Linguistics and Philosophy (2022) 45:1119–1164
https://doi.org/10.1007/s10988-021-09344-9

ORIGINAL RESEARCH

Transparent quantification into hyperpropositional attitudes de dicto

Bjørn Jespersen[1,2] · Marie Duží[2]

Accepted: 11 October 2021 / Published online: 26 March 2022

Abstract We prove how to validly quantify into hyperpropositional contexts de dicto in Transparent Intensional Logic. Hyperpropositions are sentential meanings and attitude complements individuated more finely than up to logical equivalence. A hyperpropositional context de dicto is a context in which only co-hyperintensional propositions can be validly substituted. A de dicto attitude ascription is one that preserves the attributee's perspective when one complement is substituted for another. Being an extensional logic of hyperintensions, Transparent Intensional Logic validates all the rules of extensional logic, including existential quantification. Yet the rules become more exacting when applied to hyperintensional contexts. The rules apply to only some types of entities, because the existence of only some types of entities is entailed by a hyperpropositional attitude de dicto. The insight that the paper offers is how a particular logic of hyperintensions is capable of validating quantifying-in in a principled and rigorous manner. This result advances the community-wide understanding of how to logically manipulate hyperintensions.

Keywords Quantifying-in · Hyperintensional context · Ramified type theory · Transparent Intensional Logic · Extensional logic of hyperintensions

✉ Bjørn Jespersen
b.t.f.jespersen@uu.nl

Marie Duží
marie.duzi@vsb.cz

[1] Department of Philosophy and Religious Studies, Utrecht University, Janskerkhof 13, 3512 BL Utrecht, The Netherlands

[2] Department of Computer Science FEI, VŠB-Technical University of Ostrava, 17. listopadu 15, 708 33 Ostrava, Czech Republic

1 Introduction

A logic of quantifying-in comes at the end of a long story. The ability to validly perform existential quantification into hyperintensional attitude contexts is the sweet fruits a logic reaps for being properly designed and having a number of attractive meta-theoretical features. We will demonstrate how one particular broadly Fregean theory pulls off quantifying-in.

We adopt Cresswell's negative definition of a *hyperintensional context* as a context in which the substitution of necessarily equivalent terms fails. We share the received conception of *hyperpropositions* as sentential meanings and attitude complements individuated more finely than up to analytic equivalence. For instance, an agent can believe that no bachelor is married without logic forcing the agent to believe that the arithmetic of natural numbers is not recursively axiomatizable, despite both complements being necessary truths. Note that in our theory the substitution of terms (linguistic objects) is a reflection of an underlying substitution of one hyperintension for another (logical objects).

To frame the problem, let χ be a hyperintensional operator that represents a binary relation between an agent b and a hyperproposition A to which agent b has adopted an attitude, whether factive or non-factive. The question we raise and answer in this paper can be schematised preliminarily as follows: is the following inferential schema valid ('*QI*' for 'quantifying-in')?

$$\frac{\chi[b, A(a)]}{\exists y \chi[b, A(y)]} QI$$

Our view is this. In general, which positions within A can be quantified into depends on two factors. The first one is what can be quantified over, something which depends on the *ontology* of the logic being deployed. The second one is whether $\exists y$ is capable of reaching across χ so as to bind occurrences of y within the scope of χ. The feasibility of doing so depends on the *syntax* of the logic.[1] For instance, if b believes that 5 is the sum of 2 and 4 (where $a = 5$), then *QI* can be validly applied to infer that there exists some number y such that b believes that y is the sum of 2 and 4, as this conclusion is entailed by the premise. On the other hand, if b believes that $tg(\pi/2) = 0$ (where $a = tg(\pi/2)$), then *QI* cannot be applied to infer that there is a number y such that b believes that y is equal to 0, because there is no such number. Yet *QI* can be validly applied to infer that there is a *function* y (where $a = tg$) such that b believes that y takes the number $\pi/2$ to 0.

Our answer to the question about validity is a *qualified yes*. Answering in the affirmative is predicated on providing satisfactory solutions to both the ontological and the syntactic problem. As for the first problem, it is key that *QI* be qualified so as not to conjure entities into existence. For instance, if the premise is that b believes that Vulcan causes Mercury to have an erratic orbit, then the conclusion must not be that there is a physical object y such that b believes that y causes Mercury's orbit to be erratic. But this restriction still does not preclude quantifying over non-

[1] See Bealer (1982, 26) for discussion of externally quantifiable variables.

extensional entities, i.e., intensions or hyperintensions, something which, however, requires that the framework being deployed must come with a sufficiently rich ontology. For instance, in the above example, the conclusion we would recommend is that there is an *individual concept* such that *b* believes that the physical object falling under this concept causes Mercury to have an erratic orbit.[2]

As for the second problem, our answer to the question of whether $\exists y$ is capable of reaching across χ and binding occurrences of y within the scope of χ is also a *qualified yes*. The qualifications are due to the fact that a hyperintensional context, and anything located within it, is not immediately amenable to logical manipulation. Intensional contexts, by contrast, do immediately lend themselves to logical manipulation. The operation that 'raises' a context to a hyperintensional one shields it from manipulation from the outside. Therefore, further operations are required in order to, nonetheless, reach into positions located within a hyperintensional context.

The bulk of this paper is devoted to demonstrating, in full detail, how a particular theory is capable of quantifying into hyperpropositional attitude contexts de dicto. The theory in question is Transparent Intensional Logic (TIL), which has two features that enable it to validate *QI*:

- *A fully transparent semantics.* Opacity and shift of semantic value are eschewed in favour of terms and expressions having the same meaning and the same semantic value in all contexts. The same uniform semantics applies to all the different kinds of context.[3]
- *An extensional logic of hyperintensions.* The laws of extensional logic, including existential generalisation, also apply to hyperintensional contexts. However, stronger requirements are placed on the operands than in the case of an extensional logic of intensions and of extensions. What varies with the context is not the validity of the rules themselves, but the types of objects these rules are applicable to.

Transparency and extensionality are two necessary conditions for our strictly compositional theory of quantifying-in. More specifically, what connects the transparent semantics with the extensional logic of hyperintensions is that the combination contributes to making it technically feasible to, first, denote, or present, a hyperintension as an entity in its own right and, next, to manipulate either the entire hyperintension or parts of it. The way we achieve this is by making it feasible for a hyperintension to occur as a functional argument. When a hyperintension occurs as an argument, we shall say that it occurs in the *displayed* mode, i.e., the hyperintension occurs *presented as an argument*. Occurring in the displayed mode contrasts with occurring in the *executed* mode, which involves descending to a lower-type entity. Once a hyperintension is able to occur as an argument, it can

[2] Here we use the term 'individual concept' in an intuitive sense. Below we are going to distinguish and rigorously define *individual role* or *office*, i.e., *individual-in-intension*, in opposition to *individual hyperoffice*, which is a hyperintension presenting an office.

[3] A fully transparent semantics qualifies as 'semantically innocent' according to the letter (if not spirit) of Davidson's characterisation, but we arrive at semantic innocence via the opposite route than Davidson's. He attempts to make each context extensional; we generalise from hyperintensional contexts to all other contexts. See Duží et al. (2010, 12).

figure as the complement of an attitude (thus making hyperintensional attitude contexts possible), it can be quantified *over*, it can be quantified *into*, and a part within it can be replaced by another so as to generate a new hyperintensional complex.

Quantifying-in is, above all, *technically* demanding and nowhere close to being trivial. Attempting to quantify into hyperintensions puts pressure on any theory's account of nested contexts, its variables and quantifiers. Therefore, the demonstrated ability of a logic of hyperintensions that comes with an elaborate technical machinery to successfully perform quantifying-in is evidence that the logic is justified in adopting this machinery. Furthermore, we have a theory-internal reason to demonstrate exactly how to quantify-in. TIL makes the strong claim for itself that it is an extensional logic of hyperintensions, and this is one reason why we must be able to preserve the validity of the extensional rule of existential generalisation also when applied to hyperintensional contexts.[4]

On top of that, it is also *philosophically* challenging, hence enlightening, to figure out the nature of the so-called intentional objects that are at the receiving end of an intentional act such as maintaining an attitude. For instance, if the premise is that *b* believes that the last decimal of the expansion of π is an even number, then what would be a suitable quantificational range? The premise may well be true but, necessarily, the complement of the attitude fails to be true for want of a last decimal of the expansion. Or if the premise is that *b* believes that the King of Denmark is balding, then what would be a suitable quantificational range? As a matter of contingent fact, the premise may well be true, but the complement is not, for it so happens that nobody is presently the King of Denmark (writing in 2021). Since we cannot quantify over individuals here, then what can we quantify over? Answering this question reveals which kinds of objects one is prepared to embrace in one's ontology. We will show how the ontology of TIL makes it possible to infer the existence of such entities as are logically presupposed, hence also entailed, by the premises. For instance, if *b* believes that the King of Denmark is balding then there is an individual concept such that *b* believes that its occupant is balding. Or, if *b* believes that the last decimal of the expansion of π is an even number then there is the concept of a number such that *b* believes that the number falling under this concept is even.

[4] Morton (1969, 163) says, "treatments of non-truth-functional contexts have assimilated them to intensional contexts, either to shade them with the same dark incorrigibility [i.e., the fact that it is obscure how to calculate the truth-value of an intensional context, thus understood. *The authors*] or to honor them with all the mathematical and philosophical sophistication that the intensional requires." Our conception of intensionality (actually, hyperintensionality) is the latter, which Bealer sums up thus: "[T]here is no genuinely intensional language; when *prima facie* intensional language is properly analysed, it turns out to be extensional language concerning intensional entities." (Bealer 1982, 148) See also Copi (1968, 244) and Klement (2002, 99–100). When these authors speak of 'intensionality' they intend intensionality as understood in mathematics, which is hyperintensionality. The coarse-grained intensionality of possible-world semantics equates co-intensionality with necessary co-extensionality, thus yielding (in a logic of total functions) but one necessary proposition, but one impossible proposition, failure to distinguish between inverse relations, etc., etc.

This paper is continuous with previously published research, starting with Tichý (1986), Materna (1997) and followed by Duží and Jespersen (2015, 2012) and Jespersen (2015a, 2015b), which cover quantifying into hyperpropositional attitudes de re (e.g., *believing of* π that *it* has an infinite decimal expansion, or π being such that *it is believed* to have an infinite decimal expansion) and hyperintensional objectual (i.e., non-propositional) attitudes (e.g., *calculating* the ninety-ninth decimal in the expansion of π, or *seeking* a yeti without seeking an abominable snowman). In this paper we concentrate on hyperpropositional attitudes de dicto.[5] To fix ideas, here are some examples of hyperpropositional attitudes de dicto:

- Tilman knows that 1+1=2, but he does not know that arithmetic is not recursively axiomatizable.[6]
- Tilman is trying to prove that the last decimal of the expansion of π is an even number.
- Tilman knows that Francis is the Pope, but not that he is the Head of the Catholic Church.[7]
- Tilman believes that whatever does not kill him makes him stronger, but not that whatever does not make him stronger kills him.[8]
- Tilman believes that no bachelor is married, but he does not believe that whales are mammals.[9]

Each of these attitudes relates an attributee to a hyperproposition, which offers an idiosyncratic perspective on an empirical or analytical state-of-affairs. An attitude de dicto reproduces exclusively the attributee's own perspective, whereas an attitude de re blends the attributer's and the attributee's respective perspectives. This explains why the content of an attitude de dicto is fully specified and the content of an attitude de re is only partially specified. Linguistically, the difference

[5] Yalcin (2015, 207) asks, "what should the semantic analysis of attitudes de re look like from a Fregean perspective—a perspective according to which attitude states are generally relations to structured Fregean thoughts, themselves composed of senses?". Yalcin (2015, 208) claims that "the Fregean position is underdeveloped" and left with a 'lacuna', because no Fregean position has so far specified how to compositionally derive truth-conditions for attitudes de re. We beg to disagree. Both Duží et al. (2010, §5.1.2.2) and Duži and Jespersen (2012) answer Yalcin's question and address his complaint. The Quinian problem of 'double vision' (i.e., the Ralph/Ortcutt case; see Sect. 2.2 below) which Yalcin brings up in (2015, § 4) is solved in Jespersen (2015a, 2015b).

[6] Since all true mathematical sentences denote the truth-value **T,** on an intensional reading the sentence would be a contradiction. On an intensional reading, any true mathematical sentence can be substituted for the complement, and we end up with the paradox of mathematical omniscience.

[7] Since, by definition, the Pope and the Head of the Catholic Church are one and the same office, the sentence would be contradictionary on an intensional reading. Yet, since the sentence can be true, the attitude must be hyperintensional.

[8] Again, on an intensional reading, the sentence would be contradictory; hence, the attitude must be a hyperintensional one.

[9] The attitude must be hyperintensional, because on an intensional reading the sentence is contradictory. If Tilman believes that no bachelor is married then on an intensional reading he must believe any necessarily true proposition, like, e.g., that whales are mammals, and we end up with the paradox of analytical omniscience. The other undesirable extreme is the paradox of analytical idiocy, so to speak. If one believes a necessarily false proposition (e.g., that a forged banknote is a valid banknote) then one would have to believe any necessarily false proposition.

is that a report of an attitude de re includes an anaphor that points outside the embedded context, as in "that *it* is a planet". A report de dicto would have "that *Pluto* is a planet". Moreover, a hyperpropositional attitude de dicto is impervious to the (contingent or necessary) inexistence of, e.g., Vulcan or the last digit of the expansion of π, in that its truth-value depends merely on whether or not it is true that the attributee believes that such-and-such is true. By contrast, attitudes de re come with existential presuppositions already at the extensional level of individuals, numbers, etc.: no *res*, no attitude de re.

The rest of the paper is organised as follows. Part 2 provides systematic background to the problem of quantifying-in, comparing TIL with other positions. Part 3 presents the relevant fragments of TIL. Part 4 presents and proves our rules for quantifying into hyperpropositional contexts de dicto.

2 Background to transparency and quantifying-in

2.1 Quine

Quantifying-in mixes modality with quantification.[10] Quine was squarely opposed to existential quantification into alethic modal contexts, such as $\exists xFx$. His general objection was that it generates 'Aristotelian essentialism', i.e., essentialism de re, which he deemed incoherent.[11] Quine would later, in (1956), adopt a more nuanced stance on existential quantification into a different sort of modal contexts, namely those attributing attitudinal modalities, such as wishing to find a unicorn lair, or believing that the tall handsome stranger spotted on the beach is a spy. In Quine (1956), he distinguishes between *notional* and *relational* attitudes. The sentence "Ralph believes that someone is a spy" lends itself to both readings. The notional reading is that Ralph believes that there are spies. The relational reading is that there is someone of whom Ralph believes that he or she is a spy. Only in the relational case does it make sense to quiz Ralph about who it is he suspects of being a spy.[12] Or for a standard example: "The princess wants to marry a prince". The formalisation of the relational reading in first-order logic (which at least gets the scope distribution right) would be this (χ, a generic attitude operator; M, the binary relation of marrying):

[10] So far, so good. But beyond that, exactly what problem, or cluster of problems, is being discussed in Quine (1956), or his previous work on quantifying into modal contexts, is still not entirely clear. See, for instance, Crawford (2008). Bear in mind that we are not engaged in Quine scholarship as such, but rather in charting the systematic roots of the problem of quantifying-in in the light of how we find it most fruitful to frame it.

[11] One argument against quantified modal logic is his example of the 'mathematical cyclist', which is intended to show that it is both necessary and also not necessary that an individual who is a biking mathematician is rational and bipedal. See Duží et al. (2010, §4.2.1) on how to debunk this argument along the same lines as in Stalnaker and Barcan Marcus.

[12] See Kaplan (1986, App. B) on 'the syntactically de re', which is supposed to capture Quine's relational readings. The technique consists in forming a predicate in the passive voice in the vein of 'is believed by *a* to be an *F*'. This yields "The tall handsome stranger spotted on the beach is believed by Ralph to be a spy".

$$\exists x\,(Fx \wedge \chi a\,(Max))$$

The notional reading goes into:

$$\chi a\,(\exists x\,(Fx \wedge Max))$$

Quine dismisses quantifying into notional attitudes. Doing so would validate inferring *from a* believing that there is a planet orbiting between Earth and the Sun that causes Mercury's orbit to be erratic *to* there being a planet between Earth and the Sun such that *a* believes that it causes Mercury's orbit to be erratic. Thus, it is made a necessary condition for Le Verrier's hypothesis about Mercury's orbit that Vulcan exists, which goes far too far by turning believing into a factive attitude. The quantification would also misconstrue what notional attitudes are all about. When Quine wants, on a notional interpretation, a sloop then it should exactly not follow that there is a sloop such that Quine wants it. But, as Kaplan (1986, 230) is right to stress, Quine does want to quantify into relational attitudes. Quine's problem then becomes how exactly to go about that. Quine himself offers his three-place analysis as an attempt to formulate what is, by his lights, a non-opacity-inducing formalisation that makes co-referential terms substitutable. Kaplan (1968) also puts forward some inconclusive proposals, which, however, tie relational attitudes tightly together with the particular terms chosen.

Still, the general problem of variable-binding remains. There is an incongruity between the notional and the relational reading. On both readings, the ∃-bound occurrence of x is located within the scope of χ. But on the notional reading, also ∃ is within the scope of χ. For sure, the entire context induced by χ may be 'opaque' or 'intensional', but the two occurrences of x in $\exists x$ (… x …) are on the same level. Not so on the relational reading. While $\exists x$ is in a transparent or extensional position, the occurrence of x within the scope of χ is in an opaque or intensional position. And whereas the quantificational range of x is restricted, due to Quine's extensionalist predilections, to extensional entities, the occurrence of x in an opaque context demands a shift in quantificational range, most likely so as to include individuals-in-intension as well as other extensions-in-intension. However, apart from such entities being beyond the pale for Quine, the formal predicament becomes obvious if we rename the second variable:

$$\exists x\,(\ldots\chi a\ldots y\ldots)$$

The above formula is *open*, because y occurs free, and there is no semblance of contact between $\exists x$ and y. On the other hand, there is a semblance of contact between $\exists x$ and x here:

$$\exists x\,(\ldots\chi a\ldots x\ldots)$$

However, the contact between $\exists x$ and x has been severed, appearances notwithstanding. Thus, the formula exemplifies vacuous quantification, again making '$\exists x$' a dummy.

