Constraining Meaning – Category Theoretical Perspectives in Semantics

Miriam S. Schmidt
University of Göttingen

1 The Phenomena – *that*, but not *why*

Attitude verbs like *believe, know*, or *regret* embed a proposition expressing the content of the attitude. A sentence 'x believes that p', e.g., is usually depicted as true if all of x's doxastic alternatives are included in p (cf., e.g., Kratzer 2006). Following a Hintikkian (1969) possible world semantics, p is understood to be the set of worlds in which p is true. However, there are various scenarios in which this analysis of intensionality predicts solutions that aren't fine-grained enough or unintuitive.

(1) CONTEXT: Yoda believes the proposition p. He has a logically well-trained mind and knows that p and $p \wedge (q \rightarrow p)$ are logically equivalent. Ahsoka, young and stubborn, also believes p, but thinks that $p \wedge (q \rightarrow p)$ is saying something different and is skeptic about it.

a. Yoda believes that p and $p \wedge (q \rightarrow p)$.
b. Ahsoka believes that p and $p \wedge (q \rightarrow p)$.

(2) CONTEXT: Theseus' starship (TS) flies through the galaxy and exchanges all of its parts during travel due to repair work. Watto, the junk dealer, quickly detects all repairs and does not recognize the identity of the ship that just arrived on Tatooine (T). The Jedi Obi-Wan is simply happy about the return of Theseus and his crew. He recognizes the identity of the ship.

a. Watto correctly believes that Theseus' starship did not arrive on Tattoine.
b. Obi-Wan correctly believes that Theseus' starship arrived on Tatooine.

(3) CONTEXT: Anakin once was a famous Jedi and general in the Clone Wars. Then he turned to the dark side and, using dark forces, grew even more powerful. For Rey, being a hero is to be powerful and fierce. For Kylo, being a hero is to be courageous and kind-hearted.

a. Rey believes that Anakin was a hero.
b. Kylo believes that Anakin was a hero.

Analytically, Yoda and Ahsoka have logically equivalent beliefs in (1), so both a. and b. should have the same truth value. Yet we know that beliefs are highly subjective and we would like to grant Ahsoka her skepticism about $p \wedge (q \rightarrow p)$ while (inconsistently) believing p at the same time, thus allowing sentences a. and b. to have different truth values, despite the fact that the corresponding possible worlds are the same. Sentences a. and b. in (2) are a similar, yet not logical, but conceptual, version of the problem. The possible worlds analysis predicts, as the embedded propositions in a. and b. together are logically inconsistent, that $\{w \mid \text{TS arrived on T in } w\} \cap \{w \mid \text{TS did not arrive on T in } w\} = \emptyset$. Nevertheless, we are still inclined to grant both beliefs a label of truth to concede people different conceptualizations. This, of course, arises not only for the inconsistent propositions, but also for consistent, or even intensionally identical propositions as in (3), where the classical analysis predicts no difference, but applying Rey's and Kylo's concepts, the beliefs differ hyperintensionally.
Moreover, it has been observed that some of these attitude verbs, which also allow a Content DP, show different patterns of entailment. (4) holds for *believe*, but not for *know* (cf. Djärv 2019, Elliott 2020).

(4) a. $\models$ b.

a. Yoda *attitude* {the story, the claim,...} that p.
b. Yoda *attitude* that p.

But it has not yet been noted that (4) only holds for *believe*-type verbs iff p has the same subject-matter (cf. Yablo 2014, example (1)), uses the same conceptualizations (examples (2) and (3)), and has the same structural reading (*de dicto* vs. *de re*) in both the DP- and the CP-version.

