

INTRODUCTION

Bjørn Jespersen & Marie Duží

The topic of this special issue of *Synthese* is hyperintensionality. This introduction offers a brief survey of the very notion of hyperintensionality followed by a summary of each of the papers in this collection. The papers are foundational studies of hyperintensionality accompanied by ample philosophical applications.

Hyperintensionality concerns the individuation of non-extensional entities such as propositions and properties, relations-in-intension and individual roles, as well as, for instance, proofs and judgments and computational procedures, in case these do not reduce to any of the former. Hyperintensional individuation is frequently also referred to as ‘fine-grained’ or sometimes simply ‘intensional’ individuation, when ‘intensional’ is not understood in the specific sense of possible-world semantics or in the pejorative sense of flouting various logical rules of extensional logic. A principle of individuation qualifies as hyperintensional as soon as it is finer than necessary equivalence. A hyperintensional principle of individuation bars necessary equivalence from entailing identity, making logically possible the cohabitation of necessary equivalence and non-identity between a pair of fine-grained entities A, B :

$$A \Leftrightarrow B \wedge A \neq B$$

The main reason for introducing hyperintensionality was originally to block various inferences that were argued on philosophical grounds to be invalid. The theoretician introduces a notion of hyperintensional context, in which the proper substituends are hyperintensions rather than the modal intensions of possible-world semantics or extensions. The result is that far fewer substitutions go through than if the substituends were modal intensions or extensions. One task becomes how to decide which of all the various contexts are the hyperintensional ones. This is in effect the problem of determining which contexts exceed the resources of modal logic. Another task becomes how to decide which hyperintensions can be substituted for which other hyperintensions inside which hyperintensional contexts. This is in effect the problem of providing a positive definition of the calibration of hyperintensional individuation. Just how close the discrimination ought to be remains an entirely open and lively research question. Obviously, one theory may want to adopt more than one measure of hyperintensional individuation to suit different theoretical purposes.

Here is a standard way of framing the sort of problem that hyperintensionality is called upon to solve, namely in terms of which substitutions are valid:

$$\text{Know}(a, A)$$
$$A \Leftrightarrow B$$

$$\text{Know}(a, B)$$

If this (rudimentary) schema is valid, then by knowing one thing, A , agent a will, without further effort on a 's behalf, also know each and every other thing, B , that is necessarily equivalent to A . This schema turns a into a logical and mathematical genius, because if A is a logical or mathematical truth then every logical or mathematical truth is necessarily equivalent to A . And if A is an empirical truth then a will be miraculously able to perform all the logical operations turning A into a logically equivalent truth B . For instance, if what a knows is that the Czech Republic has more breweries than Belgium then a will ipso facto also know that Belgium has fewer breweries than the Czech Republic. Individuating pieces of knowledge up to necessary equivalence is the backbone of the problem of logical omniscience, which continues to haunt standard epistemic logic. A classical example (inspired by Leibniz) of an inference that needs to be blocked is that a may know that a given set of geometric figures is a set of equilateral figures without also knowing that the set is a set of equiangular figures. The underlying philosophical view being appealed to is that the objects of knowledge, and presumably of other attitudes as well, must be individuated very minutely in order to faithfully represent what the agent does, and does not, know, believe, hope, etc.

The standard move has become to argue that the operator 'knows' induces a hyperintensional context. This move makes the second premise, $A \leftrightarrow B$, irrelevant (in casu that x has more F s than y exactly when y has fewer F s than x), even when A, B are themselves hyperintensional entities. The sort of premise required to validate the conclusion must state that A, B are hyperintensions that are related according to a principle of hyperintensional equivalence still to be decided upon. The quest is for a philosophically adequate and formally impeccable criterion of hyperintensional equivalence, or co-hyperintensionality, to underpin the second premise in the schema below, where *Know* goes proxy for any hyperintensional operator:

$$\frac{\begin{array}{c} \textit{Know}(a, A) \\ A, B \textit{ are co-hyperintensional} \end{array}}{\textit{Know}(a, B)}$$