2.2 Substitutability

In the rest of this section, we will describe and critique some contemporary positions regarding transparency, opacity and substitution, all of which affect the prospects of quantifying-in. The form in which the problem of quantifying-in has been handed down is that it challenges us to make relational attitudes transparent—and the hallmark of transparency is the validity of quantifying-in.[13] Transparency would seem to validate an inference such as this:

$$\frac{\begin{array}{l} b \text{ is such that } a \text{ believes that she is an } F \\ b = c \end{array}}{c \text{ is such that } a \text{ believes that she is an } F}$$

This seems like a straightforward application of Leibniz's Law, and so ought to be uncontroversial, also because it aligns with the tenet of 'no opacity de re'. Yet Pickel (2015), Cumming (2008) and Caie et al. (2019) think otherwise. Pickel (2015, 345) says,

> What is wanted is a theory that is sensitive to both the state of the world that the agent believes to obtain […] and the peculiar take she [i.e., the agent] has concerning 'who is who' in this state of the world [so as to make room] for the possibility that sentences differing only by the substitution of co-referential names […] differ in truth-value.

The point about 'who is who' in a given context is to be captured by

> […] the fact that variable x_l associated with the name 'Lindsay' may designate different individuals relative to different assignment functions as representing the fact that Lindsay *may* be each of these individuals, where this '*may*' reflects epistemic possibility. (Pickel, 2015, 339.)

On our interpretation, the effect of adding a so-called assignment-unsaturated meaning for the sentence "… Lindsay …" is similar to the effect sententialists and inscriptionalists obtain when they make the very syntax in which an attitude is

[13] Quine states that "*no variable inside an opaque construction is bound by an operator outside*. You cannot quantify into an opaque construction." (Quine, 1960, 166) It is the right move, of course, for Quine to resist quantifying into opaque contexts. In Quine (1960) and elsewhere, he likens trying to quantify into opaque contexts to trying to quantify into quotation contexts. In Quine (1956) he gives an additional reason. In the famous Ralph/Ortcutt case, it is true that a is believed by Ralph to be a spy, that b is not believed by Ralph to be a spy, and that $a = b$. What happens is that quantification 'quantifies away' the two different guises under which Ralph has encounted a/b. There can be no individual such that it is believed, and also not believed, by someone to be a spy. So, quantifying-in would yield a paradox. See also Kaplan (1986, 269–70). But, or so we think, the fact that opacity appears to be the root cause should have given Quine pause. He ought to have reconsidered the assumptions and tenets that landed him in an (ostensibly) opaque context that suspends quantifying-in on pain of paradox. The conclusion should not have been that opacity is a fact of linguistic life, or 'intensional', i.e., anti-extensional, logic a fact of logical life, thus turning some contexts into no-go areas. In fact, the strategy pursued by TIL is to design a formal semantics that cannot generate opacity, again with provisos for quotational contexts. (We note that the line of reasoning found in the Ralph/Ortcutt example resembles that of the reasoning behind the 'mathematical cyclist'; see fn. 12; see fn. 11.)

reported part of the reported attitude. Of course, substitution of co-referential, or even synonymous, terms and expressions will not go through; but that is because the substitution context is a quotational one. The respective attitude contexts of Pickel and Cumming are not quotational. Whereas the sententialist/inscriptionalist makes it matter that their agents may know one name for an individual, but not another, Pickel and Cumming make it matter that an agent may fail to identify a as b, even though a is identical to b. This is captured by including assignments relative to which x_a, x_b take different values, although $a = b$.[14] TIL, however, does not want to model "that [the agent] is unsure of whether *Lindsay* and *Nellie* are the same person" (Pickel, 2015, 339), for there is no such thing to model. Not to put too fine a point on it, Lindsay and Nellie not being the same person is not a matter of epistemic (or doxastic) possibility, but a case of the agent being conceptually confused. TIL assumes that its agents are aware of the identity of the individuals toward which they adopt an attitude. Of course, it makes perfect sense for us that agents may be unsure of whether the sun that sets in the evening is the same sun that rises in the morning. But this should be analysed as being unsure about whether two different individual offices are co-occupied.[15] When a pair of semantically proper names are synonyms (and assume 'Lindsay' and 'Nellie' to be such a pair), no Fregean puzzle arises, and 'Millian' substitution is valid. There is no semantic or logical difference between the two names, which are merely notational variants of one another. The initial inference is valid, by the lights of TIL, and the single problem is to safeguard the anaphoric reference from 'she' to 'b', 'c'.[16] A transparent semantics is characterised by being able to do so. The substitution is subsequently validated by Leibniz's Law.

Still, we are not entirely on board with framing the problem of quantifying-in in terms of making contexts reporting relational attitudes transparent. We do agree that Quine's distinction between notional and relational readings is intuitively persuasive. In fact, anyone who has absorbed the implications of Russell's example of "I thought your yacht was longer that it is" will probably be fine with Quine's distinction. We also agree that attitude reports de re are logically distinct from

[14] The main difference between Pickel and Cumming is that Pickel assigns a more elaborate semantics to his variables. Cumming has, as it were, got only the first half right. Pickel provides an argument to the effect that Cumming is unable to distinguish between true and false beliefs. Assume that a believes, (B_a), that the value of x_b is an F. This is formalised thus: "B_aFx_b". This is a closed formula, because operator B binds the variable. Assume that $\sigma(x_b)$ = Dublin. The formula being closed, it retains its truth-value independently of any assignment functions other than the original σ. Now let an arbitrary assignment function, τ, assign a different value: $\tau(x_b)$ = Lublin. It is true, therefore, that a believes that Lublin is an F. Except, of course, it is not. Pickel's remedy is to assign a *dual* semantics to variables. Whether B_aFx_b "is true on assignment σ depends not just on the value of x_b relative to σ, but also on whether every world-assignment pair [$\langle w, \tau \rangle$], in the agent's belief set makes true [the 'quasi-open proposition' Fx_b]. The assignments in [believer a's] belief set may assign different values to x and y, even though x and y co-refer on the input assignment [σ]." (Pickel, 2015, 347). TIL goes in the opposite direction. We do not want the option to change horses in midstream, so to speak, by bringing in an alternative to the 'input assignment' in a static context. TIL does not capture an agent's idiosyncratic perspective by means of 'shiftable' assignment functions, but by means of fine-grained, structure-sensitive hyperpropositions as attitude complements.

[15] See Duží et al. (2010, §3.3.1).

[16] See Duží et al. (2010, §3.5) on anaphoric reference.

attitude reports de dicto.[17] But we do not agree with the residual distinction between *opacity* and *transparency*. First and foremost, all of our contexts are transparent. Put bluntly, if a theory ends up with a category of what it calls opaque contexts then there is something wrong with the theory. The context-*in*variant semantics of TIL is obtained by universalising Frege's denotation-shifting semantics custom-made for 'indirect' contexts. Whereas Frege's semantics for attitude contexts was located on the margins of his overall semantic theory, we locate it right in the centre of ours. The upshot is that it becomes trivially true that all contexts are transparent. All singular-term positions are 'purely referential' (to use Quine's phrase), in the sense that pairs of terms that are co-denoting outside an attitude context remain co-denoting inside an attitude context, and pairs of terms that are not co-denoting inside an attitude context do not become co-denoting outside an attitude context. Thus, although Quine's 'the man in the brown hat' and 'the man on the beach' contingently share the same extension (Bernard J. Ortcutt, as it happens) they never co-denote him. Rather they denote, in every context/independently of context, two distinct individual offices. One comes with the uniqueness condition that its occupant must be the man in the brown hat (relative to some unspecified empirical context), and the other comes with the uniqueness condition that its occupant must be the man on the beach (again relative to some unspecified empirical context).

Second, we do not agree with the (by now obsolete?) dismissive understanding of 'intensional' as 'failing to validate rules of extensional logic and invoking creatures of darkness'. 'Intensional', as we use the term, means only 'involves intensional entities identified with functions from possible worlds'.

Third, and relatedly, when discussing quantifying into non-extensional contexts, we distinguish between *intensional* and *hyperintensional* contexts. Quantifying into intensional contexts is smooth sailing.[18] Quantifying into hyperintensional contexts is technically complicated and ontologically more exacting. The problem is less to do with the fine-graining of such contexts and more to do with having to operate on logical structures, or parts of structures, as opposed to merely operating on functions or their arguments. Quine's original problem with reaching an x inside an attitude context is, thus, also ours. But labelling the problematic context as 'opaque' explains nothing and just relabels the problem. The actual problem is that this x occurs in a different fashion inside a hyperintensional context than in either an intensional or extensional context.[19]

[17] Whether *notional/relational* must map onto *de dicto/de re* is far from a foregone conclusion, though, as different theories will have different conceptions of the dicto/re distinction. For instance, should some form or other of *acquaintance* play a role in attitudes de re? [For the record: no, not in TIL. See Duží et al. (2010, 435)].

[18] See Duží and Jespersen (2015, 2012) and Duží et al. (2010, 497–99).

[19] Quine's original objection to quantified modal logic is that (what appears to be) the same variable will have both used and mentioned occurrences within the same context. See also Kaplan (1986, 262–63). On a similar note, Pickel (2015, 340) objects to Cumming (2008), "There is no coordination between the occurrences of x outside of the belief ascription and the x occurring within the belief ascription". Our distinction between *displayed* and *executed* modes of occurrence of procedures, including variables, is sort of parallel to the distinction between words occurring mentioned or used, and quantifying into a displayed procedure is sort of parallel to quantifying into a quotation context. But we do not wish to push the parallel too far. Attempting to quantify into a quotation context is a no-starter, whereas the main

Fourth, the ability or inability to quantify-in is not what sets attitudes de re apart from attitudes de dicto. In TIL, both kinds are equally susceptible to quantifying-in. Nor should their difference be captured by means of scope differences between ∃ and χ. Quantified formulas (or rather their semantic counterparts) are a logical *consequence* of both kinds of attitudes, and not definitional of either of them. Rather their difference is anchored in a difference in logical structure (see Sect. 4.1). Their logical structure reveals that attitudes de re come with an existential presupposition on the level of extensional *res*, and if the presupposition is not satisfied then the quantified sentence is neither true nor false. This is as it should be, for attitudes de re are object-dependent. For instance, in the absence of Le Verrier's intermercurial planet, the appropriate *res* is not around to instantiate properties and fail to instantiate other properties, so predications de re about Vulcan cannot acquire a truth-value. By contrast, attitudes de dicto allow flights of fancy, so to speak, because they are not restrained by existential presuppositions. Still, attitudes are inten*t*ional relations that are invariably *about* something, so also attitudes de dicto qualify as object-dependent, provided objects other than extensional ones are allowed into one's ontology. What quantifying-in brings out is exactly what type of object a given attitude de dicto is dependent upon in the sense of having it as its complement. Logic and semantics intersect with metaphysics here, because the validity of existential quantification into hyperintensional contexts presupposes both suitable quantificational ranges, a transparent semantics and an extensional logic of hyperintensions.

2.3 Transparency versus opacity

To locate TIL in the wider landscape, we are aware of four diverse avenues one might pursue when attempting to validate quantifying-in. One is contextualism as made presentable by Frege and later formally encoded by luminaries such as Church and Montague. Our problem with this is that it allows some contexts to be opaque.[20] Another invokes 'flat' (hyper-) intensions as urged by, e.g., Bealer and Turner.[21] Our problem with this is that it foregoes any notion of objectual (hence, extra-syntactic) logical structures within which to operate on constituents. We are left with manipulating symbols, which sheds no light on hyperpropositions themselves. Yet another approach turns to various variants of sententialism, which relates agents directly to inscribed or uttered tokens (of types) of sentences.[22] Our problem with this is not only its excessive fine-graining and the fact that attitude ascriptions are held hostage to a particular symbolism or spoken language, but also the absurdity of quantifying into quotation contexts. The final one would be top-down, highly expressive, context-invariant theories, which enable objectual quantification into

Footnote 19 continued
technical point we are making here is that it is both feasible and sensible to quantify into a displayed context.

[20] For further critical comments on contextualism, see Duží et al. (2010, 110–112).

[21] See Turner (1992, 165) for a 'flat version of Montague's intensional logic' developed within the untyped λ-calculus.

[22] See Berto and Nolan (2021, §2.2) for examples and discussion.

hyperintensional contexts. TIL is, to the best of our knowledge, the only theory of this very kind. It is a defining feature of this sort of position that hyperintensional contexts are continuous with the semantics for intensional and extensional contexts. This marks a departure from Frege's contextualist semantics, of course, as does the introduction of a typed universe. But we still wish to characterise TIL as a broadly Fregean semantics. One reason is that we draw liberally on senses. Another is that the notion of *function* has been chosen as a theoretical primitive in TIL. This second point explains why our general logical framework is provided by the (typed) λ-calculus and its two key operations of application and abstraction, together with its rules of conversion. When a function is a mapping, *sets* and *relations* are rendered particular species of functions, namely such as are identified with their respective characteristic functions.[23]

Interestingly, there is another theory which also employs the typed λ-calculus (except in the vein of Montague Grammar), though with a view to developing a logic of opacity. Caie et al. (2019) sets out to develop such a logic compatible with higher-order classical logic.[24] Classical opacitists (as they call themselves) and TIL ('classical transparentists', presumably) agree on the semantics for the connectives and the identity predicate, but part company over the unconditional validity of Leibniz's Law. This 'law' being unconditionally valid is a necessary condition for Substitution being unconditionally valid. This is the schema of Substitution:

Substitution. $$a = b \rightarrow (\varphi \leftrightarrow \varphi[b/a])$$

Opacity is defined as a false instance of Substitution:

Opacity. $$a = b \wedge \neg(\varphi \leftrightarrow \varphi[b/a])$$

The metaphysics that goes together with opacity is this:

> We take it to be obvious that Hesperus and Phosphorus are identical, though we are exploring views on which Hesperus and Phosphorus have different properties. (Caie et al., 2019, 12, fn. 31.)

[23] 'Function' is historically ambiguous between Frege's *Funktion* (generation of mapping) and *Wertverlauf* (mapping). Church (1956, 16) is clear on this: "If the way in which a function-in-extension yields or produces its values for its arguments is altered without causing any change either in the range of the function or in the value of the function for any argument, then the function remains the same; but the associated *function concept*, or concept determining the function …, is thereby changed." In modern-day parlance, function concepts qualify as hyperintensions. Our notion of hyperintensions is rooted in Church's *function-in-intension* or *functional concept*, Frege's *Funktion* and *Sinn*, as well as *Turing machine*. *Application* and *abstraction* are theoretical primitives, which are central to the definition of two of our hyperintensions. See Definition 1 (iii), (iv), in Sect. 3.4.

[24] Caie et al. (2019) is an exceptionally rich paper, which we would have liked to engage with at length. For now, we are just scratching the surface and confining ourselves to the core question of the validity of Leibniz's Law, hence of Substitution. Thus, this comparison of 'classical opacity/transparency' is about substitution specifically rather than quantification. We want to stress that we find it commendable that someone should try to develop a *logic* of opacity, which will distinguish between true and false instances of Substitution. Opacity has typically been understood purely negatively as the failure to preserve transparency, but Caie et al. (2019) helps clarify what the logical implications are of adopting opacity. Still, we disagree with treating opacity as a datum (even in the explorative spirit of Caie et al. (2019)), rather than as a symptom of a wrongheaded formal semantics.

This is anathema to TIL. If the second clause is true then we reject the first clause (or *vice versa*). We are quite happy to do so, in fact, because many standard cases of "… is …", such as "Hesperus is Phosphorus" and "Water is H_2O", should not go into "… = …", as in "Hesperus = Phosphorus" or "Water = H_2O".[25] The welcome upshot is that Leibniz's Law is rendered inapplicable instead of invalid. If one, nonetheless, pushes ahead with "*a* = *b*" followed by substitution, one ends up with an inference that is valid, for sure, but also unsound. This schema summarises the position TIL assumes as regards substitutability within hyperpropositional attitude contexts, whether de re or de dicto:

$$Transparency. \qquad a = b \rightarrow \chi\varphi a = \chi\varphi b$$

Substitutability within the scope of χ is a necessary condition for the identity, or at least hyperintensional isomorphism, of *a* and *b*.[26] This schema is perfectly trivial in TIL, which is because the threshold for being a correct instance of "*a* = *b*" is high. The constraint is that the substituends for '*a*' and '*b*' must be pairs of synonymous terms. TIL has a catalogue of identity, congruence and equivalence relations, and depending on whether the context within which one intends to perform substitution is extensional, intensional or hyperintensional, one or the other relation is required for valid substitution (see, e.g. Duží et al., 2010, §2.7.1). Obviously, self-identity guarantees substitutability even in hyperintensional contexts. As soon as transparency is adopted, one steers clear of a tangle such as the following, which opacity seems to be tasked with disentangling:

> […] although for Hesperus to be visible at night just is for Phosphorus to be visible at night, the ancients knew that Hesperus is visible at night, but did not know that Phosphorus is visible at night. (Caie et al., 2019, 15)[27]

Formally:

$$\varphi a = \varphi b \wedge (\chi\varphi a \wedge \neg\chi\varphi b)$$

We agree about what the ancients knew and did not know; we disagree that Hesperus being visible at night is the same as Phosphorus being visible at night. In TIL, Hesperus is the office of being the brightest non-lunar object in the evening/night sky, and Phosphorus is the office of being the brightest non-lunar object in the morning sky. These two offices are contingently vacant or occupied, and when occupied, then contingently co-occupied. Hence, "(the occupant of the office of) Hesperus is visible at night" and "(the occupant of the office of) Phosphorus is visible at night" denote two *different* possible-world propositions that just happen to be both true (perhaps since the origin of our Solar system). If "$\varphi a = \varphi b$" is read de dicto or intensionally (meaning that φa, φb are one and the same possible-world

[25] "Water = H_2O" has never sat well with us. How can a liquid be identified with a molecular structure? This smacks of category mistake, or type-theoretic incongruity. "Water = H_2O" feels like a throwback to the long-gone days of materialist reductionism. We would rather say that (pure) water *has* H_2O (namely, as its molecular structure).

