects, while quantities can be understood in a variety of ways consistent with the semantics developed here (see, e.g., Link, Chierchia, and Rothstein [2017]).

non-standardized unit of measurement. Whereas traditional analyses naturally predict IIs, KISS instead predicts MIs. For example, on the former, (4a) will be true only if there are four whole oranges in the punch, which is false in the scenario where Mary pours prepackaged orange pulp directly into the punch. The problem for KISS is interestingly different. Although it correctly predicts that (4a) is true in this scenario, because the same natural unit—a typical orange—will be relevant in *both* scenarios,[7] it predicts that (4a) should have the same truth-conditions in those scenarios, paraphraseable as (4c). Consequently, it has no obvious semantic means of explaining various semantic contrasts between the II and the MI, which we demonstrate in section 2. Thus, neither kind of analysis appears capable of capturing both interpretations of (4a).

The key to solving COP, we submit, is to recognize that the Answer is similarly I/M-ambiguous. It has an II paraphraseable as the Answer_{I},[8] and an MI paraphraseable as the Answer_{M}:

(Answer_{I}) There are two oranges on the table, and there is one half-orange on the table.

(Answer_{M}) There are two and a half oranges worth of orange on the table.

This is shown in section 2, where we develop and apply certain linguistic diagnostics designed to reveal I/M-ambiguities more generally. As before, the problem with both the traditional analyses and KISS is that neither appears capable of capturing both interpretations. Since KISS predicts that the Answer should have truth-conditions paraphraseable as the Answer_{M} in individuating *and* measure contexts, it has no obvious semantic means for distinguishing the Answer_{I} from the Answer_{M}. Conversely, traditional analyses appear incapable of capturing *either* interpretation. Again, cardinality expressions designate the cardinality of countable objects, and so fractional cardinalities are incoherent. Thus, it's hard to see how the Answer could have *any* coherent interpretation, let alone two.

Our primary theoretical contention is that solving COP does not require adopting KISS or somehow rendering fractional cardinalities coherent. After sketching a general analysis of I/M-ambiguities in section 3, we show in section 4 that combining independently motivated analyses of the component expressions in the Answer—'two', 'and', and 'half'—in a natural way predicts that the Answer should be interpretable as Answer_{I}, and more generally that complex fractions are I/M-ambiguous. According to the resulting analysis, we are actually counting *three* things on the II of the Answer—namely, two whole oranges and a single half-orange.

This affords a novel solution to COP. Although there are, strictly speaking, only two oranges on the table in that context, the Answer remains intuitively correct because we are individuating two whole oranges and one half-orange, while counting them separately. Since this is consistent with traditional analyses, COP presents no special threat to them.

---

[7] Otherwise, NU would not correctly return '4', thus verifying (4a).

[8] The analysis developed in section 4 implies that, at least with fractional noun phrases like 'two and a half oranges', plural marking on the noun is not a reflection of its meaning, but reflects instead morphosyntactic features. Although we say 'one orange' and 'a half orange', not 'one oranges' or 'a half oranges', we also say 'one and a half oranges', not 'one and a half orange'. A similar view is found in Krifka [1989], who observes that we say 'zero oranges' and 'one point zero oranges', not 'zero orange' or 'one point zero orange'.

2. Individuating and Measuring

In this section, we develop three general diagnostics for disambiguating IIs and MIs.[9] For each diagnostic, we begin by applying it to the uncontroversially I/M-ambiguous container phrase 'four glasses of water', showing that each diagnostic successfully distinguishes IIs from MIs. We then apply that diagnostic to the Answer, revealing it to be I/M-ambiguous.

Our first diagnostic is a natural extension of one found in Rothstein (2017). According to it, the suffix '-ful' disambiguates container phrases by forcing a MI. Consider the contrast between (5), as uttered in the M(easure)-Context provided, and (6), as uttered in the I(ndividuating)-Context.

M-Context: Mary is making soup. The recipe calls for four glasses worth of water. She fills a certain glass four times with water, pouring the contents each time into the soup. John says 'As I was passing by the kitchen, I noticed some water in the soup. How many glasses of water did you put in it?'

(5) Mary: 'Four {glasses/glassfuls}.'

I-Context: Mary is heating water for coffee. She fills four glasses with water and places them into the boiling soup. John says 'As I was passing by the kitchen, I noticed some glasses in the soup. How many glasses of water did you put in it?'

(6) Mary: 'Four {glasses/??glassfuls}.'

Rothstein proposes that the semantic function of '-ful' is to transform container nouns into measure nouns. Consequently, 'glassful' denotes an amount of water measured in terms of an *ad hoc* glass-unit[10]—that is, an amount of water that would fill a salient, contextually determined glass. Accordingly, 'four glassfuls of water' will be true of those quantities of water measuring four glasses worth—that is, the MI. Hence, the acceptability of 'glassfuls' in (5) but not (6) makes sense if M-Contexts impose MIs, I-Contexts impose IIs, and '-ful' forces MIs.