Previous work in semantics has generally focused on analysis based on type-driven selection and construction, operating on a definition of propositions as sets of possible worlds. Based on Lewis 1979, both Kratzer 2006 and Moulton 2009 include a logophoric compatibility of possible worlds with the belief of a subject in the meaning. As compatibility is based on possible worlds, the notion is too coarse-grained for the phenomena described above: the worlds are identical for (1a) and (1b), as well as for (3a) and (3b). Moreover, this notion of compatibility focuses on the content as an absolute unit only; therefore, it does not capture the structural relationship towards worlds, i.e., whether something in a belief is understood *de re* or *de dicto*. Interestingly, Djärv 2023 discusses the importance of sources when talking about beliefs, but only for combinatorial reasons and does not acknowledge how content must be anchored to a source in multiple dimensions. Kratzer 2022 proposes a hyperintensional account giving up compositionality in intensional contexts, meanwhile, there is no overall indication in semantics that the core problem consists of compositionality.
Some finer-grained analyses are also available, one strand making use of a (infinitely) rich type system (e.g., Cooper 2023), another depending on (a variation of) Truthmaker Semantics (cf. Yablo 2014, Fine 2017, or Moltmann 2020a, Moltmann 2020b), which, to some extent, capture the difference in (1), but not in (2) or (3). None of them discusses *de re* vs. *de dicto* readings for content nouns. Flexible scoping for capturing intensional effects as in (2) and (3) have been provided by Giorgolo & Asudeh 2011 and Asudeh & Giorgolo 2020, but in turn this approach misses the difference between (1a) and (1b), i.e. the distinction of subject-matter discussed by Yablo 2014.
What all of the analyses of (hyper-)intensionality in attitudinal contexts have in common, is a transfer of additionally needed complexity to the input attitudinal functions then operate on. Currently there is no approach investigating a shift of that complexity into computations rather than to inputs.

The common link in formality between the phenomena and the related complex of problems has hitherto been overlooked – using inherently extensional functions to model inherently intensional phenomena. The extensionality of functions guarantees that $f(2+2) = f(8-4)$, but does not reflect that the steps of the computation might differ. Extensional functions only tell us *that*, not *why*; but sometimes ways of being do indeed matter, which extensionality cannot reflect. Semantic functions understood on a set-theoretic basis exhibit the same shortcomings: Meaning components like *hero* or *Rey* are consumed by semantic functions and, as arguments, out of each other's scope; thus, a relational analysis including the 'context' in a computation is not attainable extensionally. Truth-determining functions are the extremal point in that perspective. An output is 0 or 1, not reflecting the steps of the computation, i.e. not reflecting the ways why a proposition is true.
Human computations and reasoning can probably be assumed to be fundamentally different – humans would not only know a result to be equal to another, i.e., that $2+2=8-4$ and therefore $f(2+2) = f(8-4)$, but they would also know that the steps of the computations are different. The same applies to language and computation of meaning, as speakers do not forget previously computed arguments of meaning functions, but can access them again, and are able to use them to reason why something is true. W.r.t. beliefs, we might, for example, want to have a flexible way to include the subject of an attitude into the scope of semantic computation of content whenever people have different concepts of things. A thorough semantic analysis should therefore reflect not only that something is true or false but also why.
For obtaining flexibility for including a contextual scope, on the basis of Mazur 2008, Król 2019 and Asudeh & Giorgolo 2020, we follow the considerations in Peruzzi 2006 of a set-free intensionality and argue that semantic analysis benefits from an epistemic notion of truth based on category theory

rather than a realist one based on set theory. By abstracting from sets, we are able to formalize constraints on the relationship between semantic elements and their environment flexibly, which is necessary to formally reflect a (hyper-)intensional interpretation while maintaining compositionality. This constraint-based semantics allows for a novel analysis of (hyper-)intensionality, where complexity is put on the computations themselves rather than on the input of these.