[26] See Sect. 4.1.1 on procedural isomorphism, which defines co-hyperintensionality.

[27] See also Duži et al. (2010, 3): *Propositional Hesperus/Phosphorus.*

proposition), then it is false. If "$\varphi a = \varphi b$" is read de re or extensionally (meaning that φa, φb are both true), then the second conjunct, "$(\chi\varphi a \wedge \neg\chi\varphi b)$", is improper, hence without truth-value, because χ is not an attitude to a truth-value. The de dicto reading is more consonant with the quote above. The opacitist needs their semantics to accommodate a case of this form: "Something is the same, but the ancients did not know it was the same". The transparentist needs their semantics to accommodate a case of this form: "Something and something are not the same, and the ancients did not know that they were the same."

Opacity helps the opacitist to a true conjunction and the preservation of the Frege puzzles of cognitive significance. The opacitist has recourse to opacity to save the conjunction above from coming out false: it cannot be true that the ancients knew, and also did not know, that the same celestial body was visible at night. The complications that opacity incurs—developing two logics, one for transparency and one for opacity, and maintaining a system of double bookkeeping, one book for transparent contexts and another book for opaque contexts—serve the purpose of maintaining a fairly simple semantics for "*a* is *b*", "φa" and "φb".

Transparentists will have to discard the first conjunct in order to preserve the Frege puzzles. To see this, if the conjunction were instead "$a = b \wedge (\chi\varphi a \wedge \neg\chi\varphi b)$" then the conjunction would be necessarily false. If the conjunction were "$a = b \wedge (\chi\varphi a \wedge \chi\varphi b)$" then the conjunction would be necessarily true, with the conjuncts of "$\chi\varphi a \wedge \chi\varphi b$" being mere notational variants. In order to preserve both transparency and the non-triviality of both conjuncts, the transparentist develops a more elaborate semantics for "$a = b$", "φa" and "$\varphi a = \varphi b$" that remains the same whether occurring within the scope of χ or not. Let the first conjunct be "$a = b$". In prose, the result becomes: the office of Hesperus and the office of Phosphorus happen to share the same occupant, and the ancients knew that the occupant of the office of Hesperus was visible at night, but the ancients did not know that the occupant of the office of Phosphorus was visible at night. This is a true conjunction. Let the first conjunct now be as above: "$\varphi a = \varphi b$". In prose, the result becomes: the occupant of the office of Hesperus being visible at night is identical to the occupant of the office of Phosphorus being visible at night, and the ancients knew that the occupant of the office of Hesperus was visible at night, but the ancients did not know that the occupant of the office of Phosphorus was visible at night. As we have argued, "(the occupant of the office of) Phosphorus is visible at night" is a different proposition than "(the occupant of the office of) Hesperus is visible at night"; therefore, "$\varphi a = \varphi b \wedge (\chi\varphi a \wedge \neg\chi\varphi b)$" comes out false. This outcome is in keeping with Transparency; if one maintains that "$\neg(\chi\varphi a = \chi\varphi b)$" then it is no option to maintain, nonetheless, that "$\varphi a = \varphi b$".

For a general characterisation of classical transparentism, we would not hesitate to characterise TIL as a *transparent higher-order logic*, although TIL would not entirely fit the opacitists' description of such a logic in Caie et al., (2019, 8). We are only hesitant about one element in their description, though. This is their list of what characterises a transparent higher-order logic (ignoring the logical relationships between the various principles, as some entail others), together with further principles from their catalogue that also suit TIL (we follow their formalisation):

- *Equivalence.* $a = a \wedge (a = b \wedge a = c \rightarrow b = c)$
- *Material Equivalence.* $p = q \rightarrow p \leftrightarrow q$
- *Beta-Eta Equivalence.* $\varphi \leftrightarrow \psi$, provided φ and ψ are $\beta\eta$-equivalent
- *Lift Congruence.* $a = b \rightarrow (\lambda X.Xa) = (\lambda X.Xb)$
- *Application Congruence.* $F = G \rightarrow Fa = Ga$
- *Substitution.* $a = b \rightarrow (\varphi \leftrightarrow \varphi\ [b/a]$
- *Universal Instantiation.* $\forall x\varphi \rightarrow \varphi[a/x]$, where a is free for x in φ
- *Leibniz's Law.* $a = b \rightarrow \forall X\ (Xa \leftrightarrow Xb)$

Our proviso remains intact: (the substituends of) '*a*', '*b*', '*c*' must be constants denoting individuals and not offices or anything else, unless we wish to express the self-identity of an office (etc.) bearing more than one name, which is the exception rather than the rule. Our only reservation is with the constraint *Beta-Eta Equivalence*, in case it is recruited for the purposes of hyperintensional individuation or for a logic of *partial* functions. If $\beta\eta$-equivalence is instead only applied to conversion of terms denoting *total* functions occurring in non-hyperintensional contexts, then *Beta-Eta Equivalence* applies to TIL.[28]

3 Transparent Intensional Logic

This part describes and defines the foundations of TIL, together with the particular devices required to operate on hyperintensions.

3.1 Function and procedure

The most fundamental distinction in TIL is between *procedures* and *functions*. Procedures are structured, higher-order entities, and on a given occasion a procedure occurs either in the *displayed* or *executed* mode within another procedure. The default is that procedures occur in executed mode. Functions are modern-day mappings from a domain to a range, $f: x \mapsto f(x)$, with the important proviso that TIL allows also functions that are only *partially* defined. Functions are set-theoretic (hence, unstructured) and first-order entities, unless the elements of a domain or range of a given function are higher-order objects, in which case the function also becomes a higher-order object in the type hierarchy. Sets and relations are defined as particular kinds of functions; a set is identified with its characteristic function, and an n-ary relation is identified with a function from n number of arguments to a truth-value. Intensional entities (as per possible-world semantics) are identified with functions from a logical space of possible worlds, and necessary co-extensionality equals co-intensionality. Extensional entities such as individuals and truth-values

[28] See Jespersen (2021) regarding β-conversion and η-conversion with regard to hyperintensional individuation. See Duží and Kosterec (2017) or Duží and Jespersen (2013) for discussion of β-conversion and η-conversion with regard to a hyperintensional logic of partial functions.

are typed as medadic, or nullary, functions, i.e., as constant values. One important thing to bear in mind is that *execution* applies to procedures while *evaluation* applies to functions.

The *syntax* of TIL is that of a typed λ-calculus enriched with the tools to operate on hyperintensions, i.e., either to execute them or to present them as arguments. Its *semantics* is a procedural one that conceives of meanings as abstract *procedures*.[29] This means that TIL λ-terms denote procedures producing functions rather than denoting the functions themselves. A procedure is structured in a manner that details which logical operations of which types apply to which operands of which types. Typically, the output of one operation will serve as input for another operation. Given the types of the operations and the operands, it can be calculated which type of object the procedure is structured to produce. A simple example: if the procedure specifies the application of a function taking two numbers to a truth-value then the procedure is typed to produce a truth-value. A truth-value is obtained by picking two numbers and applying the function to them. The fact that there may be no object of a particular type as output does not detract from there being a procedure typed to produce an object of this type. Our procedures specify what *types* of entities to manipulate in what ways in order to produce some particular *type* of entity. The *atomic* procedures are of one step and provide their products (i.e., the entities they are typed to produce) as input objects on which molecular procedures operate. The *molecular* procedures are of two steps or more and detail how to proceed from input to output (or in the direction of output, if there is none). Of the procedures we define below, two of them may fail to yield a product.

The connection between procedures and hyperpropositions is that hyperpropositions are identified with particular procedures.[30] Those of the procedures that are hyperpropositions are either those that are typed to produce truth-values or are typed to produce truth-conditions, where truth-conditions are typed as functions from possible worlds to a partial function from times to truth-values.

We account for propositional structure in terms of procedural structure. We must explain how finely individuated procedurally structured propositions are. Otherwise, we cannot know which prospective input operands are admissible in an extensional logic of hyperintensions, for we would not know which substitutions would be valid. Our principle of granularity is quite strict, contributing to an exact calibration of the entities that can be quantified over. The principle is called *procedural isomorphism* and is an obvious nod to its predecessors, Carnap's intensional isomorphism and Church's synonymous isomorphism. Procedural isomorphism will be presented formally in Sect. 4.1.1.

[29] Moschovakis (2006) characterises meanings as generalised algorithms. Our procedures likewise qualify as *generalised* algorithms, because they are procedures that need not be *effectively* computable (thereby perhaps straining the idea of an algorithm a bit).

[30] This qualifies TIL as a *reductionist* theory of hyperpropositions, because hyperpropositions are 'reduced' to instances of a more general sort of entity instead of being *sui generis*. TIL is also a reductionist theory of propositions, i.e., the truth-conditions of empirical sentences, because TIL identifies them with functions from possible worlds to functions from times to truth-values.

3.2 Invariance and transparency

The title of the paper promises that quantifying-in is *transparent*, and it is because quantifying-in does not trigger shift of denotation and, thus, does not induce opacity. What *contextualist* semantic theories have got right is that another entity than the one being salient in extensional contexts needs to be picked out. What contextualism is wrong about is its tenet that each term or expression must have a bespoke semantics to suit each different sort of context. Thus, we are not erecting a tower of increasingly indirect senses and denotations. In our context-invariant semantics, a term or expression retains its fixed sense and denotation. This explains why we need devices that can insert a term's or expression's meaning rather than its denotation into argument position so as to make the displayed parts themselves amenable to logical manipulation.

This is the semantic schema of TIL:

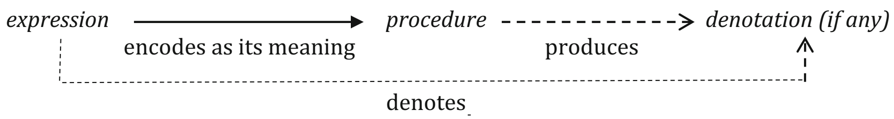

The relation of *encoding*, or *expressing*, between an expression and the procedure assigned to it as its meaning is semantically primary. Once we have the meaning procedure, we are in a position to examine which object the procedure produces, prove what is entailed by the procedure, examine its structure, etc. The semantic relation of denoting between an expression and a denotation (if any) piggybacks on the logical relation of producing between a procedure and its product (if any).

The semantics of definite descriptions, predicates and all other terms must be *top-down* for full referential transparency. This is to say that the above semantic schema that illustrates the relation between a term, the procedure that is its meaning, and its denotation (if any) in a hyperintensional context is the same schema that applies to intensional and extensional contexts. A term or expression expresses a (privileged) procedure as its meaning and denotes the entity (if any) that the procedure is typed to produce.[31] The denoted entity (if any) can be an object of one of the two basic kinds already sketched at the beginning of this part:

(i) a *non-procedural* entity, i.e., an object of a type of order 1 (see below), which comprises all partial *functions*, neither the domain nor range of which contains any procedures;
(ii) a *procedure* of a lower order in the type hierarchy (see below) than that of the relevant meaning procedure, in which case the produced (hence, denoted) procedure occurs as a functional argument.

Abstract procedures cannot be sets, tuples or aggregates of instructions, because sets, tuples or aggregates cannot be *executed*. Rather, they are structured wholes that

[31] *Indexicals* being the only exception: while the sense of an indexical remains constant (i.e., as a free variable with a type assignment), its denotation varies in keeping with its contextual embedding. See Duží et al. (2010, §3.4).

themselves can be executed.[32] Importantly, empirical terms, such as the definite description 'the Bishop of Rome' or the predicate 'is a planet', never denote their extension at any world/time pair, a fortiori not the extension in the actual world at the present time. The world/time-relative extensions (if any) fall outside the purview of the semantics. Empirical terms invariably denote the condition that an individual, a set, etc., must satisfy in order to be (in) its extension at the world/time pair of evaluation. We model these conditions as possible-world intensions.[33]

In TIL we reserve the terms 'refer' and 'reference' for the factual and extra-semantic relation between an empirical term and the value of the denoted intension at a given $\langle w, t\rangle$ pair. Thus, 'the man in the brown hat' and 'the man on the beach' co-*refer* to Mr Orcutt, as a matter of extra-semantic *fact*.[34]

3.3 Variables and Trivialisation

The two atomic 'feeder' procedures are:

- *Variable*
- *Trivialisation*

TIL deviates in four relevant respects from the version of λ-calculus made popular by Montague's Intensional Logic. First, *meanings* are not identified with (or modelled as) mappings from world/time pairs. Instead Montague-like meanings (i.e., mappings) are the products of our meaning procedures.

Second, *variables* are not linguistic items. The term 'y' expresses an atomic procedure as its meaning and picks out the entity that an assignment function has assigned to y as its value. Thus, three entities are involved: a term, a variable (a procedure), a value. We are adopting an objectual version of Tarski's conception of variables. Our objectual variables are procedures that produce entities dependently on *valuations*, i.e. assignment functions; we say that variables v-produce. Countably many variables are assigned to each type (see below). Moreover, entities of each type can be organised into sequences of countably many elements. Valuation v picks up one such sequence, and the ith variable v-produces the ith element of the sequence.

Third, the analysis of a piece of language, and this includes "b believes hyperproposition A", does not amount to translating it from some natural language into an artificial language (say, the λ-calculus), which in turn receives an interpretation, which is transferred back to the natural-language sentence. Instead our λ-calculus is an inherently interpreted formal language, which serves as a device to directly denote meaning procedures. Our λ-terms denote procedures. Meanings are studied by studying their structure and constituents as encoded in the λ-calculus of TIL in virtue of the stipulated isomorphism between formulae and meanings. It

[32] For more details on the character of these structured wholes and their mereology, see Duží (2019) and Jespersen (2019).

[33] This program of anti-actualist semantics is described in Duží et al. (2010, §2.4.1).

[34] For further details, see Duží et al. (2010, 301–11).

should be stressed again that our hyperpropositional procedures are not linguistic entities; they are higher-order abstract entities.

Fourth, TIL comes with *explicit intensionalisation and temporalisation*. See Sect. 3.5.1 below.

The ideography of TIL also comes with *constants*, which are vehicles of reference that pick out a specific entity in one go without the assistance of other terms and without invoking anything descriptive. We can potentially develop constants for any entity of any type, including hyperpropositions. The semantic counterpart of a constant is a *Trivialisation*. A Trivialisation is a procedure that picks out a specific entity in one go. Where 'Pluto' denotes Pluto (typed as an individual), the Trivialisation of Pluto, $^{0}Pluto$, is a one-step procedure for identifying Pluto. The procedure, just like a non-descriptive proper name, does not specify how to identify the object in question. Trivialisation embodies merely the procedure of 'reaching' into a particular type (here, the type of individuals) and 'extracting' a particular object (here, Pluto) from there. A variable, by contrast, embodies the procedure of 'reaching' into a particular type and 'extracting' an arbitrary object from there.

3.4 Composition, Closure, Double Execution

The two complex, or multi-step, procedures are:

- *Composition*
- *Closure*

Composition is the procedure of functional application, rather than the functional value (if any) resulting from the application. Closure is the procedure of functional abstraction, rather than the resulting function.

TIL contains a duo of explicit Execution procedures, which include these two:

- *Single Execution*
- *Double Execution*

Single Execution, ^{1}X, is part of Tichý's inductive definition of procedures (called *constructions*) in (1988, §15), but has been left out of the definition below. It is not needed for present purposes because, importantly, the default mode in which procedures occur is as executed. ^{1}X is the same procedure (though of a higher order in the type hierarchy) as X, provided X is a procedure at all.

Double Execution, ^{2}X, encodes the transitivity of descending from procedure to product (if any). Double Execution is complex, provided X is a procedure at all. Double Execution will appear in some type specifications in the interest of clarification.

Here is the inductive definition of *procedure*.

Definition 1 (*procedure*)

(i) *Variables* x, y, … are *procedures* that *produce* objects (elements of their respective ranges) dependently on a valuation v; they *v-produce*.

(ii) Where X is an object whatsoever (an extension, an intension or a *proce-*

dure), ${}^0X$ is the *procedure Trivialisation.* ${}^0X$ *produces (displays) X* without any change of X.

(iii) Let X, Y_1, …, Y_n be arbitrary *procedures*. Then *Composition* $[X\ Y_1 \ldots Y_n]$ is the following *procedure*. For any valuation v, the *Composition* $[X\ Y_1 \ldots Y_n]$ is *v-improper* if at least one of the *procedures* X, Y_1, …, Y_n is v-improper by failing to v-produce anything, or if X does not *v-produce* a function that is defined at the n-tuple of objects *v-produced* by $Y_1,\ldots,Y_n$. If X does *v-produce* such a function, then $[X\ Y_1 \ldots Y_n]$ *v-produces* the value of this function at the n-tuple.

(iv) The (λ-) *Closure* $[\lambda x_1 \ldots x_m\ Y]$ is the following *procedure*. Let x_1, x_2, …, x_m be pair-wise distinct variables and Y a *procedure*. Then $[\lambda x_1 \ldots x_m\ Y]$ *v-produces* the function f that takes any members B_1, …, B_m of the respective ranges of the variables x_1, …, x_m into the object (if any) that is $v(B_1/x_1,\ldots,B_m/x_m)$-*produced* by Y, where $v(B_1/x_1, \ldots, B_m/x_m)$ is like v except for assigning B_1 to x_1, …, B_m to x_m.

(v) The *Double Execution* ${}^2X$ is the following *procedure*. Where X is any entity, the *Double Execution* ${}^2X$ is *v-improper* if X is not itself a *procedure*, or if X does not v-produce a *procedure*, or if X *v-produces* a v-improper *procedure*. Otherwise, let X *v-produce* a *procedure* Y and Y *v-produce* an entity Z; then ${}^2X$ *v-produces* Z.