One of our primary empirical contentions is that '-ful' is similar to, but more specific than, 'worth' in this respect. Generally speaking, 'worth' imposes MIs. For example, 'four glasses worth of water' also denotes quantities of water measuring four *ad hoc* glass-units. However, unlike '-ful', 'worth' occurs acceptably with non-container nouns: whereas we cannot say '??four orangefuls of orange' or '??four grainfuls of rice', we can say 'four oranges worth of orange' and 'four grains worth of rice'.[11] In both cases, we get an amount measured in terms of an *ad hoc* unit—four times the amount of orange constituted by some salient orange, and four times the amount of rice constituted by some salient grain of rice.

Thus, our first diagnostic is this: because 'worth' imposes an MI, it is acceptable in M-Contexts but not I-Contexts. Now consider the result of applying this to the Answer:

M-Context: John is on a diet requiring him to have $2\frac{1}{2}$ oranges per day. At Smoothie World, smoothies are made with prepackaged orange pulp, and employees know how much a typical orange produces. To fulfil his dietary requirements, John orders a smoothie,

[9] For further applicable diagnostics, see Rothstein [2017].
[10] See Partee and Borschev [2012].
[11] See Rothstein [2017: ch. 8] and Snyder and Barlew [2016] for relevant discussions.

which Mary makes. John says, 'I want to make sure that smoothie will meet my dietary requirements. How many oranges are in it?'

(7) Mary: 'Two and a half {oranges/oranges worth}.'

I-Context: Three oranges are on the table. Mary takes one, cuts it in half, eats a half, and places the remaining half on the table. John says, 'As I was passing by the kitchen, I noticed some oranges on the table. How many oranges are on the table?'

(8) Mary: 'Two and a half {oranges/??oranges worth}.'

As before, (7) and (8) make sense if M-Contexts induce an MI of the Answer, I-Contexts induce an II, and 'worth' forces MIs. Notice this I-Context is just Salmon's context for COP, thus suggesting that COP imposes an *II* of the Answer. Our additional diagnostics confirm this.

Our second diagnostic involves plural pronouns. As Rothstein observes, container phrases are acceptable with these on IIs but not MIs.

M-Context: Same as (5). Pointing at the soup, Mary says,

(9) a. ?? Those are four glasses of water.
b. That is four glasses of water.

I-Context: Same as (6). Pointing at the soup, Mary says,

(10) a. Those are four glasses of water.
b. That is four glasses of water.

Generally, 'those' refers to a plurality of individuated objects, and we have a plurality of individuated glasses in the I-context but not in the M-context—hence the difference in acceptability between (9a) and (10a). Singular 'that' is also acceptable in both contexts, but plausibly for different reasons. On the II, we can think of the four glasses as constituting a single group,[12] thus making available a referent for 'that'. On the MI, we have a single quantity of water, and this supplies an appropriate referent.

We see a similar contrast with the Answer.

M-Context: Same as (7). Pointing at the smoothie, Mary says,

(11) a. ? ? Those are $2\frac{1}{2}$ oranges.
b. That is $2\frac{1}{2}$ oranges.

I-Context: Same as (8). Pointing at the table, Mary says,

(12) a. Those are $2\frac{1}{2}$ oranges.
b. That is $2\frac{1}{2}$ oranges.

Again, this makes sense if 'those' refers to some salient plurality of individuated objects. Later, we will argue that such a plurality is available in the I-Context, consisting of two whole oranges and one half-orange.

Our final diagnostic involves 'number' and 'amount'. These can be used to tease apart IIs and MIs directly, as (13) and (14) demonstrate:

Context: John and Mary are both heating water by placing glasses filled with water into boiling soup. However, John's glasses are exactly half the size of Mary's. Both place four glasses into their respective soups.

[12] See Landman [2004].

(13) a. There are four glasses of water in Mary's soup, and the same number are in John's soup. (true)

b. There are four glasses of water in Mary's soup, and the same amount is in John's soup. (potentially false)

Context: Same as for (13), only John places eight glasses into his soup.

(14) a. There are four glasses of water in Mary's soup, and the same number are in John's soup. (false)

b. There are four glasses of water in Mary's soup, and the same amount is in John's soup. (potentially true)

It has been independently argued that, whereas 'amount' relates substances to their measures, potentially including their cardinalities, 'number' disambiguates 'amount' by selecting specifically for cardinality.[13] Consequently, (13a) will be true if there are exactly as many glasses in Mary's soup as there are in John's soup, which is true. In contrast, (13b) is potentially ambiguous, depending on which dimension of measurement is relevant. If it's cardinality, then (13a,b) are equivalent. But if it's volume, then (13b) will be true only if John's soup contains the same amount of water as Mary's does, which is false. We see the reverse in (14). Since IIs involve counting the individuated members of a plurality, while MIs involve measuring a substance in terms of an *ad hoc* measure, these examples directly reveal 'four glasses of water' to be I/M-ambiguous.

We see the same pattern with the Answer.

Context: Like (8), except John's oranges are exactly half the size of Mary's. Like Mary, he has three oranges, cuts one in half, eats one half, and leaves the other on his table.

(15) a. There are $2\frac{1}{2}$ oranges on Mary's table, and the same number are on John's table. (true)

b. There are $2\frac{1}{2}$ oranges on Mary's table, and the same amount is on John's table. (potentially false)

Context: Same as (15), except John leaves five oranges on his table.

