

## REPRESENTING UNICORNS: HOW TO THINK ABOUT INTENSIONALITY

**R. M. Sainsbury**

ABSTRACT: The paper focuses on two apparent paradoxes arising from our use of intensional verbs: first, their object can be something which does not exist, i.e. something which is nothing; second, the fact that entailment from a qualified to a non-qualified object is not guaranteed. In this paper, I suggest that the problems share a solution, insofar as they arise in connection with intensional verbs that ascribe mental states. The solution turns on (I) a properly intensional or non-relational notion of representation and (II) a notion of “putting a representation on display”.

When we think about unicorns, we are thinking about something, namely unicorns. But as there are no unicorns, we aren't thinking about anything, that is, it's not the case that we're thinking about something. One feature of intensional verbs is that they seem to generate this kind of contradiction.

Such verbs also generate problems of a seemingly entirely different kind. One example is this:

How can it be that, although “She wants red shoes” entails “She wants shoes”, “She's afraid of rabid dogs” does not entail “She's afraid of dogs”?

In this paper, I suggest that the problems share a solution, insofar as they arise in connection with intensional verbs that ascribe mental states. The solution involves three elements. (I) A properly intensional or non-relational notion of representation. (II) A notion of “putting a representation on display”. (III) A study of specific intensional verbs, to explain how the representation that is put on display connects with the specific meaning of the verb.

*I/Non-relational representation*

A non-relational account of representation is one that rejects the following inference:

(R) from “ $x$  represents  $y$ ” infer “there is something such that  $x$  represents it”.

We are familiar with a non-relational notion of representation when we discuss representational painting. A painting may be *of a landscape*—as we say, it is a *landscape-painting*—without there being a landscape that it represents. It may be of a “purely imaginary” landscape. This is not a kind of landscape, for all landscapes are real. In describing a landscape as imaginary, we are simply saying that there is no such landscape, though the artist imagined that there was. We are not denying that landscapes the artist has experienced played some kind of causal role in the production. Even so, for any landscape, the painting is not, and is not supposed to be, of it; and it may well be that, for any landscape, the painting does not resemble it at all closely.

Natural as this non-relational conception of representation is for pictorial representation, it has been problematic both for perception and for language. In discussions of perception, it is associated with denials of the mind’s capacity “directly” to connect with the world. Similarly, in connection with language, it has been associated with “descriptivism”, and in turn with denials of “direct reference” theories. In fact, there is no such tension, as I’ll illustrate with the case of language. The non-relational thesis is that the mere fact that some expression  $x$  represents  $y$  does not ensure that there is something,  $y$ , that  $x$  represents. This is consistent with there being species of

## Representing Unicorns

expressions, object-involving ones, such that, for all such expressions,  $x$ , of this species, if  $x$  represents  $y$  then there is something,  $y$ , that  $x$  represents. The inference springs from the nature of some privileged subclass of expressions, and not from the very nature of representation. Representation in its own nature can be non-relational, even if, as cannot possibly be disputed, there are some things such that some representations represent them.

That the mind *represents* the world does not entail that its contact with the world is other than direct. Rather, the distinction between direct and indirect is a distinction within kinds of representation, not a distinction between representing something and accessing it in some supposed non-representational way. If John caused the mess, the representation *John* may count as a direct representation of John, and *the person who caused the mess* as an indirect representation of John. Specifying the distinction in general terms may be difficult, but here it is enough to stress that it is a distinction within kinds of representation. “Represented” does not entail “accessed indirectly”.

One should start with a non-relational notion of representation, and then consider whether some representations are relational. Words like “unicorn” and “Pegasus” will not be good candidates for relational representations; ones like “red” and “London” will be better. But there is also another question: in describing the semantic properties of relational representations, should their relationality figure in the semantics? I think there are good reasons to think not. At one point the majority of users of “witch” took it to be a relational representation. The consensus now has gone the other way. This change is not a change of meaning or semantics. On the contrary, we have to mean the

same by “witch” as the witch-hunters, when we say they were mistaken in believing that there were witches. Otherwise there is no real disagreement, merely a change of subject. It’s probably a good division of labor to spare semantic theorists the task of determining whether or not there are things that “witch” or “Homer” represent. The default option for semantic description will thus be non-relational. This is consistent, adding a relational notion as an overlay, one that might emerge by reflecting on how words are learned, or on Twin-Earth cases. Whether the overlay is regarded as belonging to “semantics” or not seems largely terminological. Given the connection between semantics and understanding, and that one can understand an expression without knowing whether there is anything it represents, it seems best to say that the overlay would not belong to semantics.