(vi) Nothing is a *procedure*, unless it so follows from (i) through (v). □

3.5 Typed universe

TIL comes with a thoroughly typed universe. The ground floor is populated by first-order, non-procedural entities. First-order entities are functions, whether 'typical' functions such as mathematical functions, or characteristic functions and the functions of possible-world semantics, or nullary functions such as individuals and truth-values. Note that intensions—functions from possible worlds—are typed as *first*-order entities, unless they are functions that contain in their range or domain a procedure such as a hyperpropositional attitude, which is typed as a relation-in-intension of an individual to a procedure.[35]

Definition 2 (*simple type*) Let B be a *base*, where a base is a collection of pair-wise disjoint, non-empty sets. Then:

(i) Every member of B is an elementary *type of order 1 over B*.

(ii) Let α, β_1, ..., β_m ($m > 0$) be types of order 1 over B. Then the collection $(\alpha\ \beta_1 \ldots \beta_m)$ of all m-ary partial mappings from $\beta_1 \times \ldots \times \beta_m$ into α is a functional *type of* order *1 over B*.

(iii) Nothing is a *type of order 1 over B* unless it so follows from (i) and (ii). □

[35] First-*order*, higher-*degree* intensions are defined as functions from intensions to functions that contain an intension in their domain or range.

For the purposes of natural-language analysis, we are currently assuming the following base of ground types, which form part of the ontological commitments of TIL:

ο: the set of truth-values {**T**, **F**}
ι: the set of individuals (the universe of discourse)
τ: the set of real numbers (doubling as times)
ω: the set of logically possible worlds (the logical space)

Hence, we are able to type both extensional and intensional entities. *Intensions* are polymorphous functions of type $(\beta\omega)$, where β is frequently a *chronology* of α-objects of type $(\alpha\tau)$; thus α-intensions are frequently of type $((\alpha\tau)\omega)$, which will be abbreviated as '$\alpha_{\tau\omega}$'. An object of type $\alpha_{\tau\omega}$ is a function from worlds to a function from times to α-typed objects.[36]

Some important extensional and intensional entities include:

- Characteristic function (here, *set of individuals*)/$(\omicron\iota)$
- Set-in-intension (here, *property of individuals*)/$(\omicron\iota)_{\tau\omega}$
- Individual-in-intension (*individual office* or *role*)/$\iota_{\tau\omega}$
- Truth-value-in-intension (*truth-condition* or *PWS proposition*)/$\omicron_{\tau\omega}$
- Binary relation-in-intension (here, *attitude*)/$(\omicron\iota\alpha)_{\tau\omega}$

Note that if the complement of the attitude is a truth-condition, as per standard possible-world semantics, then the type of the attitude is $(\omicron\iota\omicron_{\tau\omega})_{\tau\omega}$, because $\alpha = \omicron_{\tau\omega}$. If the complement of the attitude is a hyperproposition, as in this paper, then $\alpha = *_n$, hence $(\omicron\iota *_n)_{\tau\omega}$ is the type of the attitude. The higher-order types $*_n$ fall outside the purview of Definition 2 and are the province of the ramified type hierarchy, as per Definition 3 below.

3.5.1 Explicit intensionalisation and temporalisation

One of the key features of how TIL analyses natural-language discourse is what we call explicit intensionalisation and temporalisation. It is developed in opposition to Montague's intensional logic. Montague's IL comes with the 'half-hearted' type *s*, where *s* is the type of world/time pairs. *s* is neither a ground type nor a functional type, but only occurs as a fragment of functional types, as in $(s\rightarrow t)$. TIL, by contrast, contains a full-fledged type ω for worlds and a full-fledged type τ for instants of time. This enables formulas like "λw [… w …]", "λt [… t …]", "λwt [… w … t …]", and "$\lambda w \lambda t$ [… w … t …]", where it is explicit from the syntax that the evaluation takes place at empirical indices.[37] These indices are not relegated to a meta-language, as is common in other formal-semantic frameworks for natural-language discourse. The non-vacuous occurrence of at least one '*w*' or '*t*', denoting

[36] See Duží et al. (2010, §2.5).

[37] For critical comments on Montague's IL and a comparison with TIL, see Duží et al. (2010, §2.4.3).

a modal or temporal variable, in the syntax is what marks the difference between empirical and non-empirical language.

Empirical languages incorporate an element of *contingency*, because they denote *empirical conditions* that may or may not be satisfied at some world/time pair of evaluation. For instance, it is only at some world/time pairs that a given individual entertains a given attitude. Our explicit intensionalisation and temporalisation enables us to encode procedures producing possible-world intensions directly in the logical syntax in virtue of terms for world and time variables. Where variable w ranges over possible worlds (type ω) and t over times (type τ), the following logical form essentially characterises the logical syntax of empirical language:

$$\lambda w \lambda t\, [\ldots w \ldots . t \ldots]$$

The above schematic Closure is typed to produce a condition satisfiable by world/time pairs.[38] Here is the (privileged or canonical) form of the hyperproposition that Pluto is a planet:[39]

$$\lambda w \lambda t\, [{}^0 Planet_{wt}\; {}^0 Pluto]$$

The above Closure is a hyperproposition that produces a truth-condition/$o_{\tau\omega}$. The Closure does what it does by abstracting over the respective values of the variables w, t. The purpose is to isolate exactly those worlds and times at which it is true that Pluto is an element of the respective (i.e., world-and-time-relative) set of planets (assuming that a crisp definition of planethood is in place). The general flow of the procedure is to break down and then build up again. An analysis will spell out what is going on.[40] The above procedure contains three occurrences of Composition ('breaking down'):

[1] $[{}^0 Planet\ w]$: the application of $Planet/(o\iota)_{\tau\omega}$ to a possible world v-produced by $w \rightarrow \omega$ to obtain a function of type $((o\iota)\tau)$, a chronology which inputs instants of time and outputs the respective sets of planets at those particular times.

[2] $[[{}^0 Planet\ w]t]$, or ${}^0 Planet_{wt}$, for short: the application of the chronology obtained at [1] to a time v-produced by $t \rightarrow \tau$ to obtain a set of individuals/$(o\iota)$.

[3] $[[[{}^0 Planet\ w]t]\ {}^0 Pluto]$, or $[{}^0 Planet_{wt}\ {}^0 Pluto]$, for short: the application of the set obtained at [2] to Pluto/ι to obtain a truth-value/o.

The truth-value obtained is relativised to worlds and times by means of two instances of Closure ('building up') to obtain a truth-condition.

[38] When speaking of 'world/time pairs', we are allowing ourselves to pretend that a function from worlds to a function from times to entities is equivalent to a binary function from world/time pairs to entities. This pretence is innocuous in this essay, because here we are not considering the modal and the temporal dimension separately. See Duží et al. (2010, §2.5). Moreover, in a logic of partial functions, such as TIL, *schönfinkelisation* fails to always preserve equivalence: see Duží et al. (2010, 204–05).

[39] By 'privileged' or 'canonical' form we intend the *literal* analysis of a sentence, where syntactically simple terms like 'Pluto' and 'planet' are paired off with a Trivialisation of the denoted object, here ^{0}*Pluto*, ^{0}*Planet*. For the notion of literal analyses, see Duží et al. (2010, 105, Defs. 1.10, 1.11).

[40] This exposition relies on Duží et al. (2010, §2.4.2).

3.5.2 Ramified type hierarchy

The relevant logical feature of our ramified hierarchy of types is that we are guaranteed to always have a hyperintension of a higher order at our disposal that will present a hyperintension of a lower order. Without this possibility, we would be falling short of the expressive power required to pull off quantifying-in, once hyperintensions are construed as higher-order entities rather than primitive first-order entities.

The definition of the ramified hierarchy of types decomposes into three parts: firstly, simple types of order 1; secondly, procedures of order n; thirdly, types of order $n+1$.

Definition 3 (*ramified hierarchy of types*)
T_1 (*types of order 1*). See Definition 2.
C_n (*procedures of order n*)

(i) Let x be a variable ranging over a type of order n. Then x is a *procedure of order n over B.*
(ii) Let X be a member of a type of order n. Then ${}^0X$, ${}^2X$ are *procedures of order n over B.*
(iii) Let $X, X_1, \ldots, X_m$ ($m > 0$) be *procedures* of order n over B. Then $[X\, X_1 \ldots X_m]$ is a *procedure of order n over B.*
(iv) Let $x_1, \ldots, x_m, X$ ($m > 0$) be *procedures* of order n over B. Then $[\lambda x_1 \ldots x_m\, X]$ is a *procedure of order n over B.*
(v) Nothing is a *procedure of order n over B* unless it so follows from $\mathbf{C}_n$ (i–iv).

T_{n+1} (*types of order n*+1). Let $*_n$ be the collection of all procedures of order n over B. Then:

(i) $*_n$ and every type of order n are *types of order* $n + 1$.
(ii) If $m > 0$ and $\alpha_1, \ldots, \alpha_m$ are types of order $n + 1$ over B, then $(\alpha_1 \ldots \alpha_m)$ (see T_1 (ii)) is a *type of order* $n + 1$ *over B*.
(iii) Nothing is a *type of order n + 1 over B* unless it so follows from T_{n+1} (i) and (ii). □

Notational conventions. '$y \rightarrow \alpha$' means that variable y ranges over the type α. If C is a procedure, then '$C \rightarrow \alpha$' means that C is typed to produce an entity of type α. That an object a is of a type α is denoted 'a/α'. Thus, for instance, '$C/*_n \rightarrow \iota$' means that the procedure C is of order n (i.e. belongs to type $*_n$) and is typed to produce an individual. Throughout this paper we use variables $w \rightarrow \omega$ and $t \rightarrow \tau$. If $C \rightarrow \alpha_{\tau\omega}$ then the frequently used Composition $[[C\ w]\ t] \rightarrow \alpha$ will be written as 'C_{wt}' for short.

Definition 4 (*existential and universal quantifiers*). The *existential quantifier* $\exists^{\alpha}/(o(o\alpha))$ is a total polymorphic function that takes a set of α-typed elements to the truth-value **T** if the set is non-empty and otherwise to **F**. The *general quantifier* $\forall^{\alpha}/(o(o\alpha))$ is a total polymorphic function that takes a set S of α-typed elements to the truth-value **T** if S contains all the elements of type α and otherwise to **F**. □

'∃', '∀' are categorematic terms in TIL, namely functors that denote functions of the above type. Once a set produced by, e.g., λy [… y …] is inputted as an argument to ∃ or ∀, the quantifier returns a truth-value as value. The strings '$\exists y$', '$\forall y$' count as ill-formed in TIL, because all binding is λ-binding or ^{0}-binding (see below for the definition). The proper notation is '$[^0\exists\lambda y$ [… y …]]', '$[^0\forall\lambda y$ [… y …]]'. Anyway, for the sake of simplicity, we may sometimes stick to '$\exists y$ [… y …]', '$\forall y$ [… y …]', when no confusion can arise.[41]

3.6 Displayed versus executed; free versus bound; valid substitution

We define here what it means for procedures to occur in the *displayed* mode and to occur in the *executed* mode, and we explain what it means for variables to have *free* and to have *bound* occurrences.

When a procedure occurs in the displayed mode, the *procedure* itself becomes an object on which other procedures can operate. We also say that the context of its occurrence is *hyperintensional*, because all the sub-procedures of a displayed procedure occur neither intensionally nor extensionally; they are displayed as well. When a procedure occurs in the executed mode, the *product* (if any) of the procedure is susceptible to being operated on. In this case, the executed procedure is a *constituent* of its super-procedure, and an additional distinction crops up at this level. A constituent producing a function may occur either *intensionally* (de dicto) or *extensionally* (de re). If intensionally, then the produced *function* is the object of predication; if extensionally, then the *value* (if any) of the produced function is the object of predication. The pair of distinctions between displayed/executed and intensional/extensional occurrences enables us to distinguish between three kinds of *context*. The rigorous definitions of the three kinds of contexts can be found in Duží et al. (2010, §2.6). Though the basic ideas are fairly simple, the exact details are rather complicated. For this reason, we only explain intuitively the main ideas here.

- *Hyperintensional context.* A procedure occurs in the *displayed* mode (though another procedure at least one order higher needs to be *executed* in order to produce the displayed procedure).

[41] It may be instructive to consider how TIL formalises the *Barcan Formula*: $\Diamond\exists xFx \supset \exists x\, \Diamond Fx$. There are two ways to go about this. Either we stick to $\Diamond$ or we turn to existential quantification over worlds (ignoring times). S5-possibility is typed as a property of propositions, of type (ο(οω)), which is not indexed to worlds, because the S5-modalities are analytic, being valid on equivalence frames. Sub-S5-modalities are not, so properties of propositions are indexed to worlds, and for this reason such properties are of type ((ο(οω))ω): see Materna (2005).

(i) $[[^0\exists\lambda w\ [^0\exists\lambda x\ [^0F_w\ x]]] \supset [^0\exists\lambda x\ [^0\exists\lambda w\ [^0F_w\ x]]]]$

(ii) $[[^0\Diamond\ \lambda w\ [^0\exists\lambda x\ [^0F_w\ x]]] \supset [^0\exists\lambda x\ [^0\Diamond\ \lambda w\ [^0F_w\ x]]]]$

Both formulas, on their intended interpretation, express that if some world has some individual with property F then some individual at some world has property F. It is obvious why the Barcan Formula (and its Converse) requires S5-possibility and that the domain function be constant: for all $w, w' \in W$, $D(w) = D(w')$. It is also obvious (as proved by running a type check) why neither of (i), (ii) engenders the problem that a λ-bound variable has an 'opaque' occurrence.

- *Intensional context*. A procedure occurs in the *executed* mode in order to produce a function rather than one of its values (moreover, the executed procedure does not occur within another hyperintensional context).
- *Extensional context*. A procedure occurs in the *executed* mode in order to produce a particular value of a function at a given argument (moreover, the executed procedure does not occur within another intensional or hyperintensional context).

We next turn to a definition of *sub-procedure*.

Definition 5 (*sub-procedure*). Let C be a procedure. Then:

(i) C is a *sub-procedure of* C.
(ii) If C is ${}^0X$ or ${}^2X$ and X is a procedure, then X is a *sub-procedure of* C.
(iii) If C is $[X\ X_1 \ldots X_n]$ then $X, X_1, \ldots, X_n$ are *sub-procedures of* C.
(iv) If C is $[\lambda x_1 \ldots x_n\ Y]$ then Y is a *sub-procedure of* C.
(v) If A is a *sub-procedure of* B and B is a *sub-procedure of* C then A is a *sub-procedure of* C.
(vi) A procedure is a *sub-procedure of* C only if it so follows from (i–v). □

In particular, the constituent parts of a procedure are not the particular material, or otherwise non-procedural abstract, objects the procedure operates on. Correspondingly, Mont Blanc cannot be a constituent of a procedure; ^{0}*Mont_Blanc* can. Instead, *the constituents of a procedure are exclusively those sub-procedures that occur in executed mode*. (See the 10-part decomposition below for illustration.) To define the distinction between displayed and executed mode, we must take the following factors into account. A procedure C can occur in displayed mode only as a sub-procedure within another procedure D that operates on C. Therefore, C itself must be produced by another sub-procedure C' in D. And it is necessary to define this distinction for *occurrences* of procedures, because one and the same procedure C can occur in executed mode in D and at the same time serve as an input/output object for another sub-procedure C' of D that operates on C. The distinction between displayed and execution mode of a procedure can be characterised like this. Let C be a sub-procedure of a procedure D. Then an occurrence of C *occurs in displayed mode in D* if the execution of D does not involve the execution of this occurrence of C. Otherwise, C occurs in *executed mode within D*, i.e., *C occurs as a constituent* of the procedure D.

To see how this works, consider this sentence:

$$\text{“Tilman solves the equation } Sin(x) = 0\text{”}$$

When solving the problem of seeking the numbers x such that the value of the function *Sine* at x equals zero, Tilman is not related to the set of multiples of the number π, i.e., to an object of type (οτ). If he were, Tilman would have already solved the problem, thus pre-empting the search for suitable values of x. Rather Tilman wishes to find the product of the procedure $[\lambda x\ [{}^0{=}\ [{}^0Sin\ x]\ {}^0 0]]$. In other words, the sentence expresses Tilman's relation-in-intension to this very procedure, *Solve* emerging as an object of type $(o\iota *_1)_{\tau\omega}$. Therefore, the whole sentence encodes this procedure:

$$\lambda w \lambda t\,[{}^0Solve_{wt}\,{}^0Tilman\,{}^0[\lambda x\,[{}^0{=}\,[{}^0Sin\,x]\,{}^00]]]$$

Types and type checking: ${}^0Solve \rightarrow (o\iota *_1)_{\tau\omega}$; ${}^0Solve_{wt} \rightarrow (o\iota *_1)$; ${}^0Tilman \rightarrow \iota$; ${}^0Sin \rightarrow (\tau\tau)$; ${}^0{=} \rightarrow (o\tau\tau)$; ${}^00 \rightarrow \tau$; $x \rightarrow \tau$; $[{}^0Sin\ x] \rightarrow \tau$; $[{}^0{=}\ [{}^0Sin\ x]\ {}^00] \rightarrow o$; $[\lambda x\ [{}^0{=}\ [{}^0Sin\ x]\ {}^00]] \rightarrow (o\tau)$; ${}^0[\lambda x\ [{}^0{=}\ [{}^0Sin\ x]\ {}^00]] \rightarrow *_1$; $[{}^0Solve_{wt}\ {}^0Tilman\ {}^0[\lambda x\ [{}^0{=}\ [{}^0Sin\ x]\ {}^00]]] \rightarrow o$; $\lambda w \lambda t\,[{}^0Solve_{wt}\,{}^0Tilman\,{}^0[\lambda x\,[{}^0{=}\,[{}^0Sin\,x]\,{}^00]]] \rightarrow o_{\tau\omega}$.

