The 2nd Context, Cognition and Communication Conference
Contexts, Concepts, Objects

BOOK OF ABSTRACTS

Edited by

Tadeusz Ciecierski & Paweł Grabarczyk & Maciej Sendłak

Institute of Philosophy, University of Warsaw
16-19 VI 2018

Organizers:

Sponsors:

.

Contents

POST-CONFERENCE WORKSHOP 185

KEYNOTE SPEAKERS

Robyn Carston
Word Meaning: Concepts, Polysemy and Underspecification

University College London

This talk can be seen as an exercise in trying to bring two fields of language study into a closer relationship: lexical pragmatics (as developed within relevance theory, e.g. Wilson & Carston 2007) and generative grammar. In the terms of Hauser, Chomsky, Fitch (2002), it is a move to integrate one aspect of the broad construal of language, as a communicative system, with the narrow construal of language, as an abstract computational system. The focus is on the meaning of the basic (non-compositional) units of communication and the basic units of the computational system, which, arguably, do not coincide, but must interface.

My starting assumption is that speakers use (substantive) words to express concepts and that what are described as 'occasion-specific senses' of words just are concepts, where concepts are understood in the Fodorian way: as mental particulars, atomic, compositional constituents of thoughts, with referential semantic content (Fodor 1998). Any given word can be used to express a range of distinct but pragmatically related concepts, e.g. the concepts expressed by the verb 'to back' in utterances of the following sentences:

a. I backed the car into the drive.

b. She backed the prints with cardboard.

c. He backed Clinton for president.

d. We've backed all the seats.

It is also widely assumed that a word has a 'standing meaning', a stable component which is activated whenever the word is tokened and which underpins its occasion-specific senses. The question of the nature of these meanings then arises: are they also concepts, that is, semantically contentful entities, or some other, semantically underspecified, meaning component which, perhaps, acts as a constraint on concepts expressible with the word?

A standing word meaning seems not to be a single concept, 'encoded' in the word. Polysemy (multiple related senses), variable across equally competent speakers of a language, is the norm and there is no basis for extracting a single concept as 'the meaning' (Carston 2013, forthcoming). Furthermore, there is widespread cross-categorial polysemy: for instance, the senses of the noun, verb, and adjective 'back' are clearly interrelated. So it is really the root $\sqrt{}$back that is the fundamental unit of polysemy, although the first domain of content (of concept expression) is a categorised root (noun, verb, adjective). Corresponding with the distinction between the broad (communicational) and narrow (computational) construals of language, I suggest there are two kinds of lexicon, whose stored units differ: words and memorised phrasal chunks, in the communicational case; roots, e.g. $\sqrt{}$back, in the computational case. I shall argue that there is no need to postulate a population of non-conceptual non-semantic meanings for words or for roots. The words in the communicational lexicon are associated with a family of concepts, which are pragmatically interderivable, and categoryless roots have no meaning.

References:

Carston, forthcoming. Ad hoc concepts, polysemy and the lexicon.
Fodor, 1998. *Concepts.*
Hauser, Chomsky, Fitch, 2002. The language faculty.
Wilson, Carston. 2007. A unitary approach to lexical pragmatics.

Peter Gärdenfors
Word Meanings as Coordinated Equilibria

University of Lund

Cognitively oriented theories of semantics argue that word meanings are mental representations and not entities in the world (as realist theories claim). A general problem for the cognitive approach is how we can know that we mean the same thing when we use a word. In this lecture I argue that given that our mental structures have certain geometric and topological properties, common meanings can be identified as correlated equilibria in conceptual spaces.

As a background, I build on a result by Jäger and van Rooij (2007). In a game theoretic setting, they have shown that if individuals organise word meanings in conceptual spaces and they are rewarded to the extent their communication succeeds in identifying the same referents, then their mental spaces will contain points representing prototypical word meanings that form Nash equilibria. This result has been generalised by Warglien and Gärdenfors (2013) to the case where the spaces of the individuals are not identical.

These results however, allow a multitude of equilibria as solutions. I argue that in order to pick a unique common word meaning, the communicative context can be used to select a coordinated equilibrium. This process will be illustrated by some examples of meaning coordination.

A particular problem concerns the meaning of words that refer to fictional entities. This is the problem of intentional identity as formulated by Geach (1967). I show that by using coordinated equilibria, new light can be thrown on how the references of fictional identities are determined.

Michael Glanzberg
Precise Meanings from Loose Concepts: Roles for Context and Underdetermination in Semantics

Northwestern University

Cognition allows many forms of imprecision. Our concepts categorize things in highly graded ways, and our words often seem to under-specify meaning. Context impacts thought and meaning in equally imprecise ways, as many contextual parameters seem to be set only roughly. At the same time, linguistic systems, including compositional semantics, like things to be precise. Each expression has an extension, which is sharply delimited. If we have vagueness, we often model it with some precise mechanism, like multiple values, multiple extensions, and so on. In this paper, I explore sources of underdetermination, and argue for a role for precise semantic values. I argue that our linguistic systems can work with precise values where cognition and context provide only underspecified or approximate ones. The precision we see with many semantic values reflects their linguistic roles: expressing linguistically specific content, grammar, and feeding composition. Underdetermination, I propose, comes from the interface with cognition and context.

Edward N. Zalta
Unifying Three Notions of Concepts

Stanford University

In this talk, I first outline three different notions of concepts: one derives from Leibniz, while the other two derive from Frege. The Leibnizian notion is the subject of his "calculus of concepts" (which is really an algebra). One notion of concept from Frege is what we would call a "property", so that when Frege says "x falls under the concept F", we would say "x instantiates F" or "x exemplifies "F". The other notion of concept from Frege is that of the notion of sense, which played various roles within Frege's theory. This notion of concept can be generalized and, as such, accounts for our intuitive talk of "x's concept of ...", where the ellipsis can be filled in with a name for individual, a property, or a relation, etc. After outlining these three notions, I then discuss how (axiomatic) object theory offers a distinct, precise regimentation of each of the three notions.