2 Abstracting away – Semantics without Sets

Semantics based on categories[1] rather than sets implies a shift of perspective: "A consequence of the Yoneda lemma implies that a fixed object in this category is determined up to unique isomorphism by the totality of its relationships to all other objects in the category." (Bradley, Terilla & Vlassopoulos 2022: p. 555) Therefore, the focus is more on how objects are related than on what objects are. As category theory is just an abstract description of structures such as sets (and everything based on that), the same propositions can be formed without loss of expressivity.[2] As a consequence, truth is conceptualized epistemically, since the Yoneda embedding not only states that, but also explains why, in a linguistic case, expressions are true or false. Figure 1 shows a small example for the category of linguistic meaning, adapted from Asudeh & Giorgolo 2020: p. 24.

$$\begin{tikzcd}
\mathsf{0} \arrow[r, shift left, "anakin"] \arrow[r, shift right, "rey"'] \arrow[d, shift right, "false"'] \arrow[d, shift left, "true"] & \mathtt{e} \arrow[dl, "hero"'] \arrow[d, "\tilde{\mathtt{e}} \times id_{\mathtt{t}}"] \\
\mathtt{t} & \mathtt{t}^{\mathtt{e}} \times \mathtt{e} \arrow[l, "\mathtt{eval}"']
\end{tikzcd}$$

Figure 1: Linguistic category

As in standard semantics, we assume types e and t, each represented by an object in the category. The initial object 0 acts as an ontology, saying what there is via morphisms to other objects of the respective type. As expected, the morphism *hero* then acts as a 'function' from e to t. For visual simplification, composite morphisms[3] and identity morphisms[4] are not depicted. Interestingly, in category theory, there is no difference between common objects (e-ntities, t-ruth values, b-eliefs) and expontential/function objects (e.g. $\mathtt{t}^{\mathtt{e}} \times \mathtt{e}$, the object describing the application of a function from e to t to an object e). Thus, everything can be described via morphisms between objects. A sentence *Anakin is a hero* then is true iff the morphisms *true*, *hero* ∘ *anakin* and $\mathtt{eval} \circ\ \tilde{\mathtt{e}} \times id_{\mathtt{t}} \circ$ *anakin* all commute. Truth or falsity (the object t) are then always epistemically explained by the morphisms providing the context of why and how something is true or false.

To include attitudes like *believe*, we follow considerations in Moltmann 2020b and introduce an object of type b for beliefs. Moreover, like tuples in sets, we will also look at product objects like (e,b), for which π_i is a projection morphism depicting the i-th component of the tuple. Figure 2 now shows a

[1]Cf. Mac Lane 1998 or Riehl 2016 for a general introduction to category theory or Asudeh & Giorgolo 2020, Ch. 3, for a brief, but linguistically oriented overview.

[2]According to the Curry-Howard-Lambek correspondence, we do not add extra power to the system than a simply typed lambda calculus has, if we restrict our category to be cartesian closed, which is assumed for the linguistic category at hand.

[3]For every two morphisms $f : a \to b, g : b \to c$, there is a composite morphism $g \circ f : a \to c$.

[4]For every object c, there is an identity morphism $id_c : c \to c$. Linguistically, this corresponds to self-identity as in the equative copula use in *Anakin is Anakin* or sentences like *Anakin is self-identical*.

portion of the category of linguistic meaning expanded for attitude verbs as *believe*:

$$
\begin{tikzcd}
\mathtt{0} \arrow[rr, bend left=5, "anakin"] \arrow[rr, bend right=5, "rey"'] \arrow[dr, "{(rey,belief)}"] \arrow[drr, "belief"] \arrow[dd, bend right=5, "false"'] \arrow[dd, bend left=5, "true"] & & \mathtt{e} \\
& \mathtt{e} \times \mathtt{b} \arrow[ur, "\pi_1"] \arrow[r, "\pi_2"'] \arrow[dl, "believe"'] \arrow[d, "(\tilde{\mathtt{e}} \times \tilde{\mathtt{b}}) \times id_{\mathtt{t}}"] & \mathtt{b} \arrow[d, bend left=5, "anakin^*"] \\
\mathtt{t} & \mathtt{t}^{\mathtt{e} \times \mathtt{b}} \times (\mathtt{e} \times \mathtt{b}) \arrow[l, "\mathtt{eval}"'] & \mathtt{e}^* \arrow[u, bend left=5, "hero^*"]
\end{tikzcd}
$$