The *procedure* $[\lambda x\ [{}^0{=}\ [{}^0Sin\ x]\ {}^00]]$, which is the meaning of the term "The equation *Sin*(*x*) equals 0", is *displayed* (by means of *Trivialisation*) as the second argument of the relation $Solve_{wt}$. The evaluation of the truth-conditions expressed by the sentence consists in checking, for any possible world *w* and for any time *t*, whether Tilman and this procedure occur in the extensionalised relation-in-intension of *solving* as its first and second argument, respectively. Hence the execution of the hyperproposition expressed by the sentence does not involve the execution of the procedure of solving the equation; this is something Tilman is tasked with.

The *execution* steps specified by the above Closure, i.e., its *constituents*, are as follows. Each procedure is an executed part of itself, hence the Closure (1) is a constituent of itself.

(1) $\lambda w \lambda t\ [{}^0Solve_{wt}\ {}^0Tilman\ {}^0[\lambda x\ [{}^0{=}\ [{}^0Sin\ x]\ {}^00]]]$
(2) $\lambda t\ [{}^0Solve_{wt}\ {}^0Tilman\ {}^0[\lambda x\ [{}^0{=}\ [{}^0Sin\ x]\ {}^00]]]$
(3) $[{}^0Solve_{wt}\ {}^0Tilman\ {}^0[\lambda x\ [{}^0{=}\ [{}^0Sin\ x]\ {}^00]]]$
(4) ${}^0Solve_{wt}$
(5) $[{}^0Solve\ w]$
(6) ${}^0Solve$
(7) w
(8) t
(9) ${}^0Tilman$
(10) ${}^0[\lambda x\ [{}^0{=}\ [{}^0Sin\ x]\ {}^00]]$

It might seem that in order to define rigorously the distinction between displayed and executed occurrence it would suffice to say that a procedure occurs displayed within another procedure if it occurs within the scope of a Trivialisation. Yet it is not that simple. The complicating factor is that Trivialisation has a *dual* procedure, namely Double Execution. It follows from Definition 1 that while Trivialisation of a procedure raises the context to the hyperintensional level, Double Execution cancels the effect of Trivialisation, thus lowering the level of the context. The reason is this law of ^{20}-elimination:

$$ {}^{20}\mathrm{C} = \mathrm{C} $$

for any procedure *C*.

Note, however, that ${}^{02}C$ is *not* equivalent to *C*, because ${}^{02}C$ produces ${}^2C$ for any valuation *v*. The procedure ${}^2C$, and thus also the procedure *C*, occurs displayed in ${}^{02}C$ because they occur within the scope of Trivialisation. Another complicating factor is iteration of Double Execution and Trivialisation. For instance, in ^{2200}C procedure C occurs in executed mode while in ^{2002}C procedure C occurs in displayed mode. The reason is that by applying the above law twice, we get $^{2200}C{=}C$, but the law can be applied only once to ^{2002}C, thus obtaing $^{2002}C = {}^{02}C$.

For the purposes of this paper, we shall say that a procedure occurs displayed exactly when it occurs within the scope of a Trivialisation, whose effect is not cancelled by Double Execution. To simplify the definition of displayed versus executed mode, we first define *^{20}-normal form of a procedure*.

Definition 6 (*^{20}-elimination, ^{20}-normal form*).
Let X be a procedure and let ^{20}X occur as a sub-procedure of a procedure C. Then the replacement of an occurrence of ^{20}X by the procedure X, the result of which is a procedure C', is called *^{20}-elimination*, $C \rightarrow_{20} C'$. If a procedure C can be transformed into a procedure D by a finite (or empty) sequence of ^{20}-eliminations, so that D does not contain any occurrence of a sub-procedure ^{20}X, then D is the *^{20}-normal form of C*. □

Definition 7 (*occurrence of a procedure in displayed vs. executed mode*).
Let C and D be procedures in ^{20}-normal form and let D be a sub-procedure of C. Then:

(i) If C is identical to D then D *occurs in C in executed mode.*
(ii) If C is identical to ^{0}X, and D is a sub-procedure of X, then D *occurs in C in displayed mode*.
(iii) If C is identical to $[X_1\ X_2 \ldots X_m]$ and D is a sub-procedure of the X_i, for some i, $1 \leq i \leq m$, then the *occurrence of D in C* is the same as the occurrence of *D in X_i*.
(iv) If C is identical to $[\lambda x_1 \ldots x_m\ X]$ and D is a sub-procedure of X, then the *occurrence of D in C* is the same as the occurrence of *D in X*.
(v) If C is identical to ^{1}X or ^{2}X and D is a sub-procedure of X, then the *occurrence of D in C* is the same as the occurrence of *D in X*.
(vi) The *occurrence of D* in C is in *displayed* or *executed mode* only due to (i–v). □

Remark For the sake of simplicity, we define the occurrence in displayed or executed mode only for procedures in their ^{20}-normal form. This means that we first eliminate those pairs of Double Execution and Trivialisation (in this order) that can be validly eliminated, and then the occurrence of a procedure D in C is in displayed mode if D occurs within the scope of Trivialisation (see Definition 7 (ii)). Thanks to this simplification, Definition 7 leaves undefined the occurrences of, for instance, ^{200}X in ^{2200}X, ^{00}X in ^{2200}X, or ^{0}X in ^{2200}X. All of these sub-procedures of ^{2200}X occur in executed mode, as does the procedure X in ^{2200}X (according to Definition 6 (i)). Yet, for the purpose of this paper, this simplification is harmless. □

Definition 8 (*constituent of a procedure*). Let C be a procedure and D a sub-procedure of C. The executed occurrences of D are *constituents* of C. □

Corollary *Each procedure C is a constituent of itself, namely its improper constituent. All the other constituents of C are its proper constituents. (The notions of being proper or improper constituents should not be confused with the notions of being v-proper or v-improper procedures.)*

Analogously to formal languages, variables can occur free or bound within a procedure. Importantly, though, it is *occurrences* of variables that occur free or bound. This is because the same variable can have both free and bound occurrences within the same context. An occurrence can also be doubly bound, as when a λ-bound occurrence is also ^{0}-bound. Yet in the case of doubly-bound variables, we say that the variable is simply ^{0}-bound, because if a procedure C is displayed then all its sub-procedures, including its variables, are displayed as well and therefore ^{0}-bound. The general rule is that a higher context is dominant over a lower one. Thus, we define:

Definition 9 (*free variable, bound variable, open/closed procedure*).
Let C be a procedure with at least one occurrence of a variable x. Then:

(i) If the occurrence of x in C is in the displayed mode, then this *occurrence of x is ^{0}-bound in C*.
(ii) Let the occurence of x in C be in the executed mode and let the Closure $[\lambda x_1 \ldots x_m\ X]$ be a sub-procedure of C. If this occurrence of x is a sub-procedure of X and x is identical to one of the variables x_i, $1 \le i \le m$, then this *occurrence of* x is *λ-bound in C*.
(iii) If the occurrence of x is neither ^{0}-bound nor λ-bound in C, then this *occurrence of x is free in C*.
(iv) An *occurrence* of x is *free*, *λ-bound*, *^{0}-bound* in C only due to (i–iii).

A procedure with at least one occurrence of a free variable is an *open procedure*. A procedure without any occurrences of free variables is a *closed procedure*. □

Corollary *If a procedure D occurs in the displayed mode in C, then all the variables occurring in D are ^{0}-bound in C.*

Definition 10 (*correct substitution*).
Let x be a variable and C, D procedures in ^{20}-normal form. If the variable x is not free in C, then the result of the substitution of D for x in C is C. Let now the variable x be free in C. Then:

(i) If C is identical to x, then the *result of the substitution of D for x in C* is D.
(ii) If C is identical to $[X\ X_1 \ldots X_n]$, then the *result of the substitution of D for x in C* is $[Y\ Y_1 \ldots Y_n]$, where $Y, Y_1, \ldots, Y_n$ are the results of the substitution of D for x in $X, X_1, \ldots, X_n$, respectively.
(iii) If C is identical to $[\lambda x_1 \ldots x_m\ Y]$, then for $1 \le i \le m$ let y_i be x_i if x_i is not free in D, otherwise let y_i be the first variable v-producing entities of the same type as x_i such that y_i does not occur in C or D and distinct from $y_1 \ldots y_{i-1}$. Then the *result of the substitution of D for x in C* is $[\lambda y_1 \ldots y_m\ Z]$, where Z is

the result of the substitution of D for x in the result of the substitution of y_i for x_i $(1 \leq i \leq m)$ in Y.

(iv) If C is identical to ^2X, where X is a procedure, then the *result of the substitution of D for x in C* is ^2Y, where Y is the result of the substitution of D for x in X. □

Remark The procedure that is the result of the substitution of D for x in C will be denoted '$C(D/x)$'.

We proceed to define two functions already adumbrated above, namely *Sub* and *Tr*, which are needed to address the technical difficulties of quantifying into hyperintensional contexts. These difficulties stem from the fact that all the variables occurring within the hyperintensional context of a displayed procedure are ^{0}-bound. Hence, the objects v-produced by such variables are irrelevant, and so is their λ-binding. Yet, in order to quantify *into* a hyperintensional context, we need a mechanism that makes it possible to operate on displayed procedures; in particular, we need to substitute for a displayed variable. To this end, we have developed a substitution method that makes use of *Sub* and *Tr*.

The polymorphous function $Sub/(*_n *_n *_n *_n)$ operates on procedures: one procedure is substituted for another within a third procedure, thus yielding a fourth procedure. Formally:

Definition 11 (Sub^n). Let $Q_1/*_{n+1} \rightarrow *_n, Q_2/*_{n+1} \rightarrow *_n, Q_3/*_{n+1} \rightarrow *_n$ v-produce procedures P_1, P_2, P_3, resp., where P_2 is a variable. Then the Composition $[^0Sub^n\ Q_1\ Q_2\ Q_3]$ v-produces the procedure P_4 that is the result of a simultaneous correct substitution of P_1 for all occurrences of P_2 in P_3. □

In what follows we will omit the superscript n whenever no confusion can arise.

Definition 12 (*Tr*). The polymorphic function $Tr^\alpha/(*_n\ \alpha)$ returns as its value the Trivialisation of its α-typed argument. □

In what follows we will omit the superscript α whenever no confusion can arise.

Examples *Tr* takes Pluto/ι to its Trivialisation of type $*_1$. Thus, the product of the Composition $[^0Tr\ ^0Pluto]$ is ^0Pluto. The Composition $[^0Tr\ y]$ *v(Pluto/y)-produces* ^0Pluto. Where $Planet/(o\iota)_{\tau\omega}$, the Composition $[^0Sub\ [^0Tr\ y]\ ^0x\ ^0[^0Planet_{wt}\ x]]$ $v(Pluto/y)$-produces the Composition $[^0Planet_{wt}\ ^0Pluto]$, which produces a truth-value (which one depends on the particular values chosen for w, t).

Notice the substantial difference between *Trivialisation*, which is a procedure, and *Tr*, which is a function. ^0y produces the variable y regardless of valuation: ^0y just displays y without executing the procedure y. Thus, the variable y is *bound* by Trivialisation in ^0y. On the other hand, $[^0Tr\ y]$ v-produces the Trivialisation of the object v-produced by y. Hence, y occurs *free* in $[^0Tr\ y]$ and can be λ-bound.

Example Let variable $y \rightarrow \tau$. Then $[^0Tr\ y]$ $v(\pi/y)$-produces $^0\pi$. The Composition $[^0Sub\ [^0Tr\ y]\ ^0x\ ^0[^0Cot\ x]]$ $v(\pi/y)$-produces the Composition $[^0Cot\ ^0\pi]$. Hence, the Composition $[^0Sub\ [^0Tr\ y]\ ^0x\ ^0[^0Cot\ x]]$ is $v(\pi/y)$-congruent with $[^0Cot\ ^0\pi]$. Importantly, the variable y is *free* for λ-binding in the former, unlike the variable x that is ^{0}-bound.

Definition 11 of the substitution function is also helpful in explaining why our objectual manner of addressing quantifying-into hyperintensional contexts does not inherit the problems of attempting to quantify into *quotational contexts*. Objectual quantifying-in contrasts with quantifying into quotations. Since our hyperintensional λ-calculus tracks syntactic structure closely (though not slavishly, with procedural isomorphism serving to soak up semantically redundant differences; see Sect. 4.1.1), why are we immune to the notorious problems of quantifying into quotation contexts? In an objectual hyperintensional context an expression E is *used* to express the procedure that is its meaning, and this procedure is, furthermore, *displayed* as an argument susceptible to logical manipulation. In a quotational context, E is just *mentioned*, which renders E semantically inert, whereby its meaning plays no role at all and so is neither displayed nor used. Whereas in our logic we can operate on displayed procedures, we never get around to operating on mentioned expressions. However, when we do operate on displayed hyperpropositional procedures, we need a technique to work around the fact that the constituent sub-procedures we want to manipulate also occur displayed. The substitution method (see Definition 11) makes it possible to enter displayed hyperpropositions, extract parts and insert other parts in their place.

3.7 Inference and entailment

Attitudinal sentences being empirical, we need to define analytical entailment between empirical hyperpropositions, i.e., procedures that produce truth-conditions of type $o_{\tau\omega}$. We first characterise entailment in prose followed by a definition. A hyperproposition P is entailed by the hyperpropositions $Q_1, \ldots, Q_n$, iff necessarily, i.e., in all possible worlds and at all times where all the assumptions $Q_1, \ldots, Q_n$ produce true propositions (i.e., satisfied truth-conditions), the hyperproposition P produces a true proposition/satisfied truth-condition as well.

TIL being a logic of *partial* functions, it is apt for dealing with presuppositions, truth-value gaps, non-referring terms and other phenomena of natural language. Yet partiality, as we all know all too well, brings about technical complications. In particular, propositions can take the truth-value **T** at some worlds and times, **F** at others, and at yet other worlds and times have truth-value gaps. Hence, if not all the assumptions of an argument are true, some may be false and others gappy. Consequently, entailment in a logic of partial functions is truth-preserving from premises to conclusion, but not falsity-preserving from conclusion to premises. To manage partiality properly, we need the empirical propositional property of *True*. For completeness, we also define two other properties, namely *False* and *Undefined*, all of type $(oo_{\tau\omega})_{\tau\omega}$. They are defined as follows ($P \rightarrow o_{\tau\omega}$):

$$[{}^0True_{wt}P] \ v\text{-produces } \mathbf{T} \text{ if } P_{wt}, \text{ otherwise } \mathbf{F};$$
$$[{}^0False_{wt}P] \ v\text{-produces } \mathbf{T} \text{ if } \neg P_{wt}, \text{ otherwise } \mathbf{F};$$
$$[{}^0Undefined_{wt}P] = \neg [{}^0True_{wt}P] \wedge \neg [{}^0False_{wt}P].$$

Definition 13 (*Analytical entailment*).

Let $P, Q_1, \ldots, Q_n \rightarrow o_{\tau\omega}$ be hyperpropositions. Then P is *entailed* by $Q_1, \ldots, Q_n$, denoted $Q_1, \ldots, Q_n \models P$, iff $\forall w \forall t\ [[[^0True_{wt}\ Q_1] \wedge \ldots \wedge [^0True_{wt}\ Q_n]] \supset [^0True_{wt}\ P]]$.□

Note that if in Definition 13 we had not applied the property *True*, and instead used simply the Composition $[[Q_{1wt} \wedge \ldots \wedge Q_{nwt}] \supset P_{wt}]$, the whole Composition $\forall w \forall t\ [[Q_{1wt} \wedge \ldots \wedge Q_{nwt}] \supset P_{wt}]$ would produce **F**. The reason is that at those $\langle w, t \rangle$-pairs where at least one of the Q_{iwt} is v-improper, the whole Composition is v-improper, due to partiality being propagated up.

The last technical devices that we need are λ-introduction and elimination of the left-most $\lambda w \lambda t$. These are applied when dealing with empirical hyperpropositions. If the assumptions are empirical hyperpropositions, our task is then to infer the hyperproposition that is logically entailed by the hyperpropositions in the premises. Entailment means that at any world w_0 and time t_0 of evaluation, the derivation sequence must be truth-preserving from premises to the conclusion. Thus, the typical sequence of derivation steps is this. We have assumptions of the form $\lambda w \lambda t$ $[\ldots w \ldots t \ldots]$ producing entities of type $o_{\tau\omega}$, and we assume that the propositions produced by these procedures are true at the world w_0 and time t_0 of evaluation. Using the detailed notation, we obtain the Composition

$$[[[\lambda w\, [\lambda t\, [\ldots w \ldots t \ldots]]] w_0]\, t_0]$$

which produces an o-object, i.e., a truth-value. By applying restricted β-reduction twice, we eliminate the leftmost $\lambda w \lambda t$, thus obtaining $[\ldots w_0 \ldots t_0 \ldots] \rightarrow o$.[42] Now we proceed with derivation steps, until the conclusion of the form $[\ldots w_0 \ldots t_0 \ldots]$, producing a truth-value/o, is derived. Since we are to derive a hyperproposition, we finally abstract over the values of the variables w_0, t_0, thus reintroducing the leftmost $\lambda w \lambda t$ to produce a proposition: $\lambda w \lambda t\ [\ldots w \ldots t \ldots] \rightarrow o_{\tau\omega}$. In order to simplify the derivations occurring in proofs and rules, in what follows we omit the initial and final steps of λ-elimination and λ-introduction, respectively.

The proof calculus we usually apply is Gentzen's system of natural deduction adjusted to TIL.[43] We follow Church and Genzten in the classical style of a proof calculus as it is applied in, for instance, HOL languages.[44] The standard rules of a proof calculus are, in TIL, applicable to the *constituents* of procedures that are typed to produce a truth-value. The rules follow the general pattern of I/E pairs. The rules handling the truth-functions are standard, as in classical propositional logic. Since the quantifiers of TIL are functions applicable to classes of objects, the rules for quantifiers are introduced in the way similar to standard λ-calculi. For instance, the rule for universal quantifier eliminiation (∀E) in TIL comes in this form. Let $x \rightarrow \alpha$, $B(x) \rightarrow o$: the variable x is free in B; $[\lambda x\, B] \rightarrow (o\alpha)$, $\forall/(o(o\alpha))$, $C \rightarrow \alpha$: a procedure that is *not v-improper*. Then:

[42] Restricted β-reduction consists merely in the substitution of variables for variables.