What does “London” refer to? London. What does “Pegasus” refer to? Pegasus. This seems an unimpugnably correct answer, even in a context in which it is well-known that there is no such thing as Pegasus. A surprising moral: even reference, philosophers’ preferred tool for describing word-world relations, is intensional, and so a non-relational notion.<sup>1</sup> Nothing of great moment follows, for one can define an extensional notion of reference (let us demarcate it reference\*) by using the intensional one. One step on this road would be to stipulate:

$x$  refers\* to  $y$  iff there is a  $z$  such that  $z=y$  and  $x$  refers to  $y$ .<sup>2</sup>

---

<sup>1</sup> This claim can be found in Chisholm (1957: 174–75).

<sup>2</sup> Substitutivity can be added as a further stipulation: if  $x$  refers\* to  $y$  and  $y = z$ , then  $x$  refers\* to  $z$ .

## Representing Unicorns

Should we explain the semantics of names in terms of reference\* or reference? A long tradition appeals to the former; but if you are moved by the point about witches, you should regard this as a mistake. We don't know whether there is any single person to whom "Homer" refers. Leverrier at one point falsely believed that there was such a planet as Vulcan. We go beyond what we know if we assign "Homer" reference\*. We know we would be mistaken to assign "Vulcan" reference\*. Yet the meanings of the words are stable and not in doubt, which is why we can raise questions about Homer, and can be sure that there is no such thing as Vulcan (a point on which Leverrier for a short period disagreed). We need semantic constancy between these variable epistemic states. This can be assured by saying that "Homer" refers to Homer and "Vulcan" to Vulcan, but cannot readily be assured if we have to use reference\* as our central notion in the classical way.

In Davidsonian truth theory, the semantic clauses for predicates are non-relational, for example:

for all  $x$ ,  $x$  satisfies "unicorn" iff  $x$  is a unicorn.

On some versions of truth-theory (e.g., McDowell 1977), axioms for names are different, for example

"Hesperus" stands for Hesperus,

where "stands for" is regarded extensionally. This disparity seems unjustified: names should be treated like predicates, predicates should be treated non-relationally, so names should be too.

It would not be desirable to implement this idea by simply helping oneself to a non-relational notion of reference within a

truth theory. As Davidson said, the logic of intensional notions has not been worked out, and in working it out one might indeed encounter “problems as hard as, or perhaps identical with, the problems our theory is out to solve” (Davidson 1984: 22). However, one can use well-understood negative free logic. This enables us to exploit a relational notion of reference\* to define (in effect) a non-relational notion of reference, as in this example:

(P) for all  $x$  (“Pegasus” refers\* to  $x$  iff  $x$  is Pegasus).

Since nothing satisfies “ $x$  is Pegasus”, the truth of the biconditional ensures that there is nothing to which “Pegasus” refers\*. This is consistent with “Pegasus” referring to Pegasus as, intuitively, it does. We could understand (P) as in effect saying that the word refers\* to Pegasus, if to anything.<sup>3</sup>

We have not yet surveyed, even at the most superficial level, all the resources this approach requires. The description of “Pegasus” just given by (P) is no more true than

(V) for all  $x$  (“Pegasus” refers\* to  $x$  iff  $x$  is Vulcan).

But “Pegasus” does not refer to Vulcan. To address this lacuna, we could prefix the extensional description with some intensional operator. My preference is for “It is a matter of meaning alone that”, which I’ll abbreviate as “M”. Prefixing this to (P) delivers a truth, but prefixing it to (V) yields a

---

<sup>3</sup> Worries about circularity may well arise. There is a conception of a semantic project upon which they are unfounded (see McDowell 1977). Treating names (and other referring expressions) as having irreducible semantics, as exemplified in (P) above, is recommended by Sainsbury (2005).

falsehood.<sup>4</sup> This is a basis for preferring (P) to (V) as a semantic axiom.

The overall plan is to use referential\* free logical axioms of a kind which remain true even when M-prefixed. They feature unprefixes in the theory, in order to facilitate ordinary extensional logic. In the Davidsonian tradition, it is a familiar idea that keeping to appropriate theorems may require some restrictions in the proof-theory (one should restrict attention to “canonically proved” theorems). The ultimate test for correctness is that the theorems should fit harmoniously into an explanatory description of the behavior of the users of the language.

*II/ Putting a representation on display*

One element in Davidson’s paratactic analysis of propositional attitude ascriptions is that they involve putting a representation on display. For Davidson, the way in which this is done is simply by referring\* to an utterance. An utterance is a species of representation, so a representation is put on display by being referred\* to.