CONTRIBUTED PAPERS

Paal Antonsen
Logic and Linguistic Evidence

University of Bergen

The aim of this talk is to defend a version the metalinguistic view of logic, according to which logical properties supervene on linguistic properties. The core of the metalinguistic view can be stated as the following claim:

Logic in natural language thesis (LNL). A natural language, as a structure with a syntax and a semantics, thereby determines a logical consequence relation.

Glanzberg (2015) has recently argued against LNL. Central to his case is the *argument from absolute semantics*. It can be summarized as follows:

(P1) A semantic theory determines absolute truth conditions.

(P2) Absolute truth conditions do not determine a logical consequence relation.

(C) A semantic theory does not determine a logical consequence relation.

We argue that the extent to which this is an argument against LNL depends on the details of ones metalinguistic view. In particular, it depends on how one understands "semantic theory" in P1. For example, we agree that Glanzberg's argument is decisive against a Davidsonian version of LNL. On this version, a semantic theory is understood narrowly, consisting solely of a recursive definition of truth. Given that the recursive definition gives absolute truth conditions (not relativized to arbitrary points of evaluation), it does not have the resources to define logical consequence (which requires truth relative to arbitrary points). On the other hand, the argument from absolute semantics is not effective against a Dummettian (1991) version of LNL. On this version, logical properties are said to be determined only by a "full blooded" theory meaning. Following the terminology of MacFarlane (2014) "a semantic theory", in this wide sense, contains three components:

(A) Presemantics. A description of the range of possible semantic values via semantic types. This is done independently of any particular expressions.

(B) Compositional Semantics. A description of how semantic values of complex expressions are determined by the semantic values of their constituent expressions.

(C) Postsemantics: A description of how the 'top level' semantic notions are related to linguistic use. Serves as a bridge theory between semantics and pragmatics.

In P1 and C, "semantic theory" is used in the narrow sense of (B). This leaves it open that LNL may be still be true if logical properties are said to supervene on presemantic or postsemantic properties. We defend the view the first option. Using the Heim & Kratzer (1998) semantics as an example we give a type-theoretic description of the presemantic properties involved. We then show how the presemantic properties are sufficient resources to define a logical consequence relation for a language fragment.

Pauline Armary
Similarity in Context, a Comparison with Tversky and Gärdenfors

Institut Jean Nicod

Similarity is considered by psychologists and philosophers as the most fundamental part of our definition of concepts since the 1970's (Tversky, 1977 ; Rosch & Mervis, 1975 ; Gärdenfors, 2000) even if some objections have been made against its unrestricted influence (Murphy & Medin, 1985) or the relevance of concidering concepts as a unified and coherent aspect of our cognition (Machery, 2009). However, Goodman highlights (Goodman, 1972) some serious flaws in the notion of similarity. He tackles the problem of its context-sensitivity and points out that two things can be similar in too many ways. Thus, similarity is unable to account for the needed and observed stability of our concepts and their uses.

In a previous work, I proposed (Armary, Dokic, Sander, 'The Problem of Context for Similarity in Analogical Reasoning and Categorization', forthcoming) a new definition of similarity based on reflexion on Goodman's objection against similarity (Goodman, 1972). Similarity is seen as the overall importance of the shared aspects of two objects in a given context: two objects a and b are more similar than c and d if:

$$\left(\sum_{i=1}^{n} p^c{}_{x_i}\right) > \left(\sum_{i=1}^{n} q^c{}_{y_i}\right) (1)$$

where p is the weight of a property x, shared by two objects a and b, depending on a context c, q the weight of a property y, shared by two objects c and d, and n the number of properties shared by each pair of objects.

The problem of the attribution of importance in a given context is solved through the reference to the analogical process. The analogical process uses saliences at three level (the sensitive, conceptual and processual levels) to identify the more important features in each situation. I intend to formalize this proposition based on Hofstadter's Copycat algorithm (D. R. Hofstadter & Fluid Analogies Research Group, 1996), Yeh and Barsalou's theory of situated concepts (Yeh & Barsalou, 2006) and Vosniadou's work on saliences (Vosniadou, 1989). My hypothesis is that the various contexts change the

way our mind prioritize some features over others, which entails different judgements of similarity.

In his article (Goodman, 1972), Goodman proposes an example to illustrate the idea of a very high variability of the context. He presents the situation of baggage at an airport and the various point of view that could apply on it (the pilot's, the owner's and the spectator's perspective). In this example, Goodman predicts that each perspective gives importance to different aspects of the baggage (its weight, its destination, its colour) in such a way that they will judge the same baggage similar to different objects.

I propose to compare the definition of similarity as the overall importance of the shared aspects of two objects with Tversky (1977) and Gärdenfors (2000) definitions through Goodman's example.

Arvid Båve
Concept Functionalism and the Nature of Concepts

Stockholm University

I here explore Wayne Davis's proposal that concepts are event types. I argue that this simple (if unobvious) proposal lends itself to a simple yet powerful theory of concepts, which, together with plausible auxiliary assumptions, immediately answers several fundamental questions about concepts. This is part of a larger project about concepts, which I here only introduce.

I assume that concepts are constituents of propositions (the objects of propositional attitudes), that they all belong to some syntactic category or other, and that they are all either syntactically simple or complex. Propositions are concepts of the propositional category. For simplicity, I assume also that the syntactic categories are those of first-order logic.

The theory under consideration holds that:

(1) Concepts are mental event types.

(2) To undergo such a mental event is to *entertain* the relevant concept.