[43] For details, see Duží and Menšík (2020), Duží and Fait (2021) or Duží (2012).

[44] The reason is that TIL is a typed λ-calculus with a Church-style semantics, in which every λ-term comes with a unique type and types are organised into disjoint layers. On the other hand, Curry systems are essentially treated as untyped λ-calculi, in which a term is associated with a set (which may be empty) of potential types. See Gordon and Melham (1993) on HOL.

$$\begin{array}{ll} [{}^0\forall \lambda x\ B] & \emptyset \\ [[\lambda x\ B]\ C] & \forall\mathrm{E} \\ B(C/x) & \beta\text{-reduction} \end{array}$$

where $B(C/x)$ arises from B by a valid (hence, correct or collision-free) substitution of the procedure C for all occurrences of the variable x in B. That C is not v-improper is a crucial condition for the applicability of this rule. Otherwise, the rule would not be truth-preserving.

For the sake of simplicity, we usually write this rule in the ordinary abbreviated form:

$$\frac{X \vdash [{}^0\forall \lambda x\, B]}{X \vdash B(C/x)}\ (\forall E)$$

The dual rule (∀I) receives this form (y being a fresh, free variable, i.e., one local to this part of the derivation):

$$\frac{X \vdash B(y/x)}{X \vdash [{}^0\forall \lambda x\ B]}\ (\forall \mathrm{I})$$

In this paper, we will, however, need the rule of ∃-introduction, (∃I). In classical extensional logics and λ-calculi of *total* functions, the rule is unproblematic. Yet, since TIL is a hyperintensional logic of *partial* functions, we must be careful not to derive that there is a value of a function at an argument when there is none.

(∃I) is valid in its classical form, provided it is applied to a *constituent* of an assumption B. Recall that a *constituent* of B is a procedure that in B occurs in executed mode. Hence, let $D \rightarrow \alpha$ be a procedure that occurs as a constituent of the procedure B, the other types as above. Since, by assumption, B produces the truth-value **T** and D is a constituent of B, procedure B is of the form of a Composition: [… D …]. Then, as per the definition of Composition, procedure D cannot be v-improper, and so the Composition $[[\lambda x\ B]\ D]$ v-produces **T** as well. Thus, the set of α-elements produced by $[\lambda x\ B]$ is non-empty, and the application of the quantifier ∃ is truth-preserving. As a result, we arrive at the classical rule:

$$\frac{X \vdash B(D/x)}{X \vdash [{}^0\exists \lambda x\ B]}\ (\exists \mathrm{I})$$

The crucial condition for the validity of (∃I) is that D must occur as a constituent of B. Hence, this rule quantifies *over constituents*; it does not quantify *into* a hyperintensional context. If it did, we might risk deriving the existence of a non-existent object, thus crossing the line between logic and magic.

Still, any hyperintensional logic worthy of the name is obliged to also explain how to quantify *into* a hyperintensional context. In TIL, this task assumes the form of explaining how to logically operate on a displayed procedure. To this end, we make use of the two functions *Sub* and *Tr* defined above (Definitions 11, 12).

Consider the sentence "There is an object such that Tilman believes (hyperintensionally) that it is a planet". To analyse the sentence, we encounter the problem of the need to λ-bind a ^{0}-bound variable, which does not work as we would like it to ($Believe^*/(o\iota *_n)_{\tau\omega}$; $Planet(o\iota)_{\tau\omega}$; $y \rightarrow \iota$):

$$\lambda w\lambda t\ [^0\exists\lambda y\ [^0Believe^*_{wt}\ ^0Tilman\ ^0[\lambda w\lambda t\ [^0Planet_{wt}y]]]]$$

The problem is this. The Closure $\lambda y\ [^0Believe^*_{wt}\ ^0Tilman\ ^0[\lambda w\lambda t\ [^0Planet_{wt}\ y]]]$ produces either the whole universe of discourse or an empty class of individuals, according as Tilman believes that the incomplete procedure $[\lambda w\lambda t\ [^0Planet_{wt}\ y]]$ produces a truth-condition satisfied in the world and at the time of evaluation, and independently of the objects assigned to the variable y by valuation v. To obtain a plausible analysis, we must extract the variable y out of the hyperintensional context to make it free for λ-binding. Here is how:

$$\lambda w\lambda t\ [^0\exists\lambda y\ [^0Believe^*_{wt}\ ^0Tilman\ [^0Sub\ [^0Tr\ y]\ ^0x\ ^0[\lambda w\lambda t\ [^0Planet_{wt}x]]]]]$$

Now everything runs like clockwork. The Composition $[^0Sub\ [^0Tr\ y]\ ^0x\ ^0[\lambda w\lambda t\ [^0Planet_{wt}\ x]]]$ v-produces a hyperproposition that some object y is a planet. Assume, e.g., that Tilman believes that Venus is a planet. Then this Composition *v(Venus/y)*-produces the hyperproposition $[\lambda w\lambda t\ [^0Planet_{wt}\ ^0Venus]]$, and the class of objects produced by the Closure $\lambda y\ [^0Believe^*_{wt}\ ^0Tilman\ [^0Sub\ [^0Tr\ y]\ ^0x\ ^0[\lambda w\lambda t\ [^0Planet_{wt}\ x]]]]$ is non-empty, as it contains at least the object Venus.

Still, there is more to the story of quantifying into hyperintensional contexts. As stated above, in hyperintensional attitudinal sentences we must fully respect the attributee's perspective when one complement is substituted for another. For this reason, we can substitute only procedurally isomorphic complements. To illustrate the problem, consider this argument:

Tilman believes that the ninth celestial body
moving in an elliptical orbit around the Earth is a planet
―――――――――――――――――――――――――――――――
There is an object such that Tilman believes that it is a planet

Abbreviate 'the ninth celestial body moving in an elliptical orbit around' as '9-body-moving-around', and analyse the term 'the ninth celestial body moving in an elliptical orbit around the Earth' simply as $\lambda w\lambda t\ [^09\text{-}body\text{-}move\text{-}around_{wt}\ ^0Earth] \rightarrow \iota_{\tau\omega}$: the individual role of the ninth celestial body moving around the Earth. *Types*: $Earth/\iota$; $9\text{-}body\text{-}move\text{-}around/(\iota\iota)_{\tau\omega}$.

Then the analysis applying the same technique as above results in an *invalid* argument:

$$\frac{\lambda w\lambda t\ [^0Believe^*_{wt}\ ^0Tilman\ ^0[\lambda w\lambda t\ [^0Planet_{wt}\ [^09\text{-}body\text{-}move\text{-}around_{wt}\ ^0Earth]]]]}{\lambda w\lambda t\ [^0\exists\lambda y\ [^0Believe^*_{wt}\ ^0Tilman\ [^0Sub\ [^0Tr\ y]\ ^0x\ ^0[\lambda w\lambda t\ [^0Planet_{wt}\ x]]]]]}$$

As the ninth celestial body moving in an elliptical orbit around the Earth is Pluto (in the actual world at the present time), consider the valuation *v(Pluto/y)*. The Composition

$$[{}^0Sub\,[{}^0Tr\ y]\,{}^0x\,{}^0[\lambda w\lambda t\,[{}^0Planet_{wt}\,x]]]$$

then *v(Pluto/y)*-produces $[\lambda w\lambda t\ [{}^0Planet_{wt}\ {}^0Pluto]]$, which, however, is not the procedure to which Tilman is related in the premise, nor is this procedure procedurally isomorphic to that procedure:

$$[\lambda w\lambda t\,[{}^0Planet_{wt}\,[{}^09\text{-}body\text{-}move\text{-}around_{wt}{}^0Earth]]]$$

4 Quantifying-in

In this part, we describe, prove and apply two rules for quantifying into hyperpropositional attitudes de dicto. The two rules validate quantifying *into* hyperpropositions. One rule quantifies over a constituent procedure; the other rule quantifies over an object produced by a constituent Trivialisation.

4.1 Hyperpropositional attitude contexts de dicto

We study first how to analyse two sample sentences, one expressing an attitude to an empirical and the other to a mathematical hyperproposition:

(E) "Tilman believes that Pluto is a planet"
(M) "Tilman believes that the cotangent of π equals zero"

The first sentence expresses that a doxastic relation-in-intension obtains between Tilman and the hyperproposition that Pluto is a planet. The analysandum does not necessitate hyperpropositional treatment per se, for a standard intensional analysis would suffice, provided we are analysing implicit (i.e., logically closed) beliefs that the agent need not be aware of having and which the agent is not going to manipulate logically, e.g., by drawing inferences. As soon as Tilman is related to a hyperproposition, his attitude is an explicit one. Matters get a good deal more complicated when Tilman believes that Vulcan is not (defined as) a planet, despite Vulcan having been defined to be a planet, and so has an attitude toward Vulcan that flies in the face of its definition. This sort of attitude demands that the complement be a hyperproposition on pain of relating the agent to a blatant contradiction.

Sentence (M) does necessitate hyperintensional treatment, as does any other mathematical attitude. Neither the necessary proposition (the one true at all worlds), nor the impossible proposition (the one true at no worlds) is a suitable complement, because there could be but two mathematical attitudes then. More specifically to the example embedded in (M), since the function *Cotg* is not defined at the argument π, no number is produced, and so there is nothing to be compared to the number zero. The procedure encoded by "*Cotg* (π) = 0" is improper in the sense of producing no truth-value.[45] What Tilman believes (wrongly) is that the very *procedure* encoded by the term "*Cotg* (π) = 0" produces **T**.

The analysis of the two sentences issues in these two Closures:

(E*) $\lambda w \lambda t\ [{}^0Believe_{wt}\ {}^0Tilman\ {}^0[\lambda w' \lambda t'\ [{}^0Planet_{w't'}\ {}^0Pluto]]]$
(M*) $\lambda w \lambda t\ [{}^0Believe_{wt}\ {}^0Tilman\ {}^0[{}^0{=}\ [{}^0Cotg\ {}^0\pi]\ {}^00]]$

Types: $Believe/(o\iota *_1)_{\tau\omega}$; $Tilman/\iota$; $Planet/(o\iota)_{\tau\omega}$; $Pluto/\iota$; $w, w' \rightarrow \omega$; $t, t' \rightarrow \tau$; $[{}^0Planet_{w't'}\ {}^0Pluto] \rightarrow o$; $[\lambda w' \lambda t'\ [{}^0Planet_{w't'}\ {}^0Pluto]] \rightarrow o_{\tau\omega}$; ${}^0[\lambda w' \lambda t'\ [{}^0Planet_{w't'}\ {}^0Pluto]] \rightarrow *_1$; $=/(o\tau\tau)$; $Cotg/(\tau\tau)$; $\pi, 0/\tau$; $[{}^0Cotg\ {}^0\pi] \rightarrow \tau$; $[{}^0{=}\ [{}^0Cotg\ {}^0\pi]\ {}^00] \rightarrow o$; the other types are obvious.

The single most important bit is the Trivialisation of the Closure $[\lambda w' \lambda t'\ [{}^0Planet_{w't'}\ {}^0Pluto]]$ and of the Composition $[{}^0{=}\ [{}^0Cotg\ {}^0\pi]\ {}^00]$, thereby generating a *hyperintensional* context by displaying the Closure and the Composition, respectively. In both analyses the context becomes, furthermore, an *attitude* context thanks to the binary relation *Believe*, which we stipulate to be a relation between doxastic agents and the hyperpropositions they believe to be true.[46] Finally, the attitudes are *de dicto* in virtue of the form of the attitude complement. In the first case, the agent is related to the hyperproposition that Pluto is a planet, and in the second case to the hyperproposition that *Cotg* takes π to 0. (E*) and (M*) are instances of the schema of the logical forms that characterise hyperpropositional attitudes de dicto:

$$\lambda w \lambda t\ [\chi_{wt}\ a\ {}^0[\lambda w' \lambda t' [\varphi_{w't'} b]]]$$
$$\lambda w \lambda t\ [\chi_{wt}\ a\ {}^0[F\ \ c]]$$

Types: $\chi \rightarrow (o\iota *_1)_{\tau\omega}$; $a, b \rightarrow \iota$; $\varphi \rightarrow (o\iota)_{\tau\omega}$; $F \rightarrow (o\tau)$; $c \rightarrow \tau$.

By contrast, hyperpropositional attitudes de re are characterised by either of these two schemas:

$$\lambda w \lambda t\ [\chi_{wt}\ a\ [{}^0Sub\ [{}^0Tr\ \ b]\ {}^0it\ {}^0[\lambda w' \lambda t' [\varphi_{w't'}\ it]]]]$$
$$\lambda w \lambda t\ [\chi_{wt}\ a\ [{}^0Sub\ [{}^0Tr\ \ c]\ {}^0it'\ {}^0[F\ \ it']]]$$

The first schema should be read as "*a* χ's of *b* that it is a φ". In the mathematical case, the second schema should be read as "*a* χ's of *c* that it is an *F*."

Note that the occurrences of the anaphoric pronoun 'it' are analysed as the ^{0}-bound variables $it \rightarrow \iota$, $it' \rightarrow \tau$. The anaphoric references 'of *b* that it is a φ' and 'of

[45] Actually, the sentence "*Cotg* of π equals 0" comes with the *presupposition* that the value of the function *Cotg* exists at π, because this is entailed both by the sentence and its narrow-scope negation, "*Cotg* of π is *not* equal to zero". See Duží (2017, 2018a, 2018b) for more on presupposition and negation.

[46] We are glossing over a slight complication here. Our hyperpropositions are not truth-bearers, so they cannot, strictly and literally, be believed to be *true* (or *false*). Rather it is truth-conditions, i.e., the propositions of possible-world semantics, that are truth-bearers: a truth-condition is true when it is satisfied. So to believe an empirical hyperproposition is to believe that the truth-condition it produces is satisfied at the given world and time of evaluation; and to believe a mathematical hyperproposition amounts to believing that the procedure produces the truth-value **T**. This technical detail is solved in Duží et al. (2010, §5.1.6).

c that it is an F' are then resolved by the Compositions $[{}^0Sub\ [{}^0Tr\ b]\ {}^0it\ {}^0[\lambda w'\lambda t'\ [\varphi_{w't'}\ it]]]$ and $[{}^0Sub\ [{}^0Tr\ c]\ {}^0it'\ {}^0[F\ it']]$, respectively.[47]

The question now arises what can be validly deduced from (E) and from (M). Since in both cases Tilman is related to a procedure, we can obviously validly infer that there is a procedure (a *hyperproposition*, in this case) to which Tilman is related by a doxastic attitude:

$$\text{(EA)}\quad \frac{\lambda w\lambda t\ [{}^0Believe_{wt}\ {}^0Tilman\ {}^0[\lambda w'\lambda t'\ [{}^0Planet_{w't'}\ {}^0Pluto]]]}{\lambda w\lambda t\ [{}^0\exists\ \lambda c\ [{}^0Believe_{wt}\ {}^0Tilman\ c]]}$$

$$\text{(MA)}\quad \frac{\lambda w\lambda t\ [{}^0Believe_{wt}\ {}^0Tilman\ {}^0[{}^0{=}\ [{}^0Cotg\ {}^0\pi]\ {}^00]]}{\lambda w\lambda t\ [{}^0\exists\ \lambda d\ [{}^0Believe_{wt}\ {}^0Tilman\ d]]}$$

Types: $c/*_2 \rightarrow *_1$; ${}^2c \rightarrow o_{\tau\omega}$; $d/*_2 \rightarrow *_1$; ${}^2d \rightarrow o$; the other types as above. The inclusion of ${}^2c$, ${}^2d$ spells out the fact that c, d v-produce procedures typed to produce empirical truth-conditions and truth-values, respectively.

In both cases the complements of the attitude, i.e. the procedures ${}^0[\lambda w'\lambda t'\ [{}^0Planet_{w't'}\ {}^0Pluto]]$ and ${}^0[{}^0{=}\ [{}^0Cotg\ {}^0\pi]\ {}^00]$ are *constituents* of the premise; hence, we simply applied the (∃I)-rule proved above. This is a simple matter, because we just quantified *over* a believed hyperproposition, i.e., over an entire hyperintensional context. Still, our main goal and the main novelty of this paper is quantifying *into* hyperintensional contexts.

First, we are going to tackle a simpler case, which is to quantify over a *procedure* that is a constituent of an attitude complement. Then we are going to show that in some special, rigorously defined cases we can also quantify over a *product* of such a procedure. For instance, (M) entails that there is a *procedure* such that Tilman believes that its product equals zero. Indeed, there is such a procedure, namely the improper Composition $[{}^0Cotg\ {}^0\pi]$, such that Tilman believes (wrongly) that this procedure produces 0.

As explained above, quantifying into displayed contexts is not a straightforward thing to do. Carelessly quantifying-in is *not* truth-preserving ($c \rightarrow *_1$; ${}^2c \rightarrow \tau$, i.e., c v-produces a procedure that is typed to produce a number):

$$\frac{\lambda w\lambda t\ [{}^0Believe_{wt}\,{}^0Tilman\ {}^0[\,{}^0{=}\ [{}^0Cotg\,{}^0\pi]\,{}^00]]}{\lambda w\lambda t\ [{}^0\exists\lambda c\ [{}^0Believe_{wt}\,{}^0Tilman\ {}^0\ [\,{}^0{=}\ {}^2c\ {}^00]]]}$$

Its invalidity is due to the Trivialisation of the attitude complement, which displays the believed hyperproposition $[{}^0{=}\ [{}^0Cotg\ {}^0\pi]\ {}^00]$, but does not execute it. As defined above (see Defs. 6, 7), all the sub-procedures of a displayed procedure are displayed as well. The Trivialisation ${}^0[{}^0{=}\ {}^2c\ {}^00]$ produces the Composition $[{}^0{=}\ {}^2c\ {}^00]$ regardless of any valuation of the variable c. In particular, a displayed occurrence of a variable does not descend to the value of the variable. For this reason, the truth of

[47] For details on how TIL analyses anaphoric resolution, see Duží (2018a).

the premise does not warrant the non-emptiness of the class of procedures *v*-produced by the Closure

$$\lambda c\ [{}^0Believe_{wt}\ {}^0Tilman\ {}^0[{}^0{=}\ {}^2c\ {}^00]]$$

This class is either the whole type $*_1$ or the empty set of first-order procedures, depending on whether or not Tilman is related to the Composition $[{}^0{=}\ {}^2c\ {}^00]$ and independently of the truth of the premise, i.e. independently of whether or not Tilman is related to the Composition in the premise, $[{}^0{=}\ [{}^0Cotg\ {}^0\pi]\ {}^00]$.