I will suggest a variant on Davidson’s idea.<sup>5</sup> Sentences express thoughts, and thoughts represent states of affairs (or sets of worlds—the choice makes no difference for present purposes). When used assertively, a sentence is not merely displayed, but is used to claim, in effect, that the thought it expresses is true. When a sentence is embedded in an

---

<sup>4</sup> Davidson considered a somewhat similar idea. Exaggerating its similarity with the present proposal, he suggested that axioms of truth theories should be understood as if prefixed by “It is a law of nature that” (Davidson 1984: xiv).

<sup>5</sup> For propositional attitude ascriptions, the idea can be found in Sainsbury and Tye (2012, section 6.4). There are similarities with the defense of Davidson by Ludwig and Ray (1998).

extensional context, what matters to the truth or falsehood of the whole is what the thought represents. When a sentence is embedded in a non-extensional context like “Galileo said that”, what matters is the thought it expresses, and this thought is put on display. “Galileo said that the Earth moves” is true just on condition that Galileo said something expressing a thought suitably related to the one displayed by “the Earth moves”.<sup>6</sup>

From this account it follows that you need to understand the embedded sentence in order to understand the whole: you need to detect what thought is being put on display. Hence the embedded sentence does not meet a certain (rather crude) condition for being mentioned but not used. Perhaps we should say, following Ludwig and Ray (1998), that it is both used and mentioned.

A similar idea can be applied to intensional verbs. If I tell you that Sally is thinking about unicorns, I put a *unicorns*-representation on display, and I tell you that Sally is exercising such a representation in thought. This is neutral about what public language, if any, Sally speaks: many public languages have *unicorns*-representations, and for all I know such representations can be used by subjects who do not speak any language. The representation is non-relational: it represents unicorns even though there are no unicorns it represents.<sup>7</sup>

The notion of displaying a representation features in at least three other familiar idioms. One is in certain uses of “as” exemplified by “She thought of the arsenic as a tonic” and “He saw the wren as simply some small brown bird”. These

---

<sup>6</sup> The “suitable relation” needs lengthy spelling out – identity is often, though not always, sufficient, and hardly ever necessary (see Sainsbury and Tye 2012, section 6.4).

<sup>7</sup> Completing the present account requires specifying the conditions under which two representations both represent unicorns.



## Representing Unicorns

sentences are true only if the subject in the former exercised a *tonic*-representation, and the subject in the latter exercised a *small-brown-bird*-representation. Another idiom is what is sometimes called “mixed quotation”, exemplified by “The scratch on his new Mercedes made him ‘totally, absolutely heartbroken’”. In these cases, we are to suppose that not only is his broken-hearted mental state being affirmed, it is also said to be how he himself thought of his state. The primary function of the representation on display in the quoted words is to tell us how the subject was thinking about his emotions, rather than simply telling us what these emotions were. The thought of the person who thinks of arsenic as a tonic targets arsenic, but not in virtue of the exercise of a *tonic*-representation. Contrast with one who, on seeing a white powder, knowingly remarks: “Ah! Arsenic!” This thinker’s use of an *arsenic*-representation ensures that her thought targeted arsenic; she thought of the arsenic, that is, her thought targeted arsenic, while also thinking of it as arsenic, and so using an *arsenic*-representation.

The third familiar idiom in which something like displaying a representation occurs is found in a use of “like” in some US dialects of English. For example, from “Alice went, like, I gotta get outta here” we can infer not only what Alice said (that she needed to leave) but also how Alice put it (the representation). Even if she did not speak English, the attribution is correct only if she spoke in a colloquial way and used a standard cliché in whatever language she spoke.

The value of the idea is to be tested by detailed applications, some of which are offered in the next section. Before turning to these, I offer some comparisons. The thought that meanings are important in intensional contexts is neither new nor (so long as we keep the claim sufficiently vague) controversial. John Buridan, in the fourteenth century, said that

when a phrase like “a unicorn” is dominated by an intensional verb, it “*appellat suam rationem*” – it invokes (or “appellates”) its sense or meaning.<sup>8</sup> Appellation belongs to a complex theory involving a number of semantic primitives, including supposition, the closest to our contemporary notion of reference\*. It is noteworthy that appellation is distinct from supposition: in our terms, the way in which a representation or its meaning figures in such contexts, according to Buridan, is not by being referred to.<sup>9</sup> Frege, by contrast, suggested that in opaque contexts, words are used to refer to their customary meanings; and Davidson, while deploring the supposed lack of semantic innocence involved in Frege’s idea, took it that in such contexts an utterance was referred to, typically proleptically.<sup>10</sup> Different as these last two suggestions are in some respects, they both invoke the notion of reference as the relation whereby a meaning or utterance comes into salience. This contrasts with the present suggestion according to which the relevant relation is *displaying*.