(16) a. There are $2\frac{1}{2}$ oranges on Mary's table, and the same number are on John's table. (false)

b. There are $2\frac{1}{2}$ oranges on Mary's table, and the same amount is on John's table. (potentially true)

Again, this makes sense if 'number' relates pluralities of individuated objects to their cardinality, while 'amount' potentially relates quantities of orange-stuff to their volume. Hence, these examples directly reveal the Answer to be I/M-ambiguous.

Our diagnostics reveals two important facts. First, I/M-ambiguities are not limited to just container phrases. Even ordinary count nouns like 'orange' can be I/M-ambiguous. Second, COP imposes an II, which is semantically distinct from the MI. However, because the same natural unit—a typical orange—is relevant in both I-Contexts and M-Contexts, KISS not only predicts that the Answer should have a MI in the context of COP; it also has no obvious semantic means of distinguishing between these interpretations. In order to explain how they arise, we will first need

[13] See Snyder (2017).

a general analysis of I/M-ambiguities. We develop this in the next section, and we extend that analysis to fractional modifiers like 'two and a half' in section 4.

3. The Universal Measurer

This section sketches a general analysis of I/M-ambiguities, showing how 'four glasses of water' and 'four oranges' give rise to IIs and MIs. We begin with container nouns like 'glass'. In examples such as (17), the default assumption is that 'glass' is a monadic predicate true of individual glasses, while 'glasses' is a predicate true of pluralities of glasses [Rothstein 2017].

(17) Mary drank from {a glass/glasses}.

However, 'glass' takes on a distinctively relational character when occurring in container phrases. On the II, 'four glasses of water' means, roughly, 'four glasses *containing* water'. On the MI, it means, roughly, 'water *measuring* four glassfuls'. Thus, the question is that of how the same noun—'glass'—can serve three semantic functions—namely, a monadic predicate true of individuated glasses, a relational predicate relating glasses to substances that they contain, and a relational predicate measuring quantities of a substance in terms of an *ad hoc* glass-unit.

Our analysis partially follows a suggestion from Scontras [2014], based on Rothstein's analysis of I/M-ambiguities in Modern Hebrew. According to Rothstein, IIs of container phrases result from a type-shifting principle that we call *Relational Noun Shift* (RNS). This is given in (18), where 'R' is free, signifying that its interpretation is contextually supplied:

(18) $\lambda P.\lambda Q.\lambda x.\exists y.P(x) \wedge Q(y) \wedge R(x,y)$ (RNS)

RNS shifts a monadic predicate like 'glass' into a relational modifier, which can combine with a phrase like 'of water' to denote quantities of water bearing some contextually determined relation to glasses. With container nouns, this will be the being-filled-with relation. Hence, 'glasses of water' will be a predicate true of pluralities of glasses filled with water.

Assuming either Scontras' or Rothstein's analysis of cardinality expressions,[14] 'four glasses of water' will thus denote a predicate true of pluralities consisting of exactly four glasses, each filled with water. Accordingly, (3a) will be true if there is at least one such plurality in the soup—namely, (3b).

(3) b. There's a plurality of glasses x such that x consists of four individual glasses, each of which is filled with water and is in the soup.

Thus, IIs result from shifting 'glass' from a monadic predicate true of individual glasses to a relational noun true of filled glasses, and from then applying the meaning of a cardinality modifier to the result.

To capture MIs, Rothstein proposes that '-ful' is silently suffixed to the noun 'glass', transforming it into a measure noun which measures a substance in terms of an *ad hoc* glass-unit. Hence, (19) is true if there is an amount of water in the soup equal to that which would be obtained by filling a certain glass four times:

(19) There are four glassfuls of water in the soup.

[14] See note 1.

This is equivalent to the MI paraphrased in (3c):

(3) c. There's a quantity of water x such that x measures four glasses worth and x is in the soup.

So, MIs are straightforwardly predicted if '-ful' is silently suffixed to 'glass' within M-Contexts.

Rothstein's analysis is designed to account for MIs of container phrases, as only container nouns can be suffixed with '-ful'. In order to account for the generality of I/M-ambiguities, we follow Scontras in proposing that MIs result instead from type-shifting. More specifically, they result from the type-shifting operation given in (20), which we dub *the Universal Measurer* (UM):[15]

(20) $\lambda P.\lambda Q.\lambda n.\lambda x.\gamma(Q)(x) \wedge \mu_P(x) = n$ (UM)

Here, γ is a 'grinding'-like[16] operation taking a predicate Q and returning all (possibly improper) parts of the individuals in its extension, and μ_P an *ad hoc* measure—namely, a contextually determined function from individuals x to numbers n representing a potentially non-canonical measure of x based on P.[17] Applying UM to 'glass' returns a relational modifier that combines with a substance-denoting phrase like 'of water' to return a predicate true of quantities of water equal in amount to that which would fill a certain glass n-times. Assuming 'four' functions as a numeral in measure phrases,[18] 'four glasses of water' will denote quantities of water measuring four glassfuls, or four glasses *worth* of water,[19] thus leading to the MI paraphrased in (3c).

Positing UM makes it possible to account for the generality of I/M-ambiguities. Consider (21):

Context: John and Mary work in a recycling plant specializing in recycling aluminium cans and tires. Once the items are ground-up, they are sifted and packaged into boxes. John and Mary find some unmarked boxes. After dumping the contents into two sorted piles, Mary points at one and says,

(21) a. That's four boxes of tires.

b. That's four boxes of tire.