Fortunately, we have a way out, or rather a way *in*. It is our substitution method that makes it possible to operate on displayed procedures. This inference is valid:

$$(\mathrm{MA}_0)\ \ \frac{\lambda w \lambda t\ [{}^0Believe_{wt}\ {}^0Tilman\ {}^0[{}^0{=}\ [{}^0Cotg\ {}^0\pi]\ {}^00]]}{\lambda w \lambda t\ [{}^0\exists^*\ \lambda c\ [{}^0Believe_{wt}\ {}^0Tilman\ [{}^0Sub\ c\ {}^0d\ {}^0[{}^0{=}\ d\ {}^00]]]]}$$

Additional types: $\exists^*/(o(o*_n))$; $c/*_2 \rightarrow *_1$; ${}^2c{\rightarrow}\tau$; $d/*_1{\rightarrow}\tau$. The inclusion of ${}^2c$ makes it explicit that *c* *v*-produces a procedure that is typed to produce numbers.

The variable *c* occurs free in the Composition $[{}^0Sub\ c\ {}^0d\ {}^0[{}^0{=}\ d\ {}^00]]$, and the Composition $v([{}^0Cotg\ {}^0\pi]/c)$-produces exactly what it should produce, namely the procedure $[{}^0{=}\ [{}^0Cotg\ {}^0\pi]\ {}^00]$, which Tilman is related to as per the premise above. Recall that $v([{}^0Cotg\ {}^0\pi]/c)$ is a valuation just like *v*, except for assigning the Composition $[{}^0Cotg\ {}^0\pi]$ to *c*.

One might wonder whether in the case of (E) we could actually infer more than that there is a *procedure* believed by Tilman to produce a condition that is satisfied, as stipulated by (EA). No Closure is *v*-improper for any *v*.[48] A Closure of the form $[\lambda x_1 \ldots x_m\ X]$ *v*-produces a function for any *v*, even when the produced function is degenerate, i.e., one that is undefined at each of its arguments, namely when *X* is a *v*-improper procedure for every *v*. Thus, it might seem that we could validly infer from (E) not only that there is a hyperproposition but also that there is a *truth-condition* $r/o_{\tau\omega}$, such that Tilman believes that *r* is true in the world *w* and at the time *t* of evaluation. Yet this *cannot* be inferred. According to the premise, Tilman is related to the hyperproposition

$${}^0[\lambda w' \lambda t' [{}^0Planet_{w't'}\ {}^0Pluto]]$$

rather than the truth-condition produced by this very hyperproposition. Tilman might, of course, be in some relation to the state-of-affairs that Pluto is a planet, and if we inferred that he *is* in a coarse-grained belief relation to this state-of-affairs, i.e. *r*, we would face the same problem as many attitude logics do, namely one or more variants of the problem of omniscience. Furthermore, no less importantly, we would not be respecting Tilman's doxastic perspective, and so the ascription would not be of a de dicto attitude.

Yet we would still want to draw further conclusions. We want to infer that in some special cases there is some *product* of the procedure which is the constituent

48 We are dealing here with the standard Closure as defined by Definition 1. We do not take into account λ^{α}-Closure, which can be *v*-improper under specific conditions. For details, see Duží and Kosterec (2017).

of attitude complement. In the empirical case, we would want to derive, for instance, that there is an individual x such that Tilman believes that x is a planet. And indeed, there is such an individual x, namely Pluto. Or, additionally, that there is a property p such that Tilman believes that Pluto has property p. And indeed, there is such a property, namely the property of being a planet. Similarly, in the mathematical case, we would like to infer, for instance, that there is a number y such that Tilman believes that the value of *Cotg* equals zero at y. And indeed, there is such a number, namely π. Furthermore, we can validly infer that there is a function f such that Tilman believes that the value of f at π is zero.

But again, this inference is *invalid* ($x \rightarrow \tau$):

$$\frac{\lambda w \lambda t\, [^0 Believe_{wt}\, ^0 Tilman\, ^0[^0{=}\ [^0 Cotg\, ^0\pi]\, ^0 0]]}{\lambda w \lambda t\, [^0\exists \lambda x\, [^0 Believe_{wt}\, ^0 Tilman\, ^0[^0{=}\ [^0 Cotg\, ^0 x]]\, ^0 0]]}$$

because the variable x is ^{0}-bound (see Definition 9). To obtain a *valid* inference, we must apply the substitution method:

$$(\mathrm{MA}_1)\quad \frac{\lambda w \lambda t\, [^0 Believe_{wt}\, ^0 Tilman\, ^0[^0{=}\ [^0 Cotg\, ^0\pi]\, ^0 0]]}{\lambda w \lambda t\, [^0\exists\ \lambda x\, [^0 Believe_{wt}\, ^0 Tilman\, [^0 Sub\, [^0 Tr^{\tau}\ x]\, ^0 y\, ^0[^0{=}\ [^0 Cotg\ y]\, ^0 0]]]]}$$

Gloss: "There is a number x such that Tilman believes that the value of *Cotg* at x is 0." *Additional types*: $\exists/(o(o\tau))$, x, $y \rightarrow \tau$.

We are substituting the Trivialisation of a number being quantified over, using the functions *Sub*, *Tr* (see Defs. 11, 12). Similarly, we can derive that there is a function $f \rightarrow (\tau\tau)$ such that Tilman believes that the value of f at π is zero:

$$(\mathrm{MA}_2)\quad \frac{\lambda w \lambda t\, [^0 Believe_{wt}\, ^0 Tilman\, ^0[^0{=}\ [^0 Cotg\, ^0\pi]\, ^0 0]]}{\lambda w \lambda t\, [^0\exists'\ \lambda f\, [^0 Believe_{wt}\, ^0 Tilman\, [^0 Sub\, [^0 Tr^{(\tau\tau)}\ f]\, ^0 g\, ^0[^0{=}\ [g\ \pi]\, ^0 0]]]]}$$

Additional types: $\exists'/(o(o(\tau\tau)))$, $g \rightarrow (\tau\tau)$.

Now everything is as it should be. The variables x, f occur *free* in

$$[^0 Sub\, [^0 Tr^{\tau} x]\, ^0 y\, ^0[^0{=}\ [^0 Cotg\ y]\, ^0 0]]$$
$$[^0 Sub\, [^0 Tr^{(\tau\tau)} f]\, ^0 g\, ^0[^0{=}\ [g\ \pi]\, ^0 0]]$$

respectively, and the first Composition $v(\pi/x)$-produces $[^0{=}\ [^0 Cotg\ ^0\pi]\ ^0 0]$, while the second Composition $v(Cotg/f)$-produces the same procedure, namely the one believed by Tilman.

Hence, concerning (MA_1), provided the Composition

$$[^0 Believe_{wt}\, ^0 Tilman\ ^0[^0{=}\ [^0 Cotg\, ^0\pi]\, ^0 0]]$$

v-produces the truth-value **T**, the class of numbers v-produced by

$$\lambda x [^0 Believe_{wt}\, ^0 Tilman\, [^0 Sub\, [^0 Tr^{\tau}\ x]\, ^0 y\, ^0[^0{=}\ [^0 Cotg\ y]\, ^0 0]]]$$

is non-empty (as it contains at least the element π) and the application of $\exists/(o(o\tau))$ is truth-preserving.

Similarly, concerning (MA_2), provided the Composition

$$[^0 Believe_{wt}\, ^0 Tilman\, ^0[^0{=}\ [^0 Cotg\, ^0\pi]\, ^0 0]]$$

v-produces the truth-value **T**, the class of functions of type (ττ) *v*-produced by

$$\lambda f\, [^0 Believe_{wt}\, ^0 Tilman\, [^0 Sub\, [^0 Tr^{\tau\tau} f]\, ^0 g\, ^0[^0{=}\ [g\, \pi]\, ^0 0]]]$$

is non-empty (as it contains at least the function *Cotg*), and the application of $\exists'/(o(o(\tau\tau)))$ is truth-preserving.

Still, from (M) we *cannot* validly infer that there is a *number n* such that Tilman believes that this number equals 0. The reason is that the function *Cotg* is not defined at the argument π; hence the Composition $[^0 Cotg\, ^0\pi]$ is improper by failing to produce anything. Deriving that there is such a number *n* would, again, cross the line from logic into magic.

Turning now to the empirical case, valid inferences that involve quantifying *into* a hyperpropositional context are, for instance, these ($x, y \rightarrow \iota$; $p, q \rightarrow (o\iota)_{\tau\omega}$):

$$(EA_1)\quad \frac{\lambda w \lambda t\ [^0 Believe_{wt}\ ^0 Tilman\ ^0[\lambda w' \lambda t'\ [^0 Planet_{w't'}\ ^0 Pluto]]]}{\lambda w \lambda t\ [^0\exists\ \lambda x\ [^0 Believe_{wt}\ ^0 Tilman\ [^0 Sub\ [^0 Tr^{\iota}\ x]\ ^0 y\ ^0[\lambda w' \lambda t'\ [^0 Planet_{w't'}\ y]]]]]}$$

Gloss: "There is an individual *x* such that Tilman believes that *x* is a planet."

$$(EA_2)\quad \frac{\lambda w \lambda t\ [^0 Believe_{wt}\ ^0 Tilman\ ^0[\lambda w' \lambda t'\ [^0 Planet_{w't'}\ ^0 Pluto]]]}{\lambda w \lambda t\ [^0\exists\ \lambda p\ [^0 Believe_{wt}\ ^0 Tilman\ [^0 Sub\ [^0 Tr^{(((o\iota)\tau)\omega)}\ p]\ ^0 q\ ^0[\lambda w' \lambda t'\ [q_{w't'}\ ^0 Pluto]]]]]}$$

Gloss: "There is a property *p* such that Tilman believes that Pluto has *p*."

The arguments (MA_1), (MA_2), (EA_1) and (EA_2) are *valid*, because we are quantifying over objects produced by Trivialisation, namely $^0\pi$, $^0 Cotg$, $^0 Pluto$, $^0 Planet$, and these procedures are not *v*-improper for any valuation *v*. Trivialisation just displays the object which we go on to quantify over, and when applied to this object (*v*-produced by a variable) the function *Tr* returns as its value the Trivialisation of the object. Moreover, we are fully respecting Tilman's perspective here, because our analyses are *literal* ones. This means that semantically simple terms like 'planet', 'Pluto', 'cotangent' and 'π' are analysed by their Trivialisations. Indeed, the sentences do not convey any more information about the meaning of these terms, as a definition or meaning postulate would.[49]

On the other hand, if a constituent of the attitude complement can be *v*-improper, then we cannot validly infer the existence of the respective object. Assume that instead of a Trivialisation displaying Pluto we were to conceptualise Pluto by means

[49] For *literal analysis*, see fn. 39.

of the individual role denoted by the definite description ‘the first Kuiper Belt object to be discovered’. Abbreviate this description as ‘*FKB*O’. The analysis of ‘*FKBO*’ amounts to this procedure:

$$\lambda w \lambda t\ [{}^0First\ \lambda x\ [[{}^0KuiperBeltObj_{wt}\ x] \wedge [{}^0Discovered_{wt}\ x]]] \rightarrow \iota_{\tau\omega}$$

Types: *First*/(ι(οι)): the function that picks out at most one individual from a set of individuals (namely the first one to be discovered); *KuiperBeltObj*, *Discovered*/ $(o\iota)_{\tau\omega}$, $x \rightarrow \iota$; $[{}^0KuiperBeltObj_{wt}\ x]$, $[{}^0Discovered_{wt}\ x] \rightarrow o$;
$\lambda x\ [[{}^0KuiperBeltObj_{wt}\ x] \wedge [{}^0Discovered_{wt}\ x]] \rightarrow (o\iota)$;
$[{}^0First\ \lambda x\ [[{}^0KuiperBeltObj_{wt}\ x] \wedge [{}^0Discovered_{wt}\ x]]] \rightarrow \iota$.

Since the office can possibly go vacant, the following argument similar to (EA_1) is *invalid*:

$$\frac{\begin{array}{c}\lambda w \lambda t\ [{}^0Believe_{wt}\ {}^0Tilman\ {}^0[\lambda w' \lambda t'\ [{}^0Planet_{w't'} \\ \lambda w \lambda t\ [{}^0First\ \lambda x\ [[{}^0KuiperBeltObj_{wt}\ x] \wedge [{}^0Discovered_{wt}\ x]]]_{w't'}]]]\end{array}}{\lambda w \lambda t\ [{}^0\exists \lambda x\ [{}^0Believe_{wt}\ {}^0Tilman\ [{}^0Sub\ [{}^0Tr^{\iota}\ x]\ {}^0y\ {}^0[\lambda w' \lambda t'\ [{}^0Planet_{w't'}\ y]]]]]]}$$

The conclusion that there is an *individual* *x* such that Tilman believes that *x* is a planet is *not entailed* by the premise, because the Composition (note the rightmost subscripted *w’t’*)

$$[\lambda w \lambda t\ [{}^0First\ \lambda x\ [[{}^0KuiperBeltObj_{wt}\ x] \wedge [{}^0Discovered_{wt}\ x]]]_{w't'}]$$

may be *v*-improper. Thus, though $[{}^0Sub\ [{}^0Tr^{\iota}\ x]\ {}^0y\ {}^0[\lambda w' \lambda t'\ [{}^0Planet_{w't'}\ y]]]$ $v(a/x)$-produces the procedure $[\lambda w' \lambda t'\ [{}^0Planet_{w't'}\ {}^0a]]$ for some individual *a*, it is not excluded that *a* fails to be an element of the class of individuals which Tilman believes to be a planet. In other words, the class produced by the Closure

$$\lambda x\ [{}^0Believe_{wt}\ {}^0Tilman\ [{}^0Sub_1\ [{}^0Tr^{\iota}\ x]\ {}^0y\ {}^0[\lambda w' \lambda t'\ [{}^0Planet_{w't'}\ y]]]]$$

can be empty, and when it is, applying ∃ to this class will yield **F**.

Since the Closure $\lambda w \lambda t\ [{}^0First\ \lambda x\ [[{}^0KuiperBeltObj_{wt}\ x] \wedge [{}^0Discovered_{wt}\ x]]]$ cannot be *v*-improper for any *v*, it might seem that we could validly infer that there is an *individual office* $f \rightarrow \iota_{\tau\omega}$ such that Tilman believes that the occupant of the office is a planet (though the office may be vacant or, if occupied, its occupant may fail to be a planet). Yet, again, an argument to this effect would be *invalid*:

$$\frac{\begin{array}{c}\lambda w \lambda t\ [{}^0Believe_{wt}\ {}^0Tilman\ {}^0[\lambda w' \lambda t'\ [{}^0Planet_{w't'} \\ \lambda w \lambda t\ [{}^0First\ \lambda x\ [[{}^0KuiperBeltObj_{wt}\ x] \wedge [{}^0Discovered_{wt}\ x]]]_{w't'}]]]\end{array}}{\lambda w \lambda t\ [{}^0\exists \lambda x\ [{}^0Believe_{wt}\ {}^0Tilman\ [{}^0Sub\ [{}^0Tr^{((\tau\tau)\omega)} f]\ {}^0y\ {}^0[\lambda w' \lambda t'\ [{}^0Planet_{w't'}\ g_{w't'}]]]]]]}$$

Gloss: “There is an individual office *f* such that Tilman believes the hyperproposition that the occupant of *f* is a planet.” *Additional types*: $f, g \rightarrow \iota_{\tau\omega}$.

The argument is invalid, because the office in the premise is conceptualised by means of the Closure $\lambda w \lambda t\ [{}^0First\ \lambda x\ [[{}^0KuiperBeltObj_{wt}\ x] \wedge [{}^0Discovered_{wt}\ x]]]$ rather than by its Trivialisation.

4.1.1 Procedural isomorphism

The main reason these last two arguments are not valid is this. When deriving something from a *hyperintensional* attitude *de dicto*, we must strictly respect the agent's perspective. Hence, there must be a valuation v such that the procedure resulting from the substitution is exactly the same procedure as the one to which the agent was originally related. More precisely, the so v-produced procedure must be *procedurally isomorphic* to the procedure to which Tilman is related as per the premise.[50] In other words, the derived procedure must be encoded by a sentence *synonymous* with Tilman's original attitude complement. In Duží (2019) a series of *criteria* for procedural isomorphism—hence co-hyperintensionality, hence synonymy—has been defined. These criteria are partially ordered from the strongest (most restrictive) to the weakest (most liberal) with respect to synonymy. Here we opt for the almost-strongest criterion C_1. Let us pause to reflect on why we are going for C_1 rather than C_7, which we have elsewhere labelled (A1″) to incorporate it into Church's hierarchy of Alternatives.[51]

$\boldsymbol{C_1}$. α-conversion
$\boldsymbol{C_7}$. α-conversion + β-reduction by value

C_7 is applicable to hyperintensional contexts, but only if a slightly contentious assumption is granted. The assumption is that $[[\lambda x_1 \ldots x_n\ Y]\ D_1 \ldots D_n]$ is the same procedure as

$$^2[^0Sub\ [^0Tr\ D_1]\ ^0x_1 \ldots\ [^0Sub\ [^0Tr\ D_n]\ ^0x_n\ ^0Y]]$$

Is the assumption reasonable? That depends. Our hypothesis is that C_7 is applicable to natural-language discourse, including attributions of hyperpropositional attitudes de dicto. But the problem is that C_7 does not extend to attributions of mathematical and logical attitudes de dicto. To see this, consider this example:

$$\text{Tilman believes that } [^0{=}\ [\lambda x\ [^0{+}\ [log_2\ ^0 16]\ x]\ [^0Cos\ ^0 0]]\ ^0 5]$$

Then by C_7 this should follow:

$$\text{Tilman believes that } [^0{=}\ ^2[^0Sub\ [^0Tr\ [^0Cos\ ^0 0]]\ ^0x\ ^0[^0{+}\ [log_2\ ^0 16]\ x]]\ ^0 5]$$

But what if Tilman's idiosyncratic perspective is quite another so that he would compute the equation in another way than predicted? The second equation specifies that Tilman first computes $Cos(0)$ to obtain the number 1 and afterwards substitutes 1 for x in $log_2(16){+}x$, which gives $log_2(16){+}1$. Then he computes $log_2(16)$ to obtain 4, and finally the number 5. Yet, a different computation would be to first compute $log_2(16)$ to obtain 4, and then substitute the result of computing $Cos(0)$, hence number 1, for x into $4{+}x$, and finally obtain the result 5. If we instead restrict

[50] See Jespersen (2021) for arguments in favour of procedural isomorphism, together with a concrete application.