Meanings or intensions also play a crucial role on approaches influenced by Montague, within a framework of relational representation. Intensions, functions from worlds to objects, play the role of meanings, and every expression represents by being related to an intension. An immediate problem is that “unicorn” and “centaur” have the same

---

<sup>8</sup> “[T]alia verba [viz. intensional verbs] *faciunt terminos sequentes appellare suas rationes*” (Buridan 2001: 279).

<sup>9</sup> From Buridan’s *Summulae de dialectica* (trans. G. Klima): “Appellation differs from supposition, for ... there are terms that appellate and do not supposit, for example ‘chimera’ [...]”; “terms [in intentional contexts] ... appellate their own concepts by which they signify whatever they signify” (Buridan 2001: 294 and 226).

<sup>10</sup> The attributor of “Galileo said that the Earth moves” goes on to supply a reference for “that” by uttering “the Earth moves”.

## Representing Unicorns

intension, at least on many views,<sup>11</sup> whereas some difference is required in order to register the fact that one may think about unicorns without thinking about centaurs. On the non-relational account advocated here, the difference between such expressions is registered in a very direct semantic fashion: “unicorns” refers to unicorns but not to centaurs.

### *III/ Applications*

#### *A/ Specific/unspecific*

A standard mark of intensionality in verbs is supposed to be that there are two “readings” of a sentence like this:

(1) John wants a sloop.

On one reading, (1) is true just if there is some sloop on which John has fixed his desires. On the other reading, (1) is true if John wants a sloop, but no sloop in particular. This reading, the unspecific reading, is sometimes characterized as “any old sloop will do”.

I think the data on this issue have been distorted. First, we must distinguish the unspecific case from the case marked by “any old”. It is very rare that anyone who wants an F wants any old F. John may want a sloop, and want no sloop in particular, yet not want a broken-down wreck of a sloop, or a wildly overpriced sloop, so he does not want any old sloop. “Any old” has no interesting role to play in characterizing the intensionality of desire.

---

<sup>11</sup> Kripke (1980: 156) argued that there could be no unicorns, and the arguments, if sound, would extend to centaurs. On this view, the intensions of “unicorn” and “centaur” coincide.

Secondly, it is far from clear that sentences like (1) are genuinely ambiguous. The view that they are ambiguous is encouraged by a theory according to which intensional verbs “really” (or “upon paraphrase”) don’t take ordinary noun phrases in their second position. Rather, they take sentential complements of the form “that  $s$ ”, for a complete sentence  $s$ . When thus paraphrased, there is a possibility for an indefinite in  $s$ , treated as an existential quantifier, to have either wide or narrow scope relative to the main verb (e.g. “wants”). Then the alleged ambiguity is represented by the contrast between (1a) and (1b):

(1a)  $\exists x(x \text{ is a sloop and John wants it to be the case that John possesses } x)$ .

(1b) John wants it to be the case that  $\exists x(x \text{ is a sloop and John possesses } x)$ .

There is no doubt that (1a) and (1b) differ in truth conditions and that the difference is properly characterized as a difference of scope. But we can transfer this result to (1) only on the basis of the theory according to which it *needs* to be paraphrased as (1a) or (1b). One source of the belief that (1) is ambiguous comes not directly from linguistic data, but from the application of a controversial theory. Let’s now return to the data.

If (1) is ambiguous, one should not count as having understood an utterance of it unless one has resolved the ambiguity. Yet it is plain that understanding imposes no such demands: one can know that John wants a sloop without knowing whether or not there’s a sloop upon which he has fixed his desire. (1) can be used to express what one knows, and so is neutral between the specific and the unspecific case.

The principle underlying this simple argument is that any ambiguity in “*p*” is resolved in any truth of the form “*X* knows that *p*”. For example, if *X* knows that John went to the bank, what *X* knows must be either true just if John went to the river bank or true just if John went to a financial bank. It cannot be that *X* knows that John went to the bank, but does not know whether he went to a river bank or a financial bank.<sup>12</sup> But it can be that someone knows that John wants a sloop without knowing whether there is some specific sloop that is the object of his desire.

Different readings may amount to no more than different ways in which an unambiguous sentence can be true. There is more than one way for (1) to be true. There are many ways for “Jill runs” to be true, by running north or running south, running fast or running slow, running barefoot or running in trainers. This is entirely consistent with the fact that one can understand a specific utterance of “Jill runs” while not knowing anything about the further details of some truthmaker for the sentence. It’s wrong to say that there are various “readings” of “Jill runs”, if the existence of multiple readings is something that needs to be mirrored in the semantics.