It is important to bear in mind that hyperintensional granularity was originally only negatively defined (because the main purpose was to block various arguments). The original definition does not address the question of just 'how hyper' hyperintensions are, i.e. exactly how fine-grained co-hyperintensionality is. This leaves room for various positive definitions of fine-grained granularity. Deciding on at least one technically satisfactory and philosophically adequate upper bound is crucial. Otherwise we cannot determine which substitutions within hyperintensional contexts are valid. We can only determine one class of substitutions that are invalid, namely those as in the first schema above where a mere necessary equivalent, B , is offered as a substituent.

While it is obvious that hyperintensions are needed for meanings and attitudes, it is an open research question whether hyperintensional distinctions extend beyond the sphere of conceptualization so as to encompass also at least some portions of empirical reality. Nolan (2014) argues in favour of hyperintensional distinctions in the 'worldly' domain as well.

Williamson (2013, p. 217, p. 266) maintains, on the contrary, that modal distinctions suffice for the worldly domain. The compass of hyperintensional distinctions is a rich philosophical question, translating into the formal question of which contexts qualify as hyperintensional. How a given theory of hyperintensions conceives of mathematical and logical entities will also affect whether the theory brings such entities within the purview of hyperintensionality.

For historical background, hyperintensional individuation was originally put forward by Cresswell (1975) in direct opposition to the extensional, hence coarse-grained individuation of intensional entities such as X, Y that characterizes possible-world semantics where necessary equivalence is indiscernible from identity:

$$(X \Leftrightarrow Y) \supset (X = Y)$$

Specifically, where f, g are functions defined over a logical space or domain of possible worlds, the intensions of possible-world semantics are individuated in such a way that necessary co-extensionality is indiscernible from co-intensionality (the direction from left to right being the noteworthy one):

$$\forall fg (\forall w (f_w = g_w) \supset f = g)$$

The combination of the modern-day notion of functions as mappings and possible worlds as functional arguments w_i makes for a formally precise principle of individuation of intensional entities. The criterion has the virtue of being mathematically manageable, thereby enabling an algebra of intensional entities, which makes it possible to perform calculations with and about intensions. The criterion also seems to settle the age-old philosophical question of how to individuate intensions. Possible-world semantics has, by even the most exacting standards, been one of analytic philosophy's success stories. Nonetheless, it is important to be perfectly clear about what the intensions of possible-world semantics are, and what they are not. Thus, for instance, a possible-world proposition is nothing but the modal profile of a truth-condition. It is nothing but a set of indices or points of evaluation or, equivalently, its characteristic function. What the notion does is allow us to distinguish between truth-conditions in terms of whether they are contingently or necessarily or not possibly satisfied, and in terms of whether any pair of truth-conditions are just one truth-condition; and that's it. The upshot is that if possible-world propositions are total functions, leaving no room for truth-value gaps, there is but one necessary proposition and but one impossible proposition. Furthermore, as the brewery example above showed, possible-world propositions cannot distinguish between inverse relations. While possible-world intensions are insufficient as linguistic meanings and attitude complements, they are arguably still indispensable, namely in order to discharge a host of modal tasks. But one basic insight we have acquired is that there is more to meanings and attitudes (at least those that are not logically closed) than just their modal profile.¹

¹ Stalnaker remains adamant that truth-conditions, in the form of possible-world propositions, exhaust the semantics of sentences. Any additional linguistically salient material is drawn from 'general conversational rules'. The policy Stalnaker adheres to combines an intensional semantics with a hyperintensional pragmatics. A recent reference would be his (2002).