(3) Complex concepts are act-types of *conjoining*, in a certain sense, the concepts immediately involved in the concept in question.

(4) A concept is individuated by its entertaining (i.e., undergoing) conditions, to the effect that it plays a certain inferential (conceptual, functional) role.

I argue that (1)-(4) constitute a very simple theory with considerable scope. Its simplicity is partly due to the fact that (2)-(4) can be inferred from claim (1) plus independently plausible assumptions, and hence are not separate assumptions. While they are thus motivated on the basis of (1), (1) itself can be justified the role it plays in a virtuous theory.

The assumption needed to infer (2) from (1) is (a) that propositions are event types of entertaining, as proposed also by Scott Soames, and that (b) concepts are all the same kind of entity (they are not just all event types, but the same kind of event types).

Claim (3) should be understood in such a way that it follows directly from (1) plus the claim (c) that complex concepts are "built out of" and constituted

by less complex concepts, ultimately simple ones. This is guaranteed, I argue, if we define conjoining as the multigrade relation R such that the event type of entering R with the concepts $x, y, z, \ldots$ in that order = the complex concept whose immediate constituents are $x, y, z, \ldots$ (occurring in that order). With the assumption that concepts are syntactically like the expressions of first-order logic, it is further natural to suppose that the first relatum of conjoining is always an incomplete concept, followed by concepts saturating it. This yields a handy way of referring to complex concepts. Let $E[x, y_1, \ldots, y_n]$ be the event type of entering relation x to entities $y_1, \ldots, y_n$ (in that order), let "C" abbreviate "conjoining", and let us use small caps refer to simple concepts. Now, we have identities like,

> the proposition that John loves Mary and Socrates is wise
> = E[C, AND, E[C, LOVE, JOHN, MARY], E[C, WISE, SOCRATES]]

Claim (4), finally, follows from (1) and (2) together with the plausible assumption that event types can be individuated by the conditions of undergoing them (just as, e.g., properties can be individuated by the conditions of instantiating them).

Iris Berent
Origins of Cognitive Anti-nativism

Northeastern University

Few questions are as controversial as the origins of knowledge. Here, I examine whether this stalemate could be rooted in human cognition itself. Put simply, are people (innately) biased when they reason about cognitive nativism?

Previous research suggests that humans possess core knowledge systems that distinguish minds and matter (*Dualism*), and assign living things an immutable material essence (*Essentialism*). I show these two systems collide in reasoning about cognitive nativism. If, in naïve psychology (a) cognitive traits are immaterial (per *Dualism*), whereas (b) innate traits must be material (per *Essentialism*), then (c) cognitive traits cannot be innate. Results from ten experiments are in line with these hypotheses.

Experiments 1-2 show that people view cognitive traits (e.g., moral reasoning) as less amenable to brain transplant than motor traits (e.g., walking), suggesting that cognitive traits are relatively immaterial.

Experiments 3-5 demonstrate that people view innate traits as more likely to both transfer in brain transplant and correspond to a specific brain region than acquired traits. For example, people are more likely to view a novel cognitive trait (e.g., the communication system of an alien species) as material (e.g., transferable by brain transplant, localized in the brain) when this trait is presented as innate (universal to all members and early emerging) compared to when the same trait is presented as acquired (restricted to certain members, acquired following extensive learning). Accordingly, these results are in line with the hypothesis that innate traits must be material.

If cognitive traits are viewed as immaterial and innate traits as super-material, then cognitive traits cannot be innate. The results of Experiments 6-10 are in line with this possibility. For example, people believe that a linguistic regularity (e.g., the preference for sentences like *dogs bark* over *bark dogs*) is less likely to emerge spontaneously when the trait is attributed to an abstract linguistic constraint compared to when the same trait is attributed to a motor articulatory pressure. Furthermore, the propensity of a trait to emerge spontaneously (i.e., innateness) correlates with its propensity to transfer in brain transplantation (i.e., materiality). Although these correlations do not demonstrate causation, these findings are in line

with our proposal that cognitive traits are less likely to be innate *because* they are viewed as immaterial.

These results are the first to show that adult participants view cognitive traits as less likely to be innate, and that this tendency is linked to reasoning about the materiality of cognitive processes The findings (from adults) cannot ascertain whether these biases are innate. Nonetheless, the correspondence between adult reasoning and principles of core cognition seen in children and infants opens up the possibility that the resistance to cognitive nativism might be rooted deep in human nature—in core knowledge systems that distinguish mind and matter, and assign living things an immutable material essence. Left unchecked, these systematic biases could unduly influence research in brain and cognitive sciences. Accordingly, these conclusions shed new light on human nature and suggest caution in its scientific evaluation.

Sandy Berkovski
Imperatives and To-Do Lists: Three Objections

Bilkent University

Departing from Stalnaker's theory of common ground, a recent pragmatic theory of imperatives aims to identify a concrete pragmatic mechanism through which the speaker directs (commands, requests, instructs) the hearer to perform a particular action.[1] Each participant of the conversation is portrayed as equipped not only with a body of presuppositions out of which the common ground is extracted, but also with 'To-Do Lists'. These lists specify the actions to be performed by the owner. We approach most, if not all, conversations with some idea of what needs to be done, what we have to do. So in uttering an imperative sentence the speaker seeks to update the To-Do list of the addressee. The utterance of the imperative targets the addition of its semantic value to the TDL of the addressee. For specificity, I adopt Portner's proposal that semantic values of imperatives are properties to be uniquely satisfied by the addressee. So, e.g., the semantic value of 'Leave!' addressed to a would be a property construct 'a leaves'[2]

$$[\![!\phi]\!]^C = [\lambda w \lambda x : x \text{ is the addressee} . \ x \text{ is } \phi\text{-ing in } w \ ]$$

Properties thus defined can be satisfied at a world $\mathbf{w}$ by an individual i, in which case we write $P(\mathbf{w}, \mathbf{i}) = 1$. So we can now say that, in accepting the directive, a rational hearer would try to make true as many properties on his TDL as possible. The account just described, to which I will now refer as the 'theory $\mathcal{T}$', claims for itself the virtue of identifying a concrete pragmatic mechanism of imperatives. I argue that this achievement is illusory. There are three reasons for this.