What needs to be explained is the neutrality of “a sloop” as it occurs in (1). It characterizes a desire, but remains neutral on whether or not there’s a sloop that it targets. Applying the theory of section II, the first thing to say is that (1) puts an *a-sloop*-representation on display. Now we have to say what the role of that representation is.

Here are two examples of roles the representation does not play: (a) the representation is not the “object of desire”, if an

---

<sup>12</sup> *X* might know that the sentence “John went to the bank” is true without knowing how the ambiguity should be resolved. But that is a different issue.

object of desire is an object that must be referred to in a correct answer to the question “what does John want?”. John wants a sloop, not a representation of one. (b) It’s not that his desire will be satisfied by anything of which the representation is true. That would be the “any old” case, which is rare, and is certainly not right in general.

In typical cases, the representation offers a necessary condition for the satisfaction of the attributed desire. It acts as a filter: as far as this desire goes, non-sloops are non-satisfiers.

The difference between specific and non-specific truthmakers for (1) is explained as follows. One truthmaking situation is that John’s mental state of sloop-desire contains an indefinite but no definite *sloop*-representation. That will be the unspecific truthmaker. Another truthmaking situation is that the desire contains a definite *sloop*-representation: that will be the specific truthmaker.

This initial account, as we’ll see in the coming section, offers only a sufficient condition for truth. It can be expressed thus:

X wants an F/the F/Fs if X has a desire centered on an *an F/the F/Fs*-representation; such a desire is satisfied only by an F/the F/Fs.

We’ll shortly see various ways in which this sufficient condition may fail to be necessary.

### *B/Wanting*

Can we infer that if someone’s desire for a sloop centers on a definite *sloop*-representation, the desire is also characterized by an indefinite *sloop*-representation? Yes: a subject operating with a definite *sloop*-representation must also have an indefinite

*sloop*-representation.<sup>13</sup> The “must” is not a logical “must”, but closer to a psychological one. Let’s consider only linguistic representations, and only languages which, like English, use separate lexical elements to mark indefiniteness and definiteness. In these cases, a definite *sloop*-representation is semantically complex, involving “sloop” and something like “the” or “that”. One who understands “the” or “that” understands “a”, and so understands “a sloop” and so has an indefinite *sloop*-representation. Typically, the indefinite representation will be psychologically active in behavior controlled by the subject’s *sloop* desires. If the definite *sloop*-representation is “the sloop due to be auctioned tomorrow”, a desire controlled by the representation will be one which will generate a response on these lines to the suggestion that the subject take an interest in a ketch: “I don’t want that! It’s not even a sloop.” On this picture, the inference from “She wants the sloop due to be auctioned tomorrow” to “She wants a sloop” is reliable, but not logical.<sup>14</sup>

A *sloop*-representation is not a representation such that there is some sloop it represents. This condition would be both too strong and too weak. Too strong, since there may be nothing a *sloop*-representation represents, and too weak, since someone may represent something that is in fact a sloop without exercising a *sloop*-representation. Someone who stands in the relation of desire to a sloop may not have a *sloop*-representation of any kind. To illustrate: suppose that *The Mary Jane* is a sloop that John desires, but without knowing that it is a

---

<sup>13</sup> A *sloop*-representation is not merely something that represents a sloop. It must represent something as a sloop.

<sup>14</sup> These issues are likely to take a different form for thinkers who use languages without the kind of definite/indefinite markers available in English.

sloop. Then he has a definite representation that is in fact of a sloop. He might have no *sloop*-representation at all, yet he might still want *The Mary Jane*. Given that *The Mary Jane* is a sloop, it doesn't sound wrong to say he wants a sloop—the thing he wants is a sloop, even if he doesn't know it.

Theorists sometimes contrast “needs” with “wants” on this issue. For example, Graeme Forbes (2010: section 5) says that whereas the inference from “She wants a glass of water, water is H<sub>2</sub>O, so she wants a glass of H<sub>2</sub>O” is not valid, the corresponding inference with “needs” replacing “wants” is valid. On this view, we should simply tough it out, denying that John wants a sloop.

Although it might be nice to have a language with such a clear-cut distinction, English is not like that. Rather, we accept ascriptions of desires which put on display representations not used by the desiring subject (likewise for other intensional verbs for mental states). We accept that John wants a sloop, even though he has no concept of a sloop, no *sloop*-representation. We have to allow such ascriptions as correct, unless we are to be severely inhibited in our ascriptions of mental states. For example, we all want to allow that the dog Fido wants his bone and is looking for it under the sycamore tree. Yet we certainly do not wish to commit to the view that Fido has a *my bone*-representation or a *sycamore tree*-representation. Fido represents his bone and the tree somehow or other. But it would be foolish to have any confidence in the opinion that he exploits in his desires and searches the very representations that we exploit in our ascriptions of these states.