Figure 2: Linguistic category with beliefs

The morphism *belief*, stating that there is a belief, can be related to the believer Rey, as the morphisms *belief* and $\pi_2 \circ$ (*rey*,*belief*) are commuting iff Rey is believing something.[5] That belief in turn is then characterized by the (composite) morphism $hero^* \circ anakin^*$, thus representing the belief in (3a). The object $\mathtt{e}^*$ is taken to be the type of entities believed to exist, which prevents the system to classify $anakin^*$ with an (actual) existence presupposition.
Lastly, it is worth noting that the approach itself does not take an absolute stand on which types to assume. Contrarily, taking the category to be a sceletal one, category theory can help identifying which objects are needed, as isomorphic objects would then be collapsed into one object anyway.

3 Constraining Meaning

The category of linguistic meaning as in Figure 2 caters only regular function application. To distinguish more nuanced meanings as in (hyper-)intensional contexts, we now add constraints on morphisms and their compositions. For all phenomena discussed above, we show that the context of objects, i.e., their Yoneda embeddings, is enough to attain a very fine-grained analysis without the need for additional complexity in types.

Extensional vs. intensional contexts can linguistically be discerned by testing for robust reference. If both Yoda and the storm trooper punch Anakin, they punch the same entity and therefore both punch Darth Vader (given that they are the same person). In the context of *believe* though, this identity cannot be assumed because Yoda might already know about Anakin being Darth Vader, while the storm trooper lacks that information and only knows that Darth Vader is a Sith, not Anakin.

(5) a. Yoda punches Anakin.
b. The storm trooper punches Anakin.

(6) a. Yoda believes that Anakin is a Sith. (*Anakin=Darth Vader*)
b. The storm trooper believes that Anakin is a Sith. (*Anakin≠Darth Vader*)

[5] Note that for visualization *believe* is not curried to take one object after another, but the existence of expontential objects allow for that as well.

That difference can then be described with a simple constraint:

$$(a_1, x) \neq (a_2, x) \Rightarrow \pi_2 \circ (a_1, x) = \pi_2 \circ (a_2, x) \qquad \text{extensional case}$$
$$(a_1, x) \neq (a_2, x) \nRightarrow \pi_2 \circ (a_1, x) = \pi_2 \circ (a_2, x) \qquad \text{intensional case}$$

The constraint guarantees preserved identity of objects in the extensional case, while in intensional contexts that identity might be lost. Even stronger, most of the times intensional contexts can be captured best by *monic* morphisms:[6]

$$(a_1, x) \neq (a_2, x) \Rightarrow \pi_2 \circ (a_1, x) \neq \pi_2 \circ (a_2, x)$$

The constraints thus predict when to expect intensional effects that can be captured by the monadic mechanisms provided by Asudeh & Giorgolo 2020. Thus, we expect the morphisms *anakin* and *anakin** to be of different meanings. The following constraint describes the effect of potential subjectivity on objects in intensional contexts like belief.

Conceptual meanings are found exactly in these contexts that are to be interpreted with respect to a subject potentially diverging from a standard extensional interpretation. E.g., the morphisms *anakin**, *hero** and their composite are expected to be interpreted according to Rey's perspective. The following constraint captures that:

$$c^* \circ \mathit{belief} = c^* \circ \pi_2 \circ (\mathtt{e}, \mathtt{b}) \Rightarrow c^* = c_{\mathtt{e}}$$

c^* is the (composite) morphism representing a meaning term in an intensional context (e.g., *Anakin is a hero* for *believe*, *a starship* for *seek*). These meaning terms adhere to the intensional ontology $\mathtt{e}^*$ and not the regular extensional ontology given by the initial object $\mathtt{0}$. Again, the tuple $(\mathtt{e}, \mathtt{b})$ indicates a pair of a logical subject $\mathtt{e}$ and the object $\mathtt{b}$ of an intensional predicate (here: *believe*). c_e is the monadic interpretation of c^*, where c^* is interpreted with a perspectival index of that subject e.
In summary, the constraint states that whenever a belief with the content c^* is identified as the belief of a subject $\mathtt{e}$, that content must then be interpreted from to the perspective of that subject. Thus, we can flexibly discern conceptual meaning differences for any scenario comparable to (2) or (3).