Clearly, an MI is intended here: Mary is talking about four boxes *worth of* tires, not four boxes *containing* tires. Moreover, both 'boxes of tires' and 'boxes of tire' plausibly receive a 'grinding'-like interpretation—that is, Mary intending to talk

[15] UM differs from a similar operation proposed by Scontras's in two ways. First, whereas Scontras's takes a substance-denoting kind, UM instead takes a substance-denoting predicate. Secondly, UM contains the grinding-like operation γ (see note17). Note that traditional grinding-operations return the proper material parts of an individual (see Rothstein [2017: ch. 7]). In contrast, γ returns possibly *im*proper material parts of an individual or plurality. This is necessary since, as a reviewer observes, 'two cups of beans' can be understood as possibly measuring beans parts, whereas 'four boxes of oranges' is usually understood as measuring whole oranges.

[16] See Rothstein [2017] for 'grinding'-interpretations of count nouns.

[17] Roughly following Rothstein [2017], we define γ as in (i), where '$\sqsubseteq_m$' denotes material parthood [Link 1983].

(i) $\gamma = \lambda P.\lambda x.\exists y.x \sqsubseteq_m y \wedge P(y)$

[18] Compare Scontras [2014] and Rothstein [2017]. Our analysis is intended to be neutral with respect to the lexical meaning of 'four'. There are *many* analyses of 'four' available in the literature, differing in their syntactic and semantic assumptions (see Rothstein [2017]). For concreteness, we follow Scontras and Rothstein in assuming that 'four' functions as a numeral on MIs and as a cardinality modifier on IIs, perhaps via type-shifting.

[19] UM plausibly encodes something like the meaning of 'worth', suitably restricted. For relevant discussion, see Rothstein [2017].

about tire-*stuff* rather than whole tires. On our analysis, applying γ to 'tires' returns all material parts of pluralities of tires, including all individual tires and all bits of metal and rubber constituting them. In other words, applying γ to 'tires' effectively 'massivizes' the noun, thus explaining why (21a) and (21b) are equivalent in this context.[20]

Now consider the MI of 'four oranges'. This is not directly captured by UM. That's because UM is designed for container phrases like 'glass of water', where 'glass' provides the first argument (P) and 'of water' provides the second argument (Q). Since 'orange' is monadic, MIs cannot be compositionally derived via UM. Nevertheless, it's easy to see how UM should be modified. As (4c) suggests, the MI of 'four oranges' can be paraphrased as 'four oranges worth of orange':

(4) c. There's a quantity of orange x such that x measures four oranges worth and x is in the punch.

Thus, unlike with the MI of 'four glasses of water', in this case the substance being measured and the *ad hoc* unit employed *are of the same kind*. That is, we are measuring quantities of orange in terms of an *ad hoc* orange-unit. Intuitively, what's needed is some way of guaranteeing that the first two arguments of UM are identical.

We see something similar with relational verbs like 'bathe' when occurring without overt objects. Witness (22b), which, unlike (22a), can only be understood *reflexively*—that is, as claiming that John bathed himself:

(22) a. John bathed the baby.

b. John bathed.

Following, for instance, Reinhart and Siloni [2005], we assume that the reflexivization of 'bathe' is a lexical process whereby the two arguments of 'bathe' are reset to being identical, so that the bathee and the bather are the same individual.

Returning to UM, we assume that the reflexivization of 'bathe' is one instance of a more general principle that reflexivizes relations of various types. This leads to a reflexivized version of UM, or what we call *the Reflexivized Universal Measurer* (RUM), given in (23). RUM is just UM with its first two arguments reset to be identical:

(23) $\lambda P.\lambda n.\lambda x.\gamma(P)(x) \wedge \mu_P(x) = n$ (RUM)

Applying RUM to 'orange' results in a measure of orange-parts based on an *ad hoc* orange-unit. Consequently, 'four oranges' will be a predicate true of those quantities of orange measuring four *ad hoc* orange-units, or four oranges worth of orange. Hence, (4a) will be true if there are four oranges worth of orange in the punch, or the MI in (4c).[21]

[20] (i) provides a derivation for (21a):
(i) a. UM([[box]]) $= \lambda Q.\lambda n.\lambda x.\gamma(Q)(x) \wedge \mu_{box}(x) = n$
b. UM([[box]])([[of tires]])$= \lambda n.\lambda x.\gamma(tires)(x) \wedge \mu_{box}(x) = n$
c. UM([[box]])([[of tires]])([[four]])$= \lambda x.\gamma(tires)(x) \wedge \mu_{box}(x) = 4$
$\rightsquigarrow$Material parts of tires measuring 4 boxes worth

[21] (i) provides a derivation:
(i) a. RUM([[oranges]]) $= \lambda n.\lambda x.\gamma(oranges)(x) \wedge \mu_{orange}(x) = 4$
b. RUM([[oranges]])([[four]]) $= \lambda x.\gamma(oranges)(x) \wedge \mu_{orange}(x) = 4$
$\rightsquigarrow$ Material parts of oranges measuring 4-oranges worth

To summarize, positing UM allows for a completely general analysis of I/M-ambiguities.[22] IIs result just as traditional analyses predict. For example, the II of (4a) results if 'four' functions as a cardinality modifier, and so counts how many individual oranges make up a certain plurality. An MI results instead if 'orange' is shifted into a reflexivized measure noun. Which interpretation results depends entirely on context. Specifically, I-Contexts force a cardinality interpretation of 'four', which presupposes a domain of *individuated* objects to count. M-Contexts trigger an application of UM or RUM, thus leading to MIs.