[51] See Anderson (1998).

ourselves to C_1 then if the original attitude is the same as above, what follows is only this:

Tilman believes that $[^0{=}\ [\lambda y\ [^0{+}\ [log_2\ ^0 16]\ y]\ [^0 Cos\ ^0 0]]\ ^0 5]$

as only α-conversion is admitted.[52] Here is the formal definition:

Definition 13 (*α-conversion*). Let a procedure Y contain at most $x_1, \ldots, x_n$ as free variables. Then:

$$[\lambda x_1 \ldots x_n Y] \Rightarrow_\alpha [\lambda y_1 \ldots y_n Y(y_1/x_1 \ldots y_n/x_n)]$$

where $Y(y_1/x_1 \ldots y_n/x_n)$ is the procedure that arises from Y by correct substitution of y_1 for all the occurrences of $x_1, \ldots$, and y_n for all the occurrences of x_n, is *α-conversion*. □

Thus, our bifurcation of attitude complements into those hyperpropositions that produce empirical truth-conditions and those hyperpropositions that produce truth-values is accompanied by a bifurcation of what procedural isomorphism amounts to. Our current stance is that C_1 is suitable for the former and C_7 for the latter when the relevant context is a hyperpropositional attitude de dicto.

Back to the Kuiper Belt. If *FKBO* is the office produced by the Closure

$$\lambda w \lambda t\ [^0 First\ \lambda x\ [[^0 KuiperBeltObj_{wt}\ x] \wedge [^0 Discovered_{wt}\ x]]],$$

then the Composition

$$[^0 Sub\ [^0 Tr^{((\iota\tau)\omega)} f]\ ^0 g\ ^0[\lambda w' \lambda t' [^0 Planet_{w't'}\ g_{w't'}]]]$$

$v(FKBO/f)$-produces the Closure $[\lambda w' \lambda t'\ [^0 Planet_{w't'}\ ^0 FKBO_{w't'}]]$, and this procedure differs significantly from the one Tilman was originally related to, namely the Closure

$$[\lambda w' \lambda t'\ [^0 Planet_{w't'}\ \lambda w \lambda t\ [^0 First\ \lambda x\ [[^0 KuiperBeltObj_{wt}\ x] \wedge [^0 Discovered_{wt}\ x]]]_{w't'}]]$$

Hence, we can only derive that there is a *procedure* producing the office of *FKBO* such that Tilman believes that whichever the produced office may be, its occupant is a planet:

$$(\mathrm{EA}_3)\ \frac{\lambda w \lambda t\ [^0 Believe_{wt}\ ^0 Tilman\ ^0[\lambda w' \lambda t'\ [^0 Planet_{w't'}\ \lambda w \lambda t\ [^0 First\ \lambda x\ [[^0 KuiperBeltObj_{wt}\ x] \wedge [^0 Discovered_{wt}\ x]]]_{w't'}]]]}{\lambda w \lambda t\ [^0 \exists\ \lambda c\ [^0 Believe_{wt}\ ^0 Tilman\ [^0 Sub\ c\ ^0 d\ ^0[\lambda w' \lambda t'\ [^0 Planet_{w't'}\ d_{w't'}]]]]]}$$

Additional types: $c/*_2 \rightarrow *_1$; $^2c/*_3 \rightarrow \iota_{\tau\omega}$; $d/*_1 \rightarrow \iota_{\tau\omega}$.

[52] See Salmon (2010) on arguments for considering β-conversion invalid. See Jespersen (2015a) for (favourable) discussion of Salmon's examples.

A further variation is this. Suppose we type Vulcan as an office rather than an individual (as we must, since TIL has no category of inexistent or merely possible individuals) and that the office of Vulcan comes with the constraint that its respective occupants must each be a planet (as we should in order to align with Le Verrier's specification of Vulcan).[53] Suppose also that it is a datum that Tilman believes that Vulcan is not (defined as) a planet. Then what Tilman believes is necessarily (in this case, analytically) false, because his belief about Vulcan runs afoul of the stipulated definition of Vulcan. What can be derived from his attitude is that there exists an office such that Tilman believes that this office comes without the analytic property of being a planet. This attitude must be construed as a hyperpropositional one, in order to distinguish it from other analytically impossible complements.[54]

The moral to be drawn from the above examples is this. We can quantify over and *into* hyperpropositional contexts, even though the procedure occurring as attitude complement is just *displayed* in such contexts. We can quantify over *procedural complements* of attitudes and over their *procedural constituents*. And we can also quantify over those *objects* that are produced by the constituents of an attitudinal complement, but only if the respective object is displayed by Trivialisation. In the following section, we formulate two rules for quantifying into hyperpropositional contexts de dicto.

4.2 Rules for quantifying into hyperpropositional contexts

As we both warned and promised at the outset, a logic of quantifying-in comes at the end of a long story. We have now obtained everything required to introduce and prove our *general rules for quantifying into hyperpropositional attitudes de dicto*. These are the hardest cases of quantifying-in, and they constitute the core of our novel contribution. Rule (1) quantifies over a *procedure* that is a constituent of an attitude complement, and rule (2) quantifies over an *object* presented by a Trivialisation that is a constituent of an attitude complement.

RULE 1. *Quantifying over a constituent of an attitude complement.*

$$\frac{[B_{wt}\, a\, {}^0P(X/d)]}{[{}^0\exists\, \lambda c[B_{wt}\, a\, [{}^0Sub\, c\, {}^0d\, {}^0P(d)]]]}$$

Types: $P(X/d)/*_n$: a procedure with a constituent $X/*_n \rightarrow \alpha$ that has been substituted for the variable $d/*_n \rightarrow \alpha$; $c/*_{n+1} \rightarrow *_n$; ${}^2c \rightarrow \alpha$.

Proof

(1)	$[B_{wt}\ a\ {}^0P(X/d)]$	$\varnothing$
(2)	$[B_{wt}\ a\ [{}^0Sub\ {}^0X\ {}^0d\ {}^0P(d)]] = [B_{wt}\ a\ {}^0P(X/d)]$	Definition 11
(3)	$[\lambda c\ [B_{wt}\ a\ [{}^0Sub\ c\ {}^0d\ {}^0P(d)]]\ {}^0X]$	2, β-expansion
(4)	$[{}^0\exists\ \lambda c\ [B_{wt}\ a\ [{}^0Sub\ c\ {}^0d\ {}^0P(d)]]]$	3, Definition 4

[53] Such a constraint is known as a *requisite* in TIL. See Duží et al. (2010, §4.1).

[54] See Duží et al. (2021) for our logic of analytically impossible *hyperoffices*.

Remark Step (4) is truth-preserving, because the class of procedures produced by $\lambda c\ [B_{wt}\ a\ [^0Sub\ c\ ^0d\ ^0P(d)]]$ is non-empty, as it contains at least the procedure *X*. Actually, since ^0X is a proper constituent of $[B_{wt}\ a\ [^0Sub\ ^0X\ ^0d\ ^0P(d)]]$, we could have just applied the (∃I)-rule on the left-hand side of step (2), omitting step (3). Yet, for the sake of clarity, we proceeded from (2) to (4) via (3).

Finally, we have the rule for quantifying over an object such that its Trivialisation is a constituent of a procedure that occurs as attitude complement.

RULE 2. *Quantifying over a Trivialised object.*

$$\frac{[B_{wt}\ a\ ^0P(^0b/y)]}{[^0\exists\ \lambda x\ [B_{wt}\ a\ [^0Sub\ [^0Tr^{\alpha}\ x]\ ^0y\ ^0P(y)]]]}$$

Types: $P(^0b/y)/*_n$: a procedure with a proper constituent $^0b/*_n \rightarrow \alpha$ that has been substituted for the variable $y/*_n \rightarrow \alpha$; $x/*_n \rightarrow \alpha$.

Proof

(1)	$[B_{wt}\ a\ ^0P(^0b/y)]$	∅
(2)	$[B_{wt}\ a\ [^0Sub\ [^0Tr^{\alpha}\ ^0b]\ ^0y\ ^0P(y)]] = [B_{wt}\ a\ ^0P(^0b/y)]$	Definition 11
(3)	$[\lambda x\ [B_{wt}\ a\ [^0Sub\ [^0Tr^{\alpha}\ x]\ ^0y\ ^0P(y)]]\ ^0b]$	2, β-expansion
(4)	$[^0\exists\ \lambda x\ [B_{wt}\ a\ [^0Sub\ [^0Tr^{\alpha}\ x]\ ^0y\ ^0P(y)]]]$	3, Definition 4

Step (4) is justified, because the class of α-objects produced by the Closure $\lambda x\ [B_{wt}\ a\ [^0Sub\ [^0Tr^{\alpha}\ x]\ ^0y\ ^0P(y)]]$ is non-empty, as it contains at least the object *b*.

Note that, in Rule 2, *x* occurs *free* in the Composition $[^0Sub\ [^0Tr^{\alpha}\ x]\ ^0y\ ^0P(y)]$, whereby it lends itself to being λ-bound.

5 Conclusion

The above rules are *almost* trivial, as indeed they should be. After all, quantifying over and into hyperintensional contexts should, *in principle*, if not technically, be as trivial as quantifying into extensional contexts, *as soon as we have availed ourselves of a fully transparent semantics and an extensional logic of hyperintensions* that is strictly compositional. However, the near-triviality of the rules can be obtained only by means of a theory possessing great expressive power (hence the qualification 'almost trivial'). We obtain a sufficient measure of expressive power by means of ramification of procedures so that we can *display* any given procedure by going one level higher. The *Sub* function would become inoperative if it were not for procedures occurring displayed, i.e., presented as functional arguments. The technical finesse the conclusion exhibits resides in the fact that the λ-bound variables over whose values we quantify occur *outside* the hyperintensional context of the displayed procedures.

It is critical for a theory such as Transparent Intensional Logic to demonstrate exactly how it makes quantifying-in come out valid also with respect to hyperpropositional attitudes de dicto. The theory claims to have an extensional logic of hyperintensions, so the extensional rule of existential quantification had better be valid for hyperintensional contexts too. But the impact of the results above extends beyond our particular theory. First, it is good news for the community-wide project devoted to developing and substantiating theories of hyperintensionality that quantifying-in is not unattainable. And second, the demonstrated feasibility of quantifying-in provides the hyperintensional community with a touchstone: does a theory of hyperintensionality succeed in rendering quantifying-in valid? The theory scores a point if it does.

Acknowledgements We are much indebted to two anonymous referees for *Linguistics and Philosophy* whose careful remarks helped improve the quality of this paper. The paper rounds out the quantifying-in trilogy whose other two instalments are our (2012) and (2015). This research has been supported by the University of Oxford project *New Horizons for Science and Religion in Central and Eastern Europe* funded by the John Templeton Foundation, as well as by Grant No. SP2021/87, VSB-Technical University of Ostrava, Czech Republic, *Application of Formal Methods in Knowledge Modelling and Software Engineering IV.*

References

Anderson, C. A. (1998). Alonzo Church's contributions to philosophy and intensional logic. *The Bulletin of Symbolic Logic, 4*, 129–171.

Bealer, G. (1982). *Quality and concept*. Oxford: Clarendon Press.

Berto, F., & Nolan, D. (2021). Hyperintensionality. In E. N. Zalda (Ed.), *The Stanford Encyclopedia of Philosophy* (Summer 2021 edition). https://plato.stanford.edu/archives/sum2021/entries/hyperintensionality/.

Caie, M., Goodman, J., & Lederman, H. (2019). Classical opacity. *Philosophy and Phenomenological Research*. https://doi.org/10.1111/phpr.12587.

Church, A. (1956). *Introduction to mathematical logic*. Princeton, NJ: Princeton University Press.

Copi, I. (1968). *Introduction to logic*. New York: MacMillan.

Crawford, S. (2008). Quantifiers and propositional attitudes: Quine revisited. *Synthese, 160*, 75–96.

Cumming, S. (2008). Variabilism. *Philosophical Review, 117*, 525–554.

Duží, M. (2019). If structured propositions are logical procedures then how are procedures individuated? *Synthese, 196*, 1249–1283.

Duží, M. (2018a). Negation and presupposition, truth and falsity. *Studies in Logic, Grammar and Rhetoric, 54*, 15–46.

Duží, M. (2018b). Logic of dynamic discourse and anaphora resolution. In V. Sornlertlamvanich, P. Chawakitchareon, A. Hansuebsai et al. (Eds.), *Information modelling and knowledge bases XXIX* (pp. 263–279). Amsterdam: IOS Press.

Duží, M. (2017). Presuppositions and two kinds of negation. *Logique et Analyse, 239*, 245–263.

Duží, M. (2012). Towards an extensional calculus of hyperintensions. *Organon F, 19* (supplementary issue 1), 20–45.

Duží, M., & Menšík, M. (2020). Inferring knowledge from textual data by natural deduction. *Computación y Sistemas, 24*, 29–48.

Duží, M., & Jespersen, B. (2012). Transparent quantification into hyperpropositional contexts de re. *Logique et Analyse, 220*, 513–554.

Duží, M., & Jespersen, B. (2013). Procedural isomorphism, analytic information, and β-conversion by value. *Logic Journal of the IGPL, 21*, 291–308.

Duží, M., & Jespersen, B. (2015). Transparent quantification into hyperintensional objectual attitudes. *Synthese, 192*, 635–677.

Duží, M., Jespersen, B., & Materna, P. (2010). *Procedural semantics for hyperintensional logic: Foundations and applications of Transparent Intensional Logic*. Dordrecht: Springer.

Duží, M., Glavaničová, D., & Jespersen, B. (2021). Impossible individuals as necessarily empty individual concepts. In A. Giordani & J. Malinowski (Eds.), *Logic in high definition: Trends in logic*, (pp. 177–202). Cham: Springer.

Duží, M., & Fait, M.(2021). A hyperintensional theory of intelligent question answering in TIL. In R. Loukanova (Ed.), *Natural Language Processing in Artificial Intelligence. NLPinAI 2020*, (pp. 69–104). Cham: Springer.

Duží, M., & Kosterec, M. (2017). A valid rule of β-conversion for the logic of partial functions. *Organon F, 24*, 10–36.

Gordon, M. J. C., & Melham, T. F. (1993). *Introduction to HOL*. Cambridge: Cambridge University Press.

Jespersen, B. (2021). First among equals: Co-hyperintensionality for structured propositions. *Synthese*, *199*, 4483–4497. https://doi.org/10.1007/s11229-020-02987-4.

Jespersen, B. (2019). Anatomy of a proposition. *Synthese, 196*, 1285–1324.

Jespersen, B. (2015a). Qualifying quantifying-in. In A. Torza (Ed.), *Quantifiers, quantifiers, and quantifiers*: *Themes in logic, metaphysics and language* (pp. 241–269). Cham: Springer.

Jespersen, B. (2015b). Should propositions proliferate? *Thought, 4*, 243–251.

Kaplan, D. (1968). Quantifying-in. *Synthese, 19*, 178–214.

Kaplan, D. (1986), Opacity. In L. E. Hahn & P. A. Schilpp (Eds.), *The philosophy of W.V. Quine* (pp. 229–289). La Salle, IL: Open Court.

Klement, K. C. (2002). *Frege and the logic of sense and reference*. New York: Routledge.

Materna, P. (2005). Ordinary modalities. *Logique et Analyse, 189–192*, 57–70.

Materna, P. (1997). Rules of existential quantification into 'intensional contexts'. *Studia Logica, 57*, 331–343.

Morton, A. (1969). Extensional and non-truth-functional contexts. *Journal of Philosophy, 66*, 159–164.

Moschovakis, Y. N. (2006). A logical calculus of meaning and synonymy. *Linguistics and Philosophy, 29*, 27–89.

Pickel, B. (2015). Variables and attitudes. *Noûs, 49*, 333–356.

Quine, W.v.O. (1956). Quantifiers and propositional attitudes. *Journal of Philosophy, 53*, 177–187.

Quine, W. v. O. (1960). *Word and object*. Cambridge, MA: MIT Press.

Salmon, N. (2010). Lambda in sentences with designators: An ode to complex predication. *Journal of Philosophy, 107*, 445–468.

Tichý, P. (1986). Indiscernibility of identicals. *Studia Logica*, 45, 251–273. Reprinted in V. Svoboda, B. Jespersen & C. Cheyne (Eds.), *Pavel Tichý's collected papers in logic and philosophy*. Prague/ Dunedin: Filosophia,Czech Academy of Sciences/University of Otago Press, 2004.

Tichý, P. (1988). *The foundations of Frege's logic*. Berlin: De Gruyter.

Turner, R. (1992). Properties, propositions and semantics theory. In M. Rosner & R. Johnson (Eds.), *Computational linguistics and formal semantics* (pp. 159–180). Cambridge: Cambridge University Press.

Yalcin, S. (2015). Quantifying in from a Fregean perspective. *Philosophical Review, 124*, 207–253.