Motivation and compliance. The illocutionary point of commands is the performance of the action, not a change in motivation. In the terminology of $\mathcal{T}$, the point is to saturate the property $[\![!\phi]\!]^C$ in whatever way. The addressee may indicate no acceptance of the imperative, and hence may give us (or the utterer) no reason to think that the TDL was updated accordingly. Yet the subsequent performance of the action can still count as compliance

---

[1]See Portner (2004; 2007) and also Han (2011). See Portner (2016) for a review.
[2]Following the conventions in Heim and Kratzer (1998).

with the command. This last fact is not reflected in $\mathcal{T}$.

Infinite regress. The theory $\mathcal{T}$ makes no distinction between different kinds of speech acts, such as commands, orders, requests, or instructions. This creates a problem. For suppose we grant that the speaker S intends to update the hearer H's TDL $L(H)$. Yet when S commands, he should then *command H* to adjust $L(H)$. If S makes a request for H, then he should *request H* to adjust $L(H)$. The original directive of the form 'Do A!' induces another directive of the form 'Place "doing A" on your To-Do list!' Correspondingly, these further directives should receive the same treatment as the original one. Hence we end up with an infinite sequence of directives to perform ϕ-ing, TDLs, directives to update TDLs, meta-TDLs, directives to update meta-TDLs, and so forth. A regress is under way which apparently cannot be terminated.

The disappearance of directive force. Both complaints just sketched target $\mathcal{T}$'s interpretation of the speaker's intention. The joint response here may be that the goal of $\mathcal{T}$ does not include modelling the speaker's intention at all. It is rather to model the hearer's rationality (but not the actual compliance). But now, a hearer may refuse to accept the command — i.e. to update his TDL accordingly — precisely because it is a command ('I refuse to even consider this — you have no right to boss me around!'). On other occasions, the hearer would do the opposite: 'If he simply asked me, I would have refused; but since he commanded, I feel compelled ... ' The complaint, then, is that, for the hearer's rationality to be interpreted, we have to take into account the directive force of the utterance. Whether the TDL is updated depends on the kind of directive force of the speaker's utterance, a fact ignored by $\mathcal{T}$.

References

Han, C.-H. 2011. "Imperatives." In K. von Heusinger, C. Maienborn, and P. Portner (eds.), *Semantics: An International Handbook of Natural Language Meaning*, volume 2, 1785–1804. De Gruyter.

Heim, I. and Kratzer, A. 1998. *Semantics in Generative Grammar*. Blackwell.

Portner, P. 2004. "The semantics of imperatives within a theory of clause types." *Semantics and Linguistic Theory* 14:235–252.

—. 2007. "Imperatives and modals." *Natural Language Semantics* 15:351–383.

—. 2016. "Imperatives." In M. Aloni and P. Dekker (eds.), *The Cambridge*

Handbook of Formal Semantics, 593–628. Cambridge University Press.
Stalnaker, R. C. 2014. *Context*. Oxford University Press.

Justina Berskyte
Rollercoasters Are not Fun for Mary

The University of Manchester

How should we account for the content of sentences containing Predicates of personal taste (PPTs) such as "fun" or "tasty"? PPTs seem to generate *faultless disagreement*. To illustrate, take Lasersohn's example. Mary and John are riding on a rollercoaster:

(i)   *Mary*: This is not fun.

       John: Oh, yes it is! (Lasersohn, 2005: 651)

It appears that in (i) we have an instance of two propositions expressed which contradict each other, namely 'Rollercoasters are fun' and 'Rollercoasters are not fun', yet both of these appear true. Neither speaker has uttered anything false – thus neither speaker is at fault.

Glanzberg (2007, 2009) argues that we can account for PPTs by introducing an *experiencer class (E)* among the parameters determined by context, modelling his semantic theory on Kennedy's (2007) analysis of gradable adjectives. Thus, according to Glanzberg, disagreement about whether something is fun is actual disagreement about whether something is fun for *E*.

I argue that Glanzberg's introduction of *E* does not sufficiently capture the semantic profile of PPTs. What it manages to account for is what Sundell (2011) has dubbed *contextual disagreement* where individuals can argue about what *should* count as fun or tasty in a particular context. This is not the same as faultless disagreement. My argument is twofold:

(1) PPTs generate faultless disagreement, not contextual (or any other) disagreement. Following Kᷮolbel (2004) and MacFarlane (2007) I argue that we must take the phenomenon of faultless disagreement seriously. Although Glanzberg is not directly denying the existence of faultless disagreement, I argue that he either needs to deny it or show how his theory can account for it.

(2) Glanzberg's introduction of *E* cannot account for faultless disagreement. It predicts that the individuals involved in the relevant discourse will only manage to talk about what the scale for fun or tastiness should be for that

context – rather than disagreeing about their personal tastes. For *E* to account for individual tastes, we would have to have two different *E parameters* within the discourse. This, however, does not account for disagreement as we would be dealing with distinct contexts within the same discourse, hence explaining away the apparent disagreement as mere ambiguity.

I suggest that one of the reasons why Glanzberg's account fails is that he relies too much on the analysis of *observable* gradable adjectives, such as "tall" or "rich". When disagreeing about observable adjectives, we are not in the same epistemically privileged position as we are when disagreeing about tastes. After all, John can't convince Mary that rollercoasters are fun if rollercoasters are not fun for Mary.