The resulting analysis overcomes the flaws in the traditional analyses and KISS. Whereas the former predicts that (4a) should be false in the smoothie scenario, RUM allows for a true interpretation equivalent to that predicted by KISS. However, because that interpretation results from *shifting* the lexical meaning of 'orange' from that of a predicate true of individuated oranges, unlike KISS, the present analysis also captures the II, while relating the two interpretations in a natural way. In the next section, we extend our analysis to fractional modifiers like 'two and a half', and thus the Answer.

4. How to Count $2\frac{1}{2}$ Oranges

We saw in section 2 that the Answer is I/M-ambiguous. We argue in this section that both interpretations arise in a manner parallel to those of 'four oranges'. We begin with the MI, paraphrased in the Answer$_{\mathrm{M}}$:

(Answer$_{\mathrm{M}}$) There are two and a half oranges worth of orange on the table.

It should be fairly evident how this interpretation arises on the analysis just sketched: we apply RUM to 'orange', thus shifting it into a reflexivized measure noun, where '$2\frac{1}{2}$' functions as a numeral signifying a measure of orange-stuff in terms of an *ad hoc* orange-unit. Hence, the Answer will be true if there is a quantity of orange on the table measuring two and a half oranges worth. This would be true if, for example, $2\frac{1}{2}$ oranges worth of orange slices, each from possibly different oranges, were on the table.

Crucially, however, this is *not* the interpretation of the Answer relevant to COP, as revealed in section 2. Rather, COP imposes an *II* of the Answer. This explains the contrasts in (24) and (25), in the same context:[23]

(24) There are {two (true)/three (false)} oranges on the table.

(25) a. There are two oranges on the table, and a half an orange. (true)

b. There are three oranges on the table, only one is a half. (false)

Again, IIs involve individuating the members of pluralities. With COP, the relevant plurality consists of *three* such objects, we submit—namely, two whole oranges and one half-orange. That's because 'two and a half' in the Answer is functioning as a complex modifier of 'oranges', thus leading to the II of the Answer, or the Answer$_{\mathrm{I}}$:

(Answer$_{\mathrm{I}}$) There are two oranges on the table, and there is one half-orange on the table.

[22] A reviewer inquires whether all nouns can provide an *ad hoc* measure for (R)UM. There may well be numerically modified NPs which do not give rise to MIs (consider 'Mary divided three numbers by two'), although we will leave this question to future research.

[23] We thank a reviewer for this suggestion.

Thus, on our view, 'two and a half' in the Answer is functioning similarly to 'large and small' in (26):

(26) There are large and small apples on the table.

Like the Answer$_{\mathrm{I}}$, (26) can be paraphrased as a conjunction of modified noun phrases. And, just as one would not reasonably conclude from (26) that there are apples on the table that are both large and small, one would not reasonably conclude from the Answer that there are oranges having two different cardinalities, both two and one half.

If this informal characterization is correct, then COP presents no special threat to traditional analyses of cardinality expressions. It would be threatening only if '$2\frac{1}{2}$' in the Answer specified a *fractional cardinality*. But our claim is that it fails to do so on either interpretation. On the MI, it designates a *measure* of orange-stuff, while on the II it functions as a complex modifier, thus leading to the Answer$_{\mathrm{I}}$.

Anticipating this sort of response, perhaps, Salmon [1997: .6] challenges our claim that 'two and a half' can function as a complex modifier:

> The numeral '2' occurring in [the Answer] has been separated from its accompanying fraction, and now performs as a solo quantifier. The fraction itself has been severely mutilated. The numeral '1', which appears as the fraction's numerator in [the Answer], has ascended to the status of anonymous quantifier, functioning independently both of its former denominator and of the quantifier in the first conjunct. At the same time, the word 'half' appearing in [the Answer$_{\mathrm{I}}$] has been reassigned, from quantifier position to predicate position. In effect, the mixed number expression '$2\frac{1}{2}$', occurring as a unit in [the Answer], has been blown to smithereens, its whole integer now over here, the fraction's numerator now over there, the fraction's denominator someplace else… Even the schoolboy knows that the phrase 'and a half' in [the Answer] goes with the 'two' and not with the 'orange'.

Colourful rhetoric aside, we take Salmon to be making a simple but nevertheless substantial point. Intuitively, '$2\frac{1}{2}$' seems to function as a syntactic and semantic unit. But if the Answer$_{\mathrm{I}}$ genuinely is available, then it must function as a complex modifier, with the 'two' and the 'half' each separately modifying 'oranges'.[24] Thus, Salmon's remarks pose a serious challenge: compositionally derive the Answer$_{\mathrm{I}}$ from the meanings of its component expressions.