For a clear example among language-users: suppose you see a gnu at a zoo and sincerely say “I want that”, yet you have no *gnu*-representation. (A *gnu*-representation is not just something that represents gnus; in addition it represents them



## Representing Unicorns

*as gnus.*) Later in the day, your desire cannot be reported using the very representation that you used, since we now have no demonstrative access to the gnu you saw. You can be reported as wanting a gnu, or as wanting the gnu you saw, even though your desire has never involved any sort of *gnu*-representation. We must allow this mismatch between the representation used in the attribution and the representation at work in the subject's mental states, on pain of making it impossible to give a correct report of those states.

This suggests that a necessary and sufficient condition for truth of a sentence "X wants an F/the F/Fs" is on the following lines (for some relation R among representations):

X has a desire centered on a representation, Z, which is R-related to an *an F/the F/Fs*-representation; such a desire is satisfied only by a satisfier of Z.

Identity is a special case of the R-relation, providing a sufficient but not a necessary condition. In many cases, R obtains when the *an F*-representation (used in the ascription) is not Z (the representation in the subject's mental state) but is true of that of which Z is true. For example, R holds between the *The Mary Jane*-representation (Z) and an *a sloop*-representation when John wants *The Mary Jane*, which is a sloop though he does not know it (he has no *a sloop*-representation). This representation in John's mind, in this context, verifies the ascription "John wants a sloop".

The value of R is highly sensitive to context. The co-reference\* condition mentioned in the previous paragraph is in many contexts insufficient. For example, Jane's desire for a unicorn cannot be reported as a desire for a centaur, even

though an *a unicorn*-representation is true of just the things of which any *a centaur*-representation is true, that is, nothing.

In typical cases, straightforward inferences involving “wants” are easily accounted for. If Mary wants red shoes, she typically has a desire involving a *red shoes*-representation (here R is identity), which acts as a necessary condition for the satisfaction of the desire. If she were to be sorting through candidates for satisfying this desire, she’d reject things that are not red shoes; a fortiori, she’d reject things that are not shoes. In this case, compositionality in the *red-shoes*-representation ensures that she possesses a *shoes*-representation, and that she exercises it in her desire.

Preservation of such weakening inferences places a structural constraint on R. Let’s represent adjectival modification by “+”, and suppose that John has a desire centered on a representation of the form:  $Z_1 + Z_2$ . Then he has a desire centered on a  $Z_2$  representation. Suppose that  $Z_1 + Z_2$  is R-related to  $X_1 + X_2$ , so that we can correctly report his Z-desire as a desire for an  $X_1 + X_2$ . Given weakening, he desires an  $X_2$ . It follows that  $Z_2$  is R-related to  $X_2$ . To illustrate with an example. Suppose John uses a *red shoes*-representation in his desire. In many circumstances we can report him as wanting fire-engine colored footwear. In the case, his *red shoes*-representation is R-related to the *fire-engine colored footwear* representation. Since it follows that he wants footwear, it also follows that his *shoes*-representation is R-related to the *footwear*-representation.

The present proposal, centered on the notion of putting a representation on display, does not apply to all intensional verbs. It does not apply to those which, like “needs”, may not involve a mental state. In these cases, a report is not putting on display a representation that is (R-related to one) used in a truthmaking mental state. For example, “needs” allows non-

## Representing Unicorns

specific indefinites, and so may be classified as intensional, even when there is no question of a contentful mental state being at issue, as in “The lily needs a stake”. This can be true even if there is no stake it needs. The theory proposed here evidently cannot explain such cases: no representation is put on display. The phenomenon is, however, explicable drawing on familiar resources. Let’s say that a sentence is non-specific iff it contains an indefinite like “a F”, and can be true even if nothing is F. Negation and modals can evidently induce non-specificity as in:

It doesn’t have a stake.  
It ought to have a stake.

Needing is not having something one ought to have, and so its meaning already contains (twice over, in fact) material that explains the non-specificity without appeal to putting a representation on display.<sup>15</sup>

### *C/Fearing*

Mary can fear rabid dogs without fearing dogs. Weakening holds for the example of the red shoes, but fails in the example of the rabid dogs. The explanation is that, in the case of an ascription of fear, the representation that’s put on display has a different job to do: it has to provide a *sufficient* condition. Having a representation of rabid dogs nearby, *as* rabid dogs nearby, should normally be enough to trigger fear in one who

---

<sup>15</sup> Other verbs belong with “needs”, like “buys”, “orders” (as in “ordered a bottle of wine”) and perhaps even “resembles” (“the cloud resembles a unicorn”) in having non-specific truth makers, but for which it is at least not obvious that the “displayed representation” approach will be appropriate. The taxonomy in this area remains frail.

fears rabid dogs. It does not follow that a representation of dogs nearby, *as* dogs nearby, should normally be enough to trigger fear in one who fears rabid dogs. By the same token, strengthening inferences will be valid: one who fears dogs fears rabid dogs too.<sup>16</sup> As before, these facts place structural constraints on the R-relation.

*D/Quantifiers*

The inference from thinking about unicorns to thinking about something is correct. What is incorrect is to move to there being unicorns that are thought about. It would be nice if this was a simple matter of scope, governed by the rule: never apply external (wide-scope) existential quantification to the intensional place of an intensional verb. But that does not seem quite right, for it does seem acceptable to infer from “Mary is thinking about several unicorns” to “There is something Mary is thinking about—namely, several unicorns”. The unacceptable conclusion is that there are several unicorns she is thinking about.

The singular quantifier form is appropriate in inferences from a premise involving the plural “unicorns”. Similarly, an impersonal quantifier form can be appropriate in an inference from a premise involving a personal noun phrase: from “John and Peter want the same thing, namely a wife” we can properly infer “There’s something John and Peter both want, namely a wife”. We cannot properly infer “There’s somebody/some woman John and Peter both want”. These features make the

---

<sup>16</sup> How is “too” functioning here? The rabid dogs are dogs, so no further dogs are being adduced. What is added is a new representation.

## Representing Unicorns

“there is” expression “special” in some way.<sup>17</sup> It’s not functioning as an ordinary quantifier phrase.

An explanation is that “there is” is picking up on the displayed representation. There is a representation Mary is exercising in thought, a *several-unicorns*-representation. The representation in its nature is singular, even though plural in what it represents. This is not a complete explanation, for it’s not that Mary is thinking about a representation. Rather, she is thinking about unicorns, which are not representations. The relevant fact is that she is thinking *with* a representation, indeed, with a *unicorns*-representation; but there’s no overt syntactic mark of this switch from “of” to “with”. When “unicorns” occurs after “thinks about” it tells you what she’s thinking about. When “something” occurs in “There is something she’s thinking about”, the *thing* in question is something she’s thinking with, a representation, not something she’s thinking about.

Similarly, John and Peter’s desires both involve an *a wife*-representation. What they want is not the representation, but something it represents. Their desire uses that representation: they desire with it, even though not for it, rather as Mary thinks about unicorns with her *unicorns*-representation.

In some cases, just the right contrast is explicitly marked. For example, in

He thought of a sloop *as* a sloop.

the two occurrences of “a sloop” have manifestly different roles, the first to tell you what he thought about, the second to tell you how he thought about it, that is, what representation he

---

<sup>17</sup> These subtle distinctions first came to my attention in Moltmann’s work (1997). She calls such phrases “special quantifiers”.

used to think about it with. The second occurrence puts the relevant representation on display, a token of the very representation-type attributed to the thinker. We do not always separate these roles so clearly, whether we are using the more familiar intensional transitives, or are describing human action more generally – as suggested in the next section.

We are now finally in a position to explain away the contradiction offered in the first paragraph of the paper. We could reconstruct the argument for the contradictory propositions as follows:

1. We are thinking about unicorns (assumed).
2. We are thinking about something (from 1).
3. There is something we are thinking about (from 2).
4. There are no unicorns (assumed).
5. Hence there are no unicorns we are thinking about (from 4).
6. Hence there is nothing we are thinking about (from 5).

(3) and (6) are inconsistent, but seem to follow from the indisputable assumptions (1) and (4). We must accordingly show that one of the inferences is invalid.

One candidate is the inference from (2) to (3). But in fact, as just discussed, this inference is valid. (We need to be sure to distinguish the conclusion from “There are unicorns we are thinking about”.) The remaining candidate is the inference from (5) to (6). The considerations of this paper show, perhaps surprisingly, that this is invalid. The most natural reading of the conclusion treats the “there is” as a special quantifier, one that relates to a representation. Even if there are no unicorns we are thinking about, we can be exercising a *unicorns*-representation in thought, and so be thinking about unicorns,

## Representing Unicorns

and so there be something we are thinking about, namely unicorns.