Bibliography:

Glanzberg, M., (2007), 'Context, Content, and Relativism', *Philosophical Studies*, 136:1, pp. 1-29.
Glanzberg, M., (2009), 'Not All Contextual Parameters Are Alike' *PRELIMINARY DRAFT*, Available at:
http://faculty.wcas.northwestern.edu/ mjg904/documents/parametersaredifferent.pdf, pp. 1-64.
Kennedy, C., (2007), 'Vagueness and Grammar: The Semantics of Relative and Absolute Gradable Adjectives', *Linguistics and Philosophy*, 30:1, pp. 1-45.
Kölbel, M., (2004), 'Faultless Disagreement', *Proceedings of the Aristotelian Society*, 104, pp. 53-73.
Lasersohn, P., (2005), 'Context Dependence, Disagreement, and Predicates of Personal Taste', *Linguistics and Philosophy*, 28:6, pp. 643-686.
MacFarlane, J., (2007), 'Relativism and Disagreement', *Philosophical Studies*, 132:1, pp. 17-31.
Sundell, T., (2011), 'Disagreements About Taste', *Philosophical Studies*, 155:2, pp. 267-288.

Sharon Berry
Physical Possibility and Determinate Number Theory

Polonsky Academy for Advanced Study in the Humanities and Social Sciences

It's fashionable to take Putnamian model theoretic worries seriously for mathematics. In particular, the question of how we can refer to the standard model of the natural numbers using only first order resources is often considered a serious puzzle. However, similar concerns raised about our ability to refer to physical objects and physical possibility are brushed aside (or at least judged to be already answered). In this paper I will argue this asymmetric treatment and that we can uniquely describe the standard model of the natural numbers (given some reasonable assumptions) using only first order vocabulary and claims about physical necessity.

Corinne Bloch-Mullins & Kristy Nielson
Bringing Context to the Foreground: The Roles of the Foil in Categorization

Marquette University

While similarity-based theories of concepts have a broad intuitive appeal, and have been successful in accounting for categorization based on perceptual and some functional features, it is often argued that their applicability is too narrow for a general theory of concepts, as they cannot account for the formation and application of abstract concepts (e.g., Keil 1986, Rips 1989). Moreover, their adequacy as theories of concepts has been questioned, as similarity is taken as too unconstrained to be explanatory (e.g., Goodman 1972, Medin 1989, Quine 1969).

In this paper, we argue for a richer notion of similarity (drawing, in part, on Tversky 1977 and Gärdenfors 2000), in which context takes center stage. Context—in the form of contrast classes serving as a foil—participates in the determination of the properties taken as relevant for similarity judgments. It therefore constrains such judgments, making them explanatory for the formation and application of concepts *at various levels of abstraction.* Importantly, we hold that context is not merely a modifier of similarity judgments (thereby context is seen as, in a sense, external to the relation captured by these judgments). Rather, context is an integral part of the relations grasped by similarity judgments.

We then operationalize this distinctive notion of similarity, proposing that the manipulation of context in categorization tasks can be used to study the role of similarity judgments in these tasks. We argue that this approach is able to bypass several methodological pitfalls that have led other works to underestimate the role of similarity in categorization (especially when categorization is based on 'deep features' rather than 'surface features'). We report preliminary data from our lab, showing similarity judgments at work in categorization rating tasks not only in perceptual categories, but also in 'natural kinds' categories, whose members share 'deep' causal properties.

We argue that the notion of similarity proposed here is not vulnerable to some of the traditional objections against similarity, and that it has explanatory value for a theory of concepts. The similarity-based account of concepts that we offer should not be seen as a competitor of theory-based accounts, which emphasize the role of causal knowledge in categorization. Rather,

similarity judgments should be understood as the apparatus that facilitates the application of causal knowledge to particular targets in categorization. Therefore, the present version of the similarity-based approach can deliver a unified theory of concepts that could be applied to a broad range of concepts.

Gärdenfors (2000) *Conceptual Spaces: The Geometry of Thought*, MIT Press.

Goodman (1972) 'Seven strictures on similarity', in Goodman (ed.), *Problems and Projects*, Bobbs-Merrill., pp. 437-47.

Keil (1986) 'The acquisition of natural kind and artifact terms', in Dimopoulos and Marras (eds), *Language Learning and Concept Acquisition: Foundational Issues*, Ablex., pp. 133-53.

Medin (1989) 'Concepts and conceptual structure', *American Psychologist*, pp. 1469-81.

Quine (1969) 'Natural kinds' 'Natural kinds', *Ontological relativity and other essays*, Columbia University Press, pp. 114-38.

Rips (1989) 'Similarity, Typicality, and Categorization', in Vosniadou and Ortony (eds), *Similarity and Analogical Reasoning*, Cambridge University Press, pp. 21-59.

Tversky (1977) 'Features of similarity', *Psychological Reviews*, pp. 327-52.

Krystian Bogucki
Context Dependence and the Principle of Compositionality

University of Warsaw

The Principle of Compositionality is the principle that the meaning of an expression is a function of (and only of) the meanings of its parts together with the method by which those parts are combined. In a compositional language, it seems, the meaning of an expression depends upon the meanings of its parts in a bottom-up fashion, so some philosophers (e.g. Fodor & Lepore) argued that compositionality places a context-independence constraint on the properties of lexical meanings. However, some others philosphers (e.g. Szabo, Recanati) haven't seen context-dependence as a challenge to compositionality. The idea goes as follows: the fact that natural languages contain indexicals forces us to distinguish between two notions of meaning: standing meaning and occasion meaning. The standing meaning is the meaning which the word has in isolation, in virtue of the convention of the language. The occassion meaning is the meaning which an occurrence of the word takes on in a particular lingustic context. Context-dependent expressions are expresions whose occassion meaning sometimes deviates from their standing meaning. An effective argument for incoherence context-dependence and the principle of compositionality has to show that there is at least one complex expression in L whose occasion meaning varies with context, while the occasion meanings of its constituents all remain the same, but the usual considerations against compositionality typically omit the second part. This strategy (proposed by Z. Szabo) take it for granted that the context-dependence of a complex expression ultimately derives from the contextdependence of its lexical items. F. Recanati refuted this assumption and argued that alternatively we can achieve a weak form compositionality: let's the composistion function take the context as extra argument: we can say that the content of a complex expression A*B is a function of the contents of its parts and the way they are put together plus the context which provides for the modulation of theses contents. Recanati's solution suggests that the content of complex expressions is not compositionally determined and what we really need to account for e.g. the systemacity of language is not the principle of the compositionality, but only some recursive mechanism underlying lingustic