Subject-matter as defined by Yablo 2014 is constructively given by morphisms and their composite morphisms. Possible QUDs are given by collecting the corresponding non-computational morphisms (of any daughter of a syntactic tree) in a set, e.g. $\{$*anakin,hero,hero* $\circ$ *anakin*$\}$ as possible subject-matters for the sentence *Anakin is a hero*. Therefore, we can easily differentiate the beliefs and respective truth evaluation in (1). This could be modeled in the semantics directly by using a `writer`-monad or in an interface to pragmatics (and syntax) via a projectable SUBJECT-MATTER feature.
Interestingly, both approaches would allow for a natural notion of parthood of content (cf. Fine 2013, Moltmann 2017) without any additional semantic 'machinery'. Sentences a. and b. in (1) are extensionally equivalent and have the same subject-matters, but only the part p of the proposition is believed by both Yoda and Ahsoka, as it is the subset of the subject-matters that both evaluate as being true, i.e., for both Yoda and Ahsoka, there could be a truthmaker for p, but not necessarily for the rest of the subject-matter.

$$\text{SUBJECT-MATTER} \quad \{p, q, \rightarrow \circ(q, p), \wedge \circ (p, \rightarrow \circ(q, p))\} \qquad \text{Subject-matter for (1)}$$

The morphisms and their composition then naturally derive potential parts of propositions that can be believed (or not believed) individually. The analysis therefore also predicts non-trivial parts of content, where PWS alone can just rely on supersethood.

[6] A *monomorphism* (or a *monic* morphism) $f : x \rightarrow y$ is a morphism s.t. for any parallel morphisms $h, k : w \rightrightarrows x$, $f \circ h = f \circ k$ implies $h = k$, which is the contraposition of the statement above.

Factivity in category theoretical perspective is the necessary relation of content with the object t via a *true*-morphism. The constraint should therefore guarantee the entailment of embedded content regardless of the truth of the matrix predicate. The necessary truth of the embedded proposition can be captured by a constraint on all factive verbs f:

$$\begin{aligned} & f \circ (\mathtt{e},\mathtt{k}) = \mathtt{eval} \circ h \circ c^* \circ \pi_2 \circ (\mathtt{e},\mathtt{k}) \\ & \Rightarrow \exists c \in Hom(0,\mathtt{t}) : \eta(c) = c^* \& \ c = true \end{aligned}$$

$h = \tilde{b} \times id_t, \eta =$ unit for `reader`-monad

Any factive predicate f, e.g., *know* or *regret*, that is applied to a tuple of a logical subject e and an intensional object, here depicted as k, the content c^* describing that object must be met by a composite morphism c which commutes with the morphism *true*. In effect, this ensures that content of factive predicates is necessarily true, independently of whether f applied to $(\mathtt{e},\mathtt{k})$ is true or not.

Entailment in Content DPs as in (4) for *believe*-type verbs and the blocked entailment for *know*-type verbs can be traced back to *believe* being relatively transparent towards its content, while *know* isn't:

(7) a. Yoda believes the story that p.
b. There is a story that p and Yoda believes that p.

(8) a. Yoda knows the story that p.
b. There is a story that p and Yoda knows that p.

A *belief* is transparent towards the content of a Content DP like *story*, which is reflected in both a. $\models$ b. and b. $\models$ a. for *believe*, but not for *know*.[7] This transparency in *believe* then leads to a simple formal description of the entailment behavior: *belief* and *story* are isomorphic w.r.t. transparency of content, while *knowledge* and *story* are not.[8] Thus the isomorphism 'collapses' *story* and *belief* into a single object in the category, and therefore predicts the entailment of content in both ways (a. $\models$ b. and b. $\models$ a.). As there is no such isomorphism between *knowledge* and *story*, there is also no prediction of any entailment between the DP objects and their associated contents. This also ensures that for *believe*-type verbs, we do not only get entailment of content, but also that this content is believed with the same conceptual interpretation, has identical subject-matters, and the same structural reading (*de dicto* vs. *de re*) for both the CP and the Content DP version. This follows immediately from the objects being isomorphic. The entailment behavior of different attitudes can therefore categorically be explained via their relations to other objects, i.e., once again, via their Yoneda embedding.

De re vs. de dicto readings could be modeled in two different ways using a category theoretical semantics. The first option would be to stick to interpreting the distinction as a structural ambiguity, with quantor-morphisms (type t^e) in different positions w.r.t. to the propositional content p:

$(e, t^e \circ (h \circ p))$	*de re*, quantor outside of belief
$(e, p \circ (t^e \circ h))$	*de dicto*, quantor inside of belief

A non-structural solution would be to interpret de re vs. de dicto readings as just another perspectival indexation and to model it via a monad just as in examples like (2) or (3)[9].

[7] Note that the distinction w.r.t. entailment of content in DPs crosscuts the factive-non factive distinction, as, e.g., *regret* patterns with *believe* and *fear* patterns with *know*.

[8] Categorically, two objects a and b are isomorphic, if there is an isomorphism $f : a \to b$ with an inverse $f^{-1} : b \to a$ s.t. $f^{-1} \circ f = id_a$ and $f \circ f^{-1} = id_b$. Note that the existence of a specific isomorphism is enough, as it does not have to be unique, although we may sometimes want to establish a canonical isomorphism.

[9] Compatible monads are "stackable", cf. Asudeh & Giorgolo 2020, ch. 7.

Constraining category theory-based meaning can offer flat descriptions of (hyper-)intensional contexts. Understanding semantic computation relationally allows for flexible incorporation of meaning-determining factors in different contexts and of various sizes. As the semantics makes use of an ontologically fixed type system and a set of constraints, it is possibly easier paired with constraint-based frameworks of grammar such as HPSG. It remains open whether other, e.g., transformational frameworks could accommodate a category theoretical semantics with the same ease.
Coming back to the Yoneda Lemma once again, category theory can also offer perspectives on comparing semantic frameworks themselves.

Theorem 1 (Yoneda Lemma[10]) *For all functors $F : \mathscr{C} \rightarrow \mathscr{S}et$ and every object c in $\mathscr{C}$, there is a natural bijection (natural in F and in c) s.t.*

$$Nat(\mathscr{C}(c,-),F) \approxeq Fc$$

Assuming c to be an object in our category of linguistic meaning and F a functor interpreting that object as a corresponding set of possible worlds, the lemma then states that the category theoretical description and the set-based description can be translated into each other. As the category theoretical approach considers the context of an object (i.e., the Yoneda embedding $\mathscr{C}(c,-)$), the set-based approach has no choice but to rely on a point-wise description and thus transfers the complexity into Fc alone, which is reflected in previous work on the topic. In terms of reasoning about computations though, the set-based approach as of now lacks interpretations for 'computational' objects and morphisms (e.g., $t^{\mathtt{e} \times \mathtt{b}}$, or `eval`). If possible world semanticists could answer to that gap, the Yoneda Lemma could be used to prove the approaches to be isomorphic. In conclusion, category theory proves to be a fruitful base for semantic analysis while also providing useful insights for the comparison of (semantic) frameworks.

[10]Cited from Riehl 2016.

References

Asudeh, Ash & Gianluca Giorgolo. 2020. *Enriched Meanings: Natural Language Semantics with Category Theory*. Vol. 13 (Oxford Studies in Semantics and Pragmatics). Oxford: Oxford University Press.

Bradley, Tai-Danae, John Terilla & Yiannis Vlassopoulos. 2022. An Enriched Category Theory of Language: From Syntax to Semantics. *La Matematica* 1(2). 551–580.

Cooper, Robin. 2023. *From Perception to Communication: A Theory of Types for Action and Meaning*. Vol. 16 (Oxford Studies in Semantics and Pragmatics). Oxford: Oxford University Press.

Djärv, Kajsa. 2019. *Factive and Assertive Attitude Reports*. Philadelphia: University of Pennsylvania PhD Thesis.

Djärv, Kajsa. 2023. Knowing and believing things: what DP-complements can tell us about the argument structure and composition of (factive) attitudes. *Journal of Semantics* 40(2). 1–39.

Elliott, Patrick D. 2020. *Elements of Clausal Embedding*. London: University College London PhD Thesis.

Fine, Kit. 2013. A Note on Partial Content. *Analysis* 73(3). 413–419.

Fine, Kit. 2017. Truthmaker Semantics. In Bob Hale, Crispin Wright & Alexander Miller (eds.), *A Companion to the Philosophy of Language*, vol. 2, 556–577. Chichester & Malden: Wiley.

Giorgolo, Gianluca & Ash Asudeh. 2011. Multidimensional Semantics with Unidimensional Glue Logic. In Miriam Butt & Tracy Holloway King (eds.), *Proceedings of LFG '11 Conference*, 236–256. Stanford: CSLI Publications.

Hintikka, Jaakko. 1969. Epistemic Logic and the Methods of Philosophical Analysis. In *Models for Modality*, 3–19. Dordrecht: D. Reidel.

Kratzer, Angelika. 2006. *Decomposing Attitude Verbs: Honoring Anita Mittwoch on her 80th birthday at The Hebrew University of Jerusalem*. `https://semanticsarchive.net/Archive/DcwY2JkM/attitude-verbs2006.pdf`.

Kratzer, Angelika. 2022. Attitude Ascriptions and Speech Reports. In Daniel Altshuler (ed.), *Linguistics Meets Philosophy*, 17–50. Cambridge University Press.

Król, Zbigniew. 2019. Category Theory and Philosophy. In Marek Kuś & Bartłomiej Skowron (eds.), *Category Theory in Physics, Mathematics, and Philosophy*, 21–32. Cham: Springer Nature Switzerland.

Lewis, David. 1979. Attitudes De Dicto and De Se. *The Philosophical Review* 88(4). 513–543.

Mac Lane, Saunders. 1998. *Categories for the Working Mathematician*. New York & Heidelberg: Springer.

Mazur, Barry. 2008. When is One Thing Equal to Some Other Thing? In Bonnie Gold & Roger A. Simons (eds.), *Proof and Other Dilemmas* (MAA Spectrum), 221–241. Washington, DC: Mathematical Association of America.

Moltmann, Friederike. 2017. Partial Content and Expressions of Part and Whole: Discussion of Stephen Yablo: Aboutness. *Philosophical Studies* 174(3). 797–808.

Moltmann, Friederike. 2020a. Clauses as semantic predicates: Difficulties for possible-worlds semantics. In Rajesh Bhatt, Ilaria Frana & Paula Menéndez-Benito (eds.), *Making Worlds Accessible*, 101–117. Amherst: University of Massachusetts Amherst.

Moltmann, Friederike. 2020b. Truthmaker semantics for natural language: Attitude verbs, modals and intensional transitive verbs. *Theoretical Linguistics* 46(3-4). 159–200.

Moulton, Keir. 2009. *Natural Selection and the Syntax of Clausal Complementation*. Amherst: University of Massachusetts PhD Thesis.

Peruzzi, Alberto. 2006. The Meaning of Category Theory for 21st Century Philosophy. *Axiomathes* 16. 425–460.

Riehl, Emily. 2016. *Category Theory in Context*. Mineola, New York: Dover Publications Inc.

Yablo, Stephen. 2014. *Aboutness*. Princeton: Princeton University Press.