### *IV/ The larger picture: actions*

Our descriptions of our actions are subject to a similar kind of dual interest: sometimes we are concerned primarily with the representations the agents exploit in guiding what they do, sometimes primarily with the world with which they engage thanks to the representations they possess. We move between these interests with complete fluidity, and in ways that can easily be missed. Consider how the sentence

(1) They set out for Florence.

might be evaluated in each of the following scenarios:

\*We are trying to trace a missing group of tourists and we want to know where we should look. If we know that (1) is true, we should check out the road to Florence.

\*We are studying a group's psychology and powers of deduction. We leave various clues and messages in their environment. They pass the test if they figure out they should head for Florence. We know that if (1) is true they have passed the test, even if, due to an error irrelevant to our experiment, they are on the Rome road, wrongly thinking it is the road to Florence.

In the first scenario, we show a preference for the worldly side of the action, that is, with what the representation we use in the attribution represents; in the second case we show a preference for the representational side, that is, with the nature of the representation that guided the action we attribute. This corresponds to the tension we felt about whether John, who wants the sloop *The Mary Jane* without knowing what a sloop is,

wants a sloop. If we are concerned with how he relates to the world, we are inclined to say he wants a sloop. If we are concerned with how he represents the world, we may be inclined to say he doesn't want a sloop. (As would be expected on the theory advanced here, the denial is more acceptable when heavy emphasis is placed on "sloop".)

The phenomenon is widespread. Someone is counting chickens and reaches the total 13. If we are primarily interested in this action from the worldly point of view, we will regard

(2) She counted 13 chickens.

as true only if there were 13 chickens to be counted. If we are more interested in her and her mental states, and are thus concerned with the nature of the subject's representations, we will treat (2) as true if she pronounced 13 to be the number of chickens, even if there were more or fewer chickens.

Another familiar example concerns the satisfaction of desires. There is room for two opinions about whether someone kidnapped in Columbia and forcibly transported to work on a US tomato farm has thereby satisfied her desire to travel. She traveled, and that's what she wanted; from a worldly perspective, her desire was satisfied. But of course she never wanted to travel under those conditions; this was not how, in her desire, she represented traveling.

Intentions involve representations, and the same worldly thing or state can be represented in different ways. (As Davidson famously said, it's one thing to fly one's spaceship to the Evening Star, another to fly it to the Morning Star.) Moreover, a representation may fail to represent anything. Those are the platitudes needed to explain the contrast between the divergent perspectives we may take on actions. If we are



## Representing Unicorns

interested in the intentions, we are dealing with representations; representational grain is finer than worldly grain and may have no worldly correlate.

The upshot is that the kinds of considerations needed to make sense of intensional transitive verbs are required to make sense of our ascriptions of actions. Being human—acting, thinking, wanting or fearing—involves exercising representations. In describing these human states, we may well display the representations the agents used, and we transition seamlessly between a concern for what these representations were and a concern for the worldly situation in which they guided thought, emotion and action.<sup>18</sup>

University of Texas at Austin  
marksainsbury@austin.utexas.edu

## REFERENCES

- Buridan, Jean. 2001. *Summulae de Dialectica*. Ed. and trans. Gyula Klima. New Haven: Yale University Press.
- Chisholm, Roderick. 1957. *Perceiving: A Philosophical Study*. Ithaca: Cornell University Press.
- Davidson, Donald. 1967. "Truth and Meaning." *Synthese* 17 (1): 304–23. Reprinted in Davidson 1984: 17–36.
- — —. 1984. *Essays on Truth and Interpretation*. Oxford: Oxford University Press.
- Forbes, Graeme. 2010. "Intensional Transitive Verbs." Online at <http://plato.stanford.edu/archives/spr2010/entries/intensional-trans-verbs/>.

---

<sup>18</sup> Material from this paper was presented at a colloquium in the Linguistics Department, University of Texas at Austin. My thanks to participants for comments, especially David Beaver and Alex Grzankowski. My thanks also to Petr Kofátko for comments on an earlier draft.

R. M. Sainsbury

- Kripke, Saul. 1972/1980. *Naming and Necessity*. Oxford: Basil Blackwell.
- Ludwig, Kirk and Greg Ray. 1998. "Semantics for Opaque Contexts." *Philosophical Perspectives* 12 (12): 141–166.
- McDowell, John. 1977. "On the Sense and Reference of a Proper Name." *Mind* 86 (342): 159–185.
- Moltmann, Friederike. 1997. "Intensional Verbs and Quantifiers." *Natural Language Semantics* 5 (1): 1–52.
- Sainsbury, R. M. 2005. *Reference without Referents*. Oxford: Clarendon Press.
- Sainsbury, R. M., and Michael Tye. 2012. *Seven Puzzles of Thought and How to Solve Them: An Originalist Theory of Concepts*. Oxford: Oxford University Press.