competence. F. J. Pelletier proposed 'The principle of semantic groundedness' as such a recursive alternative for the principle of compositionality.

In my presentation I'm going to: i) evaluate Szabo's and Recanati's strategies to deal with context-dependent expressions in a compositional framework; ii) answer to the following question: Can we account for such phenomena as the systemacity of language due to alternative conception of semantics, one that relies on some recursive mechanism (e.g. groundedness) rather than compositionality?

Mark Bowker
Generics Without Truth-Conditions

University of St Andrews

Standard accounts of generics like 'Immigrants are dangerous' analyse them as structurally identical to explicitly quantified generalisations like 'Most immigrants are dangerous', with the role of the quantifier 'Most' played by the unpronounced operator 'GEN'. Debate then focuses on how 'GEN' determines the truth-conditions of a generic as a function of the pronounced lexical material. There are two problems with this approach. First, speakers often disagree about the truth-conditions of generics. This was powerfully illustrated by Donald Trump's description of Mexican immigrants during the USA Presidential Campaign: "They're rapists". Some saw Trump as falsely claiming that all Mexican immigrants are rapists. Others defended Trump by suggesting that some significant proportion are. The truth-conditional approach must begin analysis of such utterances by favouring one party in this disagreement over another. Second, no account of 'GEN' has so far managed to accommodate even all widely-shared intuitions about the truth and falsity of generics.

This paper will adopt an alternative approach. Generics are analysed as incomplete structures that must be assigned a quantifier to determine truth-conditions. As contexts are regularly consistent with different assignments, generics are consistent with different truth-conditional interpretations. This account treats the alternative interpretations of Trump's utterance on a par. Both are consistent with Trump's words and the context in which he spoke. In this sense, both are correct. But neither is *the* correct interpretation. In this sense, both are mistaken.

Despite the absence of truth-conditions, utterances of generics may have unique truth-values. Interpretations of generics are contextually-restricted to those that answer the question under discussion (QUD). Following Stalnaker,[1] I represent the common ground of a conversation by a set of worlds and, following Groenekdijk and Stofhoff,[2] represent the QUD as a partition on

[1]Stalnaker, R. 1978. Assertion. In Peter Cole (ed.), *Syntax and Semantics* 9, 315–332. New York: Academic Press.

[2]Groenendijk, J. & Stokhof, M. 1984. On the Semantics of Questions and the Pragmatics of Answers. In *Studies on the semantics of questions and the pragmatics of answers.* Amsterdam: Institute for Logic, Language & Computation (ILLC), University

the set. An interpretation of an utterance is *pertinent* only if it provides an answer to the QUD. I assume that interpretation is driven by the search for pertinent interpretations.[3] Suppose that the QUD is 'Should Mexicans be allowed to immigrate to the USA?' and that it is common ground that the answer is 'No' if more than half are rapists. The existential interpretation of Trump's utterance is not pertinent to this question, given that it is consistent with both answer. Both the universal and majority interpretations are pertinent and each provides a negative answer to the QUD. While its truth-conditions are underdetermined, Trump's utterance provides a negative answer to the QUD. The truth-value of the utterance depends on the truth-value of this answer.

---

of Amsterdam PhD thesis.

    [3]It remains to be seen what relation *pertinence* bares to the relevance-theoretic notion of *relevance.*

Adrian Briciu
Intentions and Improper Contexts

West University of Timisoara

Recording devices are generally taken to present problems for the standard Kaplanian semantics for indexicals. In the Kaplanian semantics every utterance of (1) "I am here now" comes out true; and given how negation works it follows that every utterance of (2) "I am not here now" comes out false. But the pervasive use of recorded messages seems to indicate that there are true uses of (2). For example if we call someone's house and (2) is played back on the answering machine, we judge it to be true if that person is not home at the time of the call.

In my talk, I'll argue that the remote utterance view (the view that recording devices allow agents to perform utterances at a distance) offers the best way for a Kaplanian semantics to handle the recalcitrant data that comes from the use of recording devices. I'll make us use the between tokens and utterances (Searle 1978, Bromberger 1992, Perry 2001) to show that recording devices can be used in two ways (a) to record our utterances, (b) to perform utterances at a distance. I develop the view beyond the initial sketch given by Sidelle (1998), and show that it gives correct predictions for all the data coming from the use of recording devices, including data (provided by Predelli 1998, 2011) which seemed problematic for the initial remote utterance view proposed by Sidelle.