We will address Salmon's challenge head-on in what follows. We will show that it is possible to derive the Answer$_{\mathrm{I}}$ compositionally from independently motivated analyses of 'two', 'and', and 'half'. In section 4.1, we address the semantic contribution of 'and' via Krifka's [1990] analysis of generalized cumulative conjunction. The contribution of 'half' is addressed via a slightly modified version of Ionin, Matushansky, and Ruys' [2006] analysis of fraction words in section 4.2. We derive the Answer$_{\mathrm{I}}$ by combining these in section 4.3.

4.1. Generalized Cumulative Conjunction

It is well-known that 'and' is ambiguous between propositional, or 'Boolean', conjunction and cumulative, or 'non-Boolean', conjunction. Familiar examples of the former are given in (27), as they are both paraphraseable as the conjunction of two propositions:

[24] See note 8.

(27) a. John sings and Mary dances.

b. Mary sings and dances.

(28a–b) are examples of cumulative conjunction, and cannot be paraphrased as conjunctive propositions:[25]

(28) a. John and Mary met at the mall.

b. That flag is entirely green and white.

For example, (28a) cannot mean that John met at the mall and (also) Mary met at the mall.

In addition to demonstrating that 'and' is plausibly ambiguous, (27) and (28) show that both forms of 'and' can coordinate expressions of different semantic types. For example, propositional 'and' coordinates two clauses in (27a) and two verbs in (27b), while cumulative 'and' coordinates two names in (28a) and two adjectives in (28b). *Generalized conjunction* refers to the theory of how both forms of 'and' can take on various semantic types, and has been developed by Krifka [1990], among others.

Now, consider (26) again:

(26) There are large and small apples on the table.

If 'and' here were propositional conjunction, then the incoherent interpretation in (29a) would result. That's because 'large and small' would distribute over 'apples', and so both adjectives would describe the same plurality of apples. On the other hand, the coherent interpretation in (29b) results if 'and' is cumulative:

(29) a. There are apples on the table that are both large and small.

b. There are large apples on the table, and there are small apples on the table.

On Krifka's analysis, 'large' and 'small' describe possibly different parts of a plurality of apples. For example, if there are five apples on the table, it could be that two of them are large and the rest are small.

Our claim is that 'and' in 'two and a half oranges' is also cumulative. If it were propositional, then 'two and a half' would distribute over 'oranges', thus resulting in the incoherent (30a):

(30) a. There are oranges on the table that number both two and one half.

b. There are oranges on the table which number two, and there is also one half-orange on the table.

The task of sections 4.2 and 4.3 is to show how the coherent interpretation in (30b) results from cumulative 'and', where 'two' and 'a half' describe different parts of a plurality consisting of possibly different kinds of things.

4.2. Fractions

To account for fraction words such as 'half', 'quarter', 'third', etc., we adopt a slight modification of the analysis due to Ionin, Matushansky, and Ruys [2006;

[25] (27) and (28) are from Krifka [1990].

henceforth IM&R]. IM&R model numerators as cardinality expressions, and denominators as 'packaging' parts of things into individuals that numerators then count—namely, halves, quarters, thirds, etc. For example, 'two thirds of an orange' counts proper parts of an orange—namely, the thirds. Similarly, 'two thirds of the oranges' counts proper parts of a plurality of oranges—namely, the thirds. In either case, fractions like 'one half', 'two thirds', etc. essentially involve *counting* parts of a whole.[26]

According to IM&R, 'third' has a fairly complicated meaning consisting of three components. These are separated into clauses (i)–(iii) in (31), for convenience. We have added clause (iv):[27]

(31) $[[\text{third}]] = \lambda n.\lambda P.\lambda x.\exists y.$

i. $P(y) \wedge \forall z[P(z) \rightarrow z \sqsubseteq y] \wedge$

ii. $\exists S[\Pi(S)(y) \wedge |S| = 3 \wedge x \in S \wedge$

iii. $\exists \mu[\mu \in M \wedge \forall s_1, s_2[(s_1 \in S \wedge s_2 \in S) \rightarrow (\mu(s_1) = \mu(s_2))]] \wedge$

iv. $\exists X[X \subseteq S \wedge x = \sqcup X \wedge |x| = n]]$ (where $\sqcup$ is the sum-operation)

To illustrate, consider (32):

(32) John ate two thirds of the oranges on the table.

According to (31), 'two thirds of the oranges' is a predicate true of subpluralities of oranges that have a cardinality of two and that divide the totality of oranges into thirds. Specifically, clause (i) guarantees that the subpluralities have a *maximal* part—namely, the whole plurality of oranges. Clause (ii) guarantees that there is a *partitioning* of that plurality into three non-overlapping subparts, thus ensuring that we do not count the same parts twice. Clause (iii) guarantees that the three parts all measure the same, according to some contextually determined measure (for example, cardinality or volume). Finally, clause (iv) counts the thirds partitioned using the fraction's numerator. Putting it all together, (32) will be true if there's a plurality of oranges on the table that are divisible into three non-overlapping and quantity-uniform subparts—namely, the thirds, two of which John ate.

Since IM&R's partitioning function 'packages' substances into countable parts, the resulting analysis predicts that we can semantically individuate thirds of an orange, in a way similar to slices of orange, for example. This is borne out in I-Contexts:

I-Context: Mary cuts an orange into thirds, leaving two on the table. John asks, 'What kind of fruit is on the table?'

[26] Following IM&R, we assume that 'of' denotes a contextually determined parthood-relation, depending on the kind of noun with which 'third' combines. In the case of singular nouns like 'an orange', it denotes material parthood [Link 1983], so that 'of an orange' denotes the quantities of orange constituting some orange, while 'of a kilo of oranges' denotes quantities of orange constituting oranges measuring one kilogram. With definite plurals, 'of' instead denotes individual parthood [ibid.]. Thus, 'of the oranges' denotes the subpluralities of oranges constituting a given maximal plurality. This might explain why, for example, 'four oranges of apple(s)' is anomalous, as pluralities of oranges are not individual or material parts of apple(s). In cases where 'of' is missing (e.g. 'one half orange'), 'half' applies directly to the noun, thus leading to a partitioning of an orange into two countable parts—the halves.

[27] By including clause (iv), we follow Ahn and Sauerland [2015: 128], who treat numerators as arguments of fraction words. This makes it possible to apply IM&R's semantics to mixed fractions such as 'two and a third', which IM&R don't consider.

(33) Mary: 'Those are two thirds of an orange.'

Of course, given RUM, we predict that '$\frac{2}{3}$ of an orange' should also have a MI. Hence the contrast witnessed in (34):

M-Context: Like the smoothie context in (7), except John's diet requires eating $\frac{2}{3}$ of an orange per day.

(34)? ? Mary: 'Those are two thirds of an orange.'

Similarly, RUM predicts that if there are six oranges worth of orange slices on the table, each from possibly different oranges, (32) will be true only if the total amount of orange that John ate measures four oranges worth.

4.3. Putting It All Together

The goal here is to informally derive Answer$_{\mathrm{I}}$, based on the analyses of cumulative conjunction and fractions just sketched:

(Answer$_{\mathrm{I}}$) There are two oranges on the table, and there is one half-orange on the table.

First, we assume that 'two' in the Answer is a cardinality modifier, one that takes a predicate like 'orange' and returns a complex predicate true of pluralities consisting of two individual oranges. Second, following the analysis in section 4.2, we assume that 'a half' (or 'one half') is also a predicate modifier. It combines with a predicate such as 'orange' to return a predicate true of individual half-oranges—namely, individuated quantities of orange partitioned into two non-overlapping quantity-uniform parts.

Since 'two' and 'a half' are both predicate modifiers, they have the same semantic type, and so can be coordinated by cumulative 'and'. The result is a complex predicate modifier—namely, 'two and a half', which can then combine with a predicate like 'oranges'. Formally, assuming Krifka [1990]'s analysis of cumulative conjunction, 'two and a half' receives the denotation in (35a), where $\sqcup$ denotes the sum-operation:

(35) a. $[[\text{two and a half}]] = \lambda P.\exists x.\exists y.\exists z.x = y \sqcup z \wedge \textit{two}(P(y)) \wedge \textit{one}\ (\textit{half}\ (P(x)))$

b. $[[\text{two and a half oranges}]] = \lambda x.\exists y.\exists z.x = y \sqcup z \wedge \textit{two oranges}\ y \wedge \textit{one half}\ (\textit{orange}\ (x)))$

Accordingly, 'two and a half oranges' will be true of pluralities consisting of two parts. The first is itself a plurality consisting of two whole oranges; the second is a single half-orange. Because 'on the table' is a distributive predicate, it will be true of a plurality only if each of its parts are on the table. Specifically, it will be true of 'two and a half oranges' only if both whole oranges and the single half-orange are all on the table.

Finally, we assume that 'there is' expresses existential quantification. Consequently, we predict the following truth-conditions for the Answer: it will be true just in case there is at least one plurality consisting of two parts, a plurality of two oranges and one half-orange, all of which are on the table. As the reader can check, these are precisely the conditions under which the Answer$_{\mathrm{I}}$ is also true. In sum, when combined in a natural way, independently motivated analyses of cumulative conjunction and fractions imply that the Answer should have an II.

We emphasize that our semantics for fractions is completely general. (36) is also predicted to have an II, for instance:

(36) There are $2\frac{2}{3}$ oranges on the table.

Following our account from section 4.2, the numerator here counts two things—namely, thirds of an orange. Consequently, (36) will be true on the II if there are *four* things on the table—namely, two whole oranges and two thirds of an orange.[28] This is significant because COP could have just as easily been formulated using '$2\frac{2}{3}$ oranges', of course. Thus, the generality of our analysis ensures that the solution that we propose is sufficiently general.

5. Conclusion: Don't KISS the COP!

We began with a simple puzzle. If we start with three oranges on the table, eat half of one, and set the remaining half back on the table, then the intuitively correct answer to the Question is the Answer:

(Question) How many oranges are there on the table?

(Answer) There are $2\frac{1}{2}$ oranges on the table.

Since half-oranges aren't oranges, it seems that there can only be two or three oranges on the table. So, how can the Answer be correct?

Our analysis of I/M-ambiguities offers an answer. Given the analysis of fractions developed above, we are clearly committed to embracing one horn of this apparent dilemma. A half-orange is definitely not an orange. So, there are only two oranges on the table, strictly speaking. But there is also a half-orange on the table. On the II of the Answer, we are individuating two whole oranges and a single half-orange, and we are counting them separately. Thus, one may truly answer the Question with the Answer even though there are, strictly speaking, only two oranges on the table.