As a matter of historical fact, at the very inception of possible-world semantics Carnap (1947, §13) noted that some contexts are neither extensional nor intensional but, as we would say nowadays, hyperintensional; Carnap's example was a belief context. In (ibid., §15) Carnap attributed to C.I. Lewis the then-recent insight that, "Not every pair of expressions having the same intension would be called synonymous". Carnap offered in effect his notion of intensional isomorphism as a generalized encapsulation of Lewis's insight. Church (1954) found Carnap's characterization of hyperintensionality flawed, proposing his own alternative, called synonymous isomorphism. Church would put forward various so-called Alternatives, each of which offering a specific calibration of synonymous isomorphism couched in the logic of functions he developed, the λ -calculus. Mates (1952) argued that, despite the synonymy between 'to chew' and 'to masticate', it is false that nobody doubts that whoever believes that x chews, believes that x masticates. Mates's puzzle, as it has become known, would appear to demand of a theory of hyperintensional individuation that it cut hyperintensions so finely so as to accommodate a logical distinction even between pairs of synonymous terms. It is important for any theory of hyperintensions to take a principled stand on whether hyperintensionality is to coincide with synonymy or exceed it. More recent examples of pairs of synonymous lexical terms include 'is a woodchuck'/'is a groundhog' and 'is furze'/'is gorse' (assuming the latter pair are not just morphological variants), to the exclusion of pairs of a lexical and a compound term like 'lasts a fortnight'/'lasts fourteen days' or 'lasts two weeks'. The latter sort fuels instead the so-called paradox of analysis much discussed by Church. The paradox of analysis raises slightly more intricate issues, such as whether the members of the pair are synonymous or just equivalent, and if synonymous then maybe in a weaker sense than applies to 'is a cougar'/'is a puma', which may arguably just be a matter of semantic redundancy (i.e. a strictly synonymous translation from one predicate to another within the same language). A theory of hyperintensionality must also take a principled stand on how to approach standard substitution examples from philosophy of language bearing on the likes of 'Superman' vs. 'Clark Kent', 'Hesperus' vs. 'Phosphorus', 'London' vs. 'Londres' (Kripke's Pierre case), or 'Paderewski'-the-statesman vs. 'Paderewski'-the-pianist. In particular, do intensional distinctions suffice for these cases, or must hyperintensional distinctions be invoked?

However hyperintensional theoreticians may position themselves, they will have to address thorny issues such as whether hyperintensional entities have parts that are arranged within a structure, whether the structure (if any) of a hyperintensional entity tracks syntactic structure, and if so, of which formalism or notational system. For instance, Carnap's intensional isomorphism draws critically on how bits of meaning are arranged within an encompassing structure. The notions of hyperintensionality and structured meaning need not go hand in hand, for primitive hyperintensions (or hyperintensional primitives) are a theoretical option, but these two notions have often been explored and developed in tandem. In fact, when Cresswell (1985) and Kaplan (1978) began, independently of one another, to reinvoke structured meaning they were in part motivated by hyperintensionality concerns, in Kaplan's case not least the un-Fregean distinction between singular and general propositions.² Still it is

² Though both Cresswell and Kaplan carved out a niche for structured meaning, they attempted to fill the niche with set-theoretic sequences. While this strategy allows them to remain within model theory,

important to stress that, narrowly understood, hyperintensionality is a matter of logic rather than philosophy of language. Semantic and linguistic issues are, strictly speaking, an overlay. The logical brief is, firstly, to rule out various inference schemas as invalid; secondly, to rule in various other schemas as valid. The theoretician must take an interest in the functional question what hyperintensions do (what functions or tasks they discharge). It is optional to take an additional interest in the structural question how hyperintensions do what they do. To address the structural question is not necessarily to put forward some enabling structure that hyperintensions would themselves have. One may be agnostic about whether hyperintensions are structured (and if so, how) or maintain that they are not, while locating structure in one's algebra or syntax that represents hyperintensions (much the same way we say that sets are structured by this or that algebra, though sets themselves lack structure). On the note of structure, the theoretician must also explain how the compositionality constraint is respected by either hyperintensions themselves or their syntactic representations, and also how sentential structure and propositional structure are correlated. In any event, addressing both the functional and the structural question makes arguably for a philosophically much richer theory of hyperintensional distinctions and hyperintensional entities by opening up a metaphysical angle as well. But it is controversial whether the hyperintensionalist is best advised to broach also metaphysical topics or remain close to the original challenge to prevent non-extensional contexts from validating unwarranted conclusions.

This collection contains five papers written especially for this special issue on the basis of a call for papers. Below follow summaries of each paper individually.

Carl Pollard's paper, 'Agnostic hyperintensional semantics', puts forward a hyperintensional semantics for natural language which, as he says, is 'agnostic about the question of whether propositions are sets of worlds or worlds are (maximal consistent) sets of propositions'. Montague's theory of intensional senses is replaced by a weaker theory, written in standard classical higher-order logic. Montague's theory can then be recovered from the proposed theory by identifying the type of propositions with the type of sets of worlds and adding an axiom to the effect that each world is the set of propositions that are true there. In Pollard's theory, propositions, worlds, and entities are primitives, which are interrelated by the axioms. The theory does not tell what these things are, no more so than an axiomatic set theory tells what sets are. Senses are in a many-to-one correspondence with intensions. Propositions form a pre-boolean algebra pre-ordered by entailment, and for each world w the set of facts of w forms an ultra-filter s , which means that the set of propositions that belong to s is closed under conjunction, closed under entailment, does not have falsity as a member, and for every proposition p , s has either p or not- p as a member. Senses are compositionally assigned to linguistic expressions by a linear categorial grammar (LCG). An explicit grammar fragment is provided that illustrates the compositional assignment of senses to a variety of constructions, including dummy-subject constructions, infinitive complements, predicative adjectives and nominals, raising to subject, 'tough-movement', and quantifier scope ambiguities. Notably, the grammar and the derivations that it licenses never make

which is a firmly set-theoretic enterprise, sequences underdetermine structure and are also unfit as propositions. For a recent critique of tuples as structured propositions, see Jespersen (2012, §3).

reference to either worlds or to the extensions of senses. It makes the composition of senses straightforward. The straightforwardness arises from the fact that the syntactic combinatorics driving the semantic composition are based on the implicative fragment of linear logic (in a sequent-style natural-deduction presentation), which has only two rules and one logical axiom scheme. The two rules, implication elimination (modus ponens) and implication introduction (hypothetical proof), correspond to the combination and abstraction of sense terms. LCG draws on five main ideas. First, syntactic analyses of linguistic expressions are logical proofs. Second, phenogrammatical structure (concrete syntax), which has to do with surface form, should be systematically distinguished from tectogrammatical structure (abstract syntax), which has to do with semantically relevant combinatorics. Third, the meaning of a complex linguistic expression is recursively composed from the meanings of its immediate syntactic constituents, with the recursion grounded in the meanings of lexical items specified by the grammar. Fourth, the phenogrammatical component of an expression is not a string, as per Montague, but rather a term which denotes a string, or a (possibly higher-order) function over strings, and is recursively composed from the phenogrammatical structures of the immediate constituents, in parallel with the meaning composition. And fifth, rather than composing references of expressions, Pollard instead composes senses, so that the grammar defines a relation between phenogrammatical structure (roughly, surface forms), tectogrammatical structures (syntactic types), and meanings (hyperintensional senses) without ever involving reference at all.

Chris Fox and Shalom Lappin's paper, 'Type-theoretic logic with an operational account of intensionality', proposes *Property Theory with Curry Typing* (PTCT), reformulated within a typed predicate logic as an alternative framework for fine-grained intensional (i.e. hyperintensional) semantic representation. PTCT has first-order formal power, but it emulates the expressive richness of higher-order systems. In a 2005 version PTCT was specified as a federated tripartite framework consisting of (i) an untyped lambda-calculus, which generates the language of terms, (ii) a rich Curry-style typing system for assigning types to terms, (iii) a first-order language of well-formed formulas for reasoning about the truth of propositional terms, where these are term representations of propositions. A tableau proof theory constrains the interpretation of each component of this unified representational language, and it relates the expressions of the different components. Restrictions on each component prevent semantic paradoxes. The identity of PTCT terms is governed by the α -, β -, and η -rules of the λ -calculus, hence identity is λ -equivalence. But in addition to this, the terms of the untyped λ -calculus are interpreted as encoding computable functions. In the re-formulation of PTCT that Fox and Lappin propose here these three components of the framework are expressed in a single unified first-order typed predicate logic through its proof theory. The authors characterize the difference between fine-grained intensions and extensions in terms of the distinction between the operational and the denotational interpretations of computable functions. Thus two propositions are extensionally equivalent if they share the same truth-conditions, and two properties are extensionally equivalent if they apply to the same terms, but identity of-truth conditions does not make them identical. While theories of fine-grained intensionality may avoid the reduction of intensional identity to provable equivalence, many of them do not go beyond a bare inscriptionalist treatment of intensional distinctions. Therefore they leave the notion of hyperintension ineffable. In PTCT intensional difference consists in the operational distinctions among computable functions, and extensional identity is

the denotational equivalence of the values that functions compute. This account grounds fine-grained intensionality in a way that naturally accommodates cases of intensional difference combined with provable denotational equivalence. The authors characterize the distinction between intensional identity and provable equivalence computationally by invoking the contrast between operational and denotational semantics in programming languages. They adduce two examples of hyperintensionally different definitions of the same set. First they provide two distinct definitions of the set of predecessors, one in terms of the predecessor relation, the other in terms of the successor relation. Second, they provide two equivalent but intensionally different grammars generating the same language $\{a^n b^n c^n \mid 1 \leq n\}$. The authors also compare their theory to Muskens's and Moschovakis's respective approaches to hyperintensionality. Both Muskens and Moschovakis identify the sense of an expression with an algorithm for computing its denotation. There are two points of difference between Muskens's theory and PTCT. First, Muskens applies the technique of logic programming to encode senses, while in PTCT the analysis is developed in the functional programming framework. Yet this is not a great difference, because the same algorithm can be formulated in any programming language. The second difference concerns whether to invoke possible worlds. While Muskens defines intensions as functions with domain in possible worlds in order to capture modalities, PTCT yields a radically non-modal view of intensions in which possible worlds play no role in their specification or their interpretation. An intension is identified directly with the sequence of operations performed in computing the value of the function that expresses it. Though there may be independent epistemic or semantic reasons for incorporating possible worlds into one's general theory of interpretation, according to the authors, possible worlds are not required for an adequate explanation of fine-grained intensionality. Concerning Moschovakis, there are two points of difference between Moschovakis's algorithmic theory of senses and the account proposed by PTCT. First, while in PTCT α -, β - and η -reduction sustain intensional identity, in Moschovakis's language β -reduction does not. Yet this is not a deep question of principle. It is possible to narrow the specification of intensional identity in PTCT to exclude β - (as well as η -, and even α -) reduction, without altering the proposed account of senses as computable functions. This would simply involve imposing a particularly fine-grained notion of intensional identity. The second point of difference is more significant. Moschovakis specifies a Kripke-frame semantics for his language which is a variant of Montague's possible worlds models (referring to them as 'Carnap states'). These are n -tuples of indices corresponding to worlds, times, speakers, and other parameters of context. Senses are characterized as algorithmic procedures for computing the denotation of a term relative to a world and the other elements of such an n -tuple. Therefore, as with Muskens's theory, Moschovakis's operational view of intensions treats them as inextricably bound up with possible worlds. According to the authors an important advantage of the proposed PTCT account is that it factors modality and possible worlds out of the specification of intensions.

Mark Jago's paper, 'Hyperintensional propositions', offers an outline of propositions that seeks to unify two of the dominant theories, namely propositions as sets of worlds and propositions as structures. The general picture that emerges is this. Let particulars, properties, relations, etc., be given; organize them as elements of tuples (which neo-Russellians take to be structures); a set of tuples is a *world*; a set of worlds is a *proposition*. Jago takes a lead from Jeffrey King's version of Russellian propositions, according to which propositional structure

tracks sentential structure. Jago goes one crucial step further by either literally identifying sentences with propositions or else having them come out isomorphic, with any residual differences being philosophically and logically irrelevant. Sentences play a representational role, which includes describing logically impossible scenarios, such as an individual both having and lacking a left foot at the same time, or two sentences being true without their conjunction being true. Jago adopts a Lagadonian conception of language (as originally introduced by Swift in *Gulliver's Travels* and later used by Lewis in his description of what he calls linguistic ersatzism), which identifies terms and expressions with their semantic values, such as particulars, properties, relations, etc. For instance, to talk about one's left foot one does not say 'my left foot' but instead exhibits one's left foot. Jago also adopts an ersatzist conception of worlds on which both possible and impossible worlds are linguistic constructions. Jago obtains fine granularity for his propositions, in that a proposition is a set of fine-grained (possible or impossible) worlds, the proposition being identical or at least isomorphic to a sentence from a Lagadonian language. Just how fine-grained the syntax of Lagadonian sentences is depends on, for instance, the ability to differentiate, at an impossible world, between particulars a and b when a is the same entity as b .

Bartosz Wieckowski's paper, 'Constructive belief reports', develops a theory of hyperintensional attitude contexts that is a variant of Martin-Löf's constructive type-theoretic semantics. Wieckowski compares his theory to Ranta's from the 1990s, which was also built on Martin-Löf's type theory, but unlike Wieckowski's shares affinities with Hintikka's coarse-grained, quantifier-based analysis of attitudes. Since Ranta uses the type-theoretic universal quantifier as belief operator, his approach turns out to be too restrictive, and also the very fine degree of hyperintensional individuation Wieckowski argues for is beyond the pale for Ranta. Since Wieckowski works within a constructive type-theoretic semantics, his task is to define formation and equality, introduction and elimination rules for the operators that generate attitude-reporting propositions. In accordance with the Curry-Howard isomorphism, a proposition is identified with the set of its proof-objects. Therefore, Wieckowski's task boils down to defining proof-theoretic rules of set membership for proof-objects that validate attitude reports. This is done by laying down what conditions are to be met in order to be entitled to make the judgement that a certain proof-object is an element of a set, i.e. that a certain proposition is true. A characteristic feature of a constructive type-theoretic account of attitude reports such as belief or knowledge reports is that their truth is explained in terms of judgments. For instance, to use one of Wieckowski's examples, when there is a unicorn about which Mary entertains the belief *de nomine* that it walks then Mary judges as true a proposition that is composed of *inter alia* a predicative concept of 'unicorn' and a predicative concept of 'walking'. Wieckowski's predicative concepts are so fine-grained that he is able to distinguish between 'groundhog' and 'woodchuck'. This is because the two concepts have two distinct subatomic (sub-propositional) proof-objects.

Marie Duží and *Bjørn Jespersen's* paper, 'Transparent quantification into hyperintensional objectual attitudes', claims that hyperintensional individuation is procedural individuation and that semantically significant structures are procedures understood as abstract objects in their own right. Duží and Jespersen motivate philosophically and develop formally a criterion of hyperintensional individuation within the framework of Tichý's Transparent Intensional Logic

(TIL). TIL incorporates tenets not least from Frege and Russell, Carnap, Church and Montague. The upshot is a theory that comes with a ramified type hierarchy encompassing both higher-order and first-order entities, both structured hyperintensions and possible-world intensions, and sufficient expressive power to provide a principled account of how to existentially quantify into hyperintensional attitude contexts and over both higher-order and first-order entities. This particular paper demonstrates how to quantify into such non-propositional attitude contexts as expressed by “Agent a calculates the last decimal of pi” and “Agent b is seeking a yeti without seeking an abominable snowman”. It follows that there is an x such that a is calculating x , and there is a y and there is a y' such that a is seeking y without seeking y' . But x , y , y' are very far from being extensional entities such as individuals or numbers. TIL is an extensional logic of intensions and hyperintensions, and the adjacent semantics is designed to preserve referential transparency in all contexts (see also Duží 2012). Therefore, all the rules of extensional logic are valid, including Leibniz’s Law and existential generalization. But which (abstract) entities qualify as proper substituends depends on whether the substitution is performed inside an extensional or an intensional or a hyperintensional context. Inside a hyperintensional context the only proper substituends are hyperintensional entities (so-called constructions) that are procedurally isomorphic to the original hyperintensional entity. Procedural isomorphism, in its most recent version as defined in this paper, is a variant of Church’s Alternative (A1). α -conversion is preserved, of course, but formulated so as to accommodate a technical detail particular to TIL. Also β -conversion is preserved, but only in the form of conversion by value.

As this batch of papers hopefully makes clear, present-day research on hyperintensionality has made the critical step from the programmatic to the nitty-gritty. This marks the progression from projects to theories. What is hopefully also evident is that different researchers cater to different if overlapping audiences, ranging from linguists over philosophers of language, philosophers of mathematics, and philosophical logicians to theoretical computer scientists. Any discipline that is alert to fine-grained distinctions is tasked with developing and cultivating a notion of hyperintensionality. This may well turn out to apply to additional branches of philosophy, such as metaphysics, epistemology and, say, value theory. Even some fragments of modal logic seem to be in need of hyperintensional distinctions to accommodate some particular modalities. One obvious example would be so-called counterpossibles, which are counterfactual conditionals whose antecedent is a necessary falsehood; without hyperintensional distinctions, counterpossibles all come out vacuously true (see Bjerring 2014).

Last, but not least, we wish to sing the praise of the anonymous heroes of academic publishing. Getting the right referee to review the right paper is a cornerstone of academic editing. We as guest editors were fortunate enough to find a number of highly competent colleagues ready, willing and able to return well-crafted, insightful and constructive reports that always revealed careful reading of and engagement with the papers put before them. Apart from a few desk rejections, each submitted paper was sent out to two referees, and in some cases even three. Some papers went through three rounds of revision. Good referees are known to be in short supply, so this is only an extra reason to thank our referees for a job well

done. Otávio Bueno was the corresponding editor for *Synthese*, always providing the right advice and often within the hour.

REFERENCES

- Bjerring, J.C. (2014), 'On counterpossibles', *Philosophical Studies*, vol. 168, pp. 327-53.
- Carnap, R. (1947), *Meaning and Necessity*, Chicago: Chicago of University Press.
- Church, A. (1954), 'Intensional isomorphism and identity of belief', *Philosophical Studies*, vol. 5, pp. 65-73.
- Cresswell, M.J. (1985), *Structured Meanings*, MIT Press.
- Cresswell, M.J. (1975), 'Hyperintensional logic', *Studia Logica*, vol. 34, pp. 25-38.
- Duží, M. (2012), 'Extensional logic of hyperintensions', in: *Conceptual Modelling and its Theoretical Foundations*, A. Duesterhoeft, M. Klettke and K.-D. Schewe (eds.), *Lecture Notes in Computer Science*, vol. 7260, pp. 268-90.
- Jespersen, B. (2012), 'Recent work on structured meaning and propositional unity', *Philosophy Compass*, vol. 7, pp. 620-30.
- Kaplan, D. (1978), 'Dthat', in: *Syntax and Semantics*, vol. 9, P. Cole (ed.), New York: Academic Press.
- Mates, B. (1952), 'Synonymity', in: *Semantics and the Philosophy of Language*, L. Linsky (ed.), Urbana, IL: University of Illinois Press.
- Nolan, D. (2014), 'Hyperintensional metaphysics', *Philosophical Studies*, vol. 171, pp. 149-60.
- Stalnaker, R. (2002), 'Common ground', *Linguistics and Philosophy*, vol. 25, pp. 701-21.
- Williamson, T. (2013), *Modal Logic as Metaphysics*, Oxford: Oxford University Press.