I'll discus two consequences of the view. One is that it requires that we give up the (Kaplanian) restriction that semantic evaluation always takes place at proper contexts (contexts at which the agent of the context is at the time, place and world of the context), and thus a reworking of Kaplan's logic of indexicals. I'll show that there are additional reasons to give up this restriction. Building on recent work by Stojanovic (2011) and Radulescu (2015) I'll show that besides the use of recording machines in conversation there are good, independent reasons for a revision of Kaplan's logic of indexicals, one in which improper contexts are available for semantic evaluation. Secondly, I'll show that contrary to initial appearances the remote utterance view is not a an intentionalist account - i.e. one according to which the reference of indexicals is partly determined by speakers' intentions. According to the remote utterance view, speakers' intentions play no role in the semantic evaluation of indexical sentences. They do play a role, though,

in determining the utterance produced, namely they determine what sentence is to be uttered, and determine when and where to utter it. And this is what we should expect since utterances are purposeful acts which execute communicative intentions. But intentions play no role in the semantics of indexicals: they do not determine, from one occasion of use to another, the character that governs that particular use as in Smith (1989), nor do they determine the context (which can be distinct from the context of use) at which indexicals are to be semantically evaluated as in Predelli (1998, 2011) or Recanati (2005).

Christopher Brown
Epistemic Contextualism Entails That We Know Nothing

City University of New York

Epistemic contextualism has gained support as a way to avoid skeptical worries without addressing skeptical challenges head-on. In its most popular form, it says that epistemic standards are context-dependent, and that the truth of a knowledge ascription depends on the context that the ascriber is in. If the ascriber has high epistemic standards, then knowledge ascriptions are less likely to be true. For instance, when a person is told what time a bus will leave, her knowledge ascription depends on the stakes at risk if she were to miss her bus—if the stakes are low, as when she is commuting home from work, but can easily take a slightly later bus, then she is in a low-standards context. Even if the ascribee has only weak justification for her belief about the bus departure time, the belief nonetheless counts as knowledge. On the other hand, if the ascriber must get to a dying lover's bedside before the curtain call, then the weak justification of the ascribee will not suffice for knowledge.

Under this framework, it is apparently correct to ascribe knowledge in everyday contexts, despite our inability to come up with a philosophically satisfying answer to radical skepticism. This is because it seems that we only enter contexts in which skepticism is relevant when we consciously consider skeptical scenarios, and thus raise our epistemic standards to unreachable heights. If this is right, then knowledge ascriptions are only affected by skeptical worries when we are doing philosophy—leaving everyday knowing safe and sound. However, this strategy to side-step skepticism has a serious flaw: given a rigorous analysis of context formation, it turns out that skeptical worries force themselves into our everyday contexts, and we are again left knowing nothing. If this is correct, then epistemic contextualism dooms everyday knowledge, rather than saving it.

My reasoning: sceptical scenarios are always highly relevant—whether we realize it or not—to our everyday interests. In the aforementioned bus example, the truth of a skeptical scenario would make any sort of bus travel impossible—in reality, there would be no bus to travel on, and no home or dying lover at the destination. Put another way, the possibility of a skeptical scenario raises stakes in all contexts impossibly high: normal humans are

always highly interested in not being brains in vats or deceived by evil demons, regardless of what activities we are engaging in. The epistemic contextualist cannot respond by saying "well, the ascriber doesn't know that skeptical concerns are relevant to her at that time, so such concerns are not context-forming". Such a response is obviously viciously circular: knowledge would depend on contexts, but contexts would also depend on knowledge. No good!

My argument depends on a distinction between occurrent and non-occurrent interests, which is like the distinction between occurrent and non-occurrent beliefs. I flesh out this distinction, and show that denying the context-determining power of non-occurrent interests—like our anti-skeptical interests—has a number of unfavorable consequences. If successful, this robs epistemic contextualism of its standard answer to skeptical challenges.

Alyona Budnikova
Lexical Meaning as a Fluctuating Structure

Minsk State Linguistic University

The view on word meaning as a mental construct highly sensitive to the context of its use and as a faceted structure with different facets profiled in different contexts (Lakoff 1987; Langacker 1987; Pustejovsky 1991; Croft & Cruse 2004; Kharitonchik 2009, etc.) poses a research question of how to model this variable and open-ended structure on the systemic level.

A case study of some common English nouns like *cat*, *tree*, *fish*, *bird*, *man*, *woman*, *teacher*, etc. reveals that even though there is a relatively clear conceptual link between these words and the designated entities (esp. compared to other word classes), their meanings vary from definition to definition, and from one use in context to another. For example, one definition of the word *bird* relies on perceptual ('appearance', 'size', etc.) and locative ('ability to move by air') features of the referent; another will stress perceptual, locative and biological properties ('reproduction'); still others will profile perceptual, biological and systematizing type of knowledge about the referent ('taxonomic name', 'size of population'), etc. The contextual use of the word *bird* shows even greater degree of fluctuation of its meaning structure with a wider variety of profiled facets. For instance, in the context *to buy a bird* the utilitarian facet ('commodity') of the concept BIRD is realized; in *the bird's skull* it's the constitutive facet ('internal body parts'), etc.

The foregrounding of some semantic features and backgrounding of others in word meaning is also noticeable when different metaphorical derivatives of the word are motivated by different features of the source concept. For example, the derivative *bird of passage* ('a person who moves from place to place frequently') highlights the locative facet of the source concept BIRD, and *bird-brained* ('a stupid person') – its behavioral facet.

In the model of word meaning suggested in the paper I try to demonstrate its variability on the systemic level. There is no rigid dividing line between core and peripheral components; rather the classifier is surrounded by features of different character (e.g. of locative, perceptual, behavioral, temporal, etc. facets) and different weight (i.e. salience judged by the frequency of activation) according to the taken perspective (definitional needs, contextual dependency, and derivational potential). The model explains how open-endedness and relative stability of lexical meaning coincide in

its structure and how such factors as level of abstraction, animateness, etc. influence the organization of semantic features in word meaning.

References

1. Croft, W. and Cruse, D.A. (2004). *Cognitive linguistics*. Cambridge University Press.

2. Kharitonchik, Z.A. (2009). Lexical meaning as a dynamic entity of variant semantic components. *In Probleme Actuale de Lingvistică, Glotodidacticăşi Ştinţă Literară: actele colocviului internaţional IV/I*, pp. 113-121.