On the MI, we are measuring two and a half oranges worth of orange. Thus, the Answer would remain true even if we were to halve the remaining two whole oranges, or to cut them and the single half-orange into slices of various sizes, or even if we were to pulverize all of the remaining orange into pulp. Although this interpretation of the Answer will always be available thanks to RUM, we have seen that it is not the interpretation relevant for COP.

Our proposed solution makes two significant theoretical contributions. First, fractional cardinalities are not required on either interpretation. On the II, we are counting three things—namely, two oranges and one half of an orange, and two and one are both whole numbers, obviously. On the MI, we are instead measuring an amount of orange-stuff, and so 'two and a half' fails to specify a cardinality. In either case, COP poses no threat to traditional analyses of cardinality expressions.

[28] RUM also predicts an MI that is true if there is an amount of orange on the table measuring two and two-thirds oranges worth. As a reviewer notes, there are some delicate morphosyntactic issues that we are brushing over here. Other fractional modifiers like 'third' are less permissive than 'half': compare 'two oranges and a half of an orange' with 'two oranges and a third of an orange', 'two oranges and half an orange' with ??'two oranges and third an orange', and 'two oranges and a half orange' with ??'two oranges and a third orange'. We will leave this to future research.

Second, the fact that the Answer is I/M-ambiguous reveals the shortcomings of both the traditional analyses and KISS. Without RUM in place, traditional analyses wrongly predict that the Answer should be false in M-Contexts—for example, when—there are $2\frac{1}{2}$ oranges worth of orange slices on the table, each from possibly different oranges. On the other hand, according to KISS, the Answer is true just in case the amount of orange-stuff on the table measures 2.5 oranges worth—that is, the MI. Hence, without a way of semantically recovering the distinction between IIs and MIs revealed in section 2, KISS would appear to wrongly predict that the Answer should only have a MI. Our conclusion is this: don't KISS the COP; have RUM instead![29]

References

Ahn, D. and U. Sauerland 2015. The Grammar of Relative Measurement, *Proceedings of Semantics and Linguistic Theory (SALT) 25*: 125–42.

Barwise, J. and R. Cooper 1981. Generalized Quantifiers and Natural Language, *Linguistics and Philosophy* 4/2: 159–219.

Chierchia, G. 1998. Reference to Kinds Across Languages, *Natural Language Semantics* 6/4: 339–405.

Frege, G. 1884. *Die Grundlagen der Arithmetik*, Breslau: Wilhelm Koebner. Trans. J.L. Austin as *The Foundations of Arithmetic* (1959). Oxford: Blackwell.

Ionin, T., O. Matushansky, and E.G. Ruys 2006. Parts of Speech: Towards a Unified Semantics for Partitives, *NELS 36: Proceedings of the 36th North East Linguistics Society*: 357–70.

Kennedy, C. and J. Stanley 2009. On 'Average', *Mind* 118/471: 583–646.

Krifka, M. 1989. Nominal Reference, Temporal Constitution, and Quantification in Event Semantics. in *Semantics and Contextual Expressions* ed. R. Bartsch, J. von Bentham, and P. von Emde Boas, Dordrecht: Foris Publications: 75–115.

Krifka, M. 1990. Boolean and Non–Boolean *and*, in *Papers from the Second Symposium on Logic and Language*, ed. L. Kálman and L. Polos, Budapest: Akadémmiai Kiadó: 161–88.

Landman, F. 2004. *Indefinites and the Type of Sets*. Oxford: Blackwell.

Liebesman, D. 2016. Counting as a Type of Measuring, *Philosopher's Imprint* 16: 1–25.

Link, G. 1983. The Logical Analysis of Plurals and Mass Terms: A Lattice-Theoretical Approach, in *Meaning, Use, and Interpretation of Language*, ed. R. Bäuerle, C. Schwarze, and A. von Stechow, Berlin: de Gruyter: 302–23.

Partee, B.H. and V. Borschev 2012. Sortal, Relational, and Functional Interpretations of Nouns and Russian Container Constructions, *Journal of Semantics* 29/4: 445–86.

Reinhart, T. and T. Siloni 2005. The Lexicon–Syntax Parameter: Reflexivization and Other Arity Operations, *Linguistic Inquiry* 36/3: 389–436.

Rothstein, S. 2017. *Semantics for Counting and Measuring*, Cambridge: Cambridge University Press.

Salmon, N. 1997. Wholes, Parts, and Numbers, *Noûs* 31/supp: 1–15.

Scontras, G. 2014. *The Semantics of Measurement*. Ph.D. dissertation, Harvard University.

Snyder, E. 2017. Numbers and Cardinalities: What's Really Wrong with the Easy Argument? *Linguistics and Philosophy* 40/4: 373–400.

Snyder, E. and J. Barlew 2016. The Universal Measurer, in *Proceedings of Sinn und Bedeutung 20* (Tübingen), ed. P. Berezovskaya and A. Schoeller: 658–75, at semanticsarchive.net.

[29] We would like to thank a number of people for very helpful conversations leading to significant improvements of this paper, including but not limited to Chris Barker, Chris Kennedy, Fred Landman, Jessica Rett, Craige Roberts, Susan Rothstein, Greg Scontras, Stewart Shapiro, and Neil Tennant.