3. Lakoff, G. (1987). *Women, fire, and dangerous things. What categories reveal about the mind*. University of Chicago Press.

4. Langacker, R. W. (1987). *Foundations of cognitive grammar: Theoretical prerequisites*. Stanford University Press.

5. Pustejovsky, J. (1991). The generative lexicon. *Computational Linguistics*. 17 (4): 409–441.

Philip Clapson
The Theory of Brain-Sign

Birkbeck, University of London

It is usually supposed that the source of the conscious mind is the brain. However it has so far proved impossible to link brain functioning with the mind in any specific sense. Yet it appears the mind exists, capable of sustaining contexts, concepts and abstract objects as conscious states. Science is impeded by this unsolved ontological problem.

The biophysical *theory of brain-sign*, however, does provide a scientific link between the brain and what has been supposed as consciousness. Brain-sign is not consciousness. It is a physical sign in the brain which acts as a communication medium/mechanism enabling joint activity in the world with other brains. Brain-sign is derived from the brain's causal orientation towards the world at each moment. Joint action between organisms requires cooperative *inter*-action as if they were one. Obviously they cannot be, but by each brain-signing the world of their causal orientations they establish the 'common world' of their interaction.

As an example, if A passes B a cup of coffee, each physical brain controls the action of each body by its neural causal states. The 'image' of the cup (the *concrete object*) in both brains derives from the causal orientations of each brain toward the cup, and thus identifies the world-target of the action. What is also signified as brain-sign is the environment – our arms and the walls of the room, etc. –, and the verbal agreement that coffee will be exchanged, and the desire for the beverage, and the ability to carry it out. In this way the *context* of the action, including what *this* cup is for, establishes the activity's meaning. But since brain-sign is not consciousness, 'we' do not know the context, nor do we know its meaning. Brain-sign is a neurobiological condition which facilitates the interaction of our bodies as purely physical organisms.

The [concept of, or *abstract object* – cup – are linguistic conditions; but brain-sign language retains the physicalist approach. Language functions to alter the causal orientation of another brain via compression waves or electromagnetic radiation. The '*sense*' that we understand words and sentences in which concepts and abstract objects are carried, is actually a communicative status with another brain whether present or not. If A asks B if she wants a cup of coffee, the coincidence of 'cup', as brain-sign,

represents our brains' mutual causal orientation, which includes 'cup' as abstract object *and* concrete object. If A is reading about coffee production, he is brain-singing the communication of author B at a different time and place; so there is a brain-sign correspondence, but not simultaneity.

'Concepts' derive from neural assemblies allowing brains to communicate about the nature of objects in the world, and their structure and functioning. Built in childhood onwards, they are interpreted by brains as brain-signs for communication. There is no conscious thinking about the world for there are no mental states. How contexts, concepts and objects occur, therefore, requires an adequate neuroscience including the operation of brain-sign. This requires major cross-discipline development, with new concepts and vocabulary.

José Pedro Correia
Do We Need a Theory of Concepts? Signaling games, Wittgenstein, and the case of color categorization

University of Porto

Concepts are usually seen as useful notions to give accounts of various aspects of human cognition. Margolis and Laurence go as far as stating that it is "relatively uncontroversial" that concepts "are crucial to such psychological processes as categorization, inference, memory, learning, and decision-making" (2014). Given this assumption, it can seem desirable to develop a deeper understanding of what concepts are and how they function, i.e. to develop a theory of concepts. I want to argue that there can be ways to understand the aforementioned processes without using the notion of concepts, and thus to do away with the need for a theory of concepts altogether.

I focus on one example: the application of Lewis-Skyrms signaling games (Lewis 1969; Skyrms 2010) to understand the task of color categorization. The World Color Survey (Kay et al. 2009) has collected a large amount of data on the assignment of color terms to each of 330 different color stimuli by an average of 24 native speakers of each of 110 unwritten languages. I discuss a signaling game model that is able to approximate important characteristics of this empirical data without any appeal to concepts. By modeling the task as involving both a speaker and a hearer, one can use evolutionary dynamics and the notion of equilibrium to explain how a useful system of color categorization can evolve and be sustained. Although such a system might look as if driven concepts, concepts are not needed to characterize it. Therefore, a theory of concepts is not necessary to explain this particular categorization process.

I believe this example provides additional motivation for a perspective on how to develop a deeper understanding of language and mental processes that is in line with Ludwig Wittgenstein's later philosophy (1953). I argue that some of Wittgenstein's remarks serve as reminders to refrain from pursuing any sort of theory of concepts. Rather, one should instead focus on developing ways to understand processes like categorization by looking into patterns of use. I argue that the framework of signaling games is one tool among others that allow us to do exactly that. Color categorization serves as a particular example of how to do so, but if one can apply the

same approach to other cognitive processes, and concepts turn out to also be unnecessary to explain them, perhaps a theory of concepts is not needed after all.

References

Lewis, David. Convention: A Philosophical Study. Blackwell, 1969.

Kay, Paul; Berlin, Maffi, Merrifield, Cook. The World Color Survey. Stanford: CSLI, 2009.

Margolis, Eric and Laurence, Stephen, "Concepts", The Stanford Encyclopedia of Philosophy (Spring 2014 Edition), Edward N. Zalta (ed.), 2014.

Skyrms, Brian. Signals: Evolution, Learning, and Information. Oxford University Press, 2010.

Wittgenstein, Ludwig. Philosophical Investigations. Edited by P. M. S. Hacker and Joachim Schulte, Revised 4th edition, Wiley-Blackwell, 1953/2010.

Sean Crawford
Falling out of Thinking

The University of Manchester

Demonstrative thoughts are context-bound perception-based thoughts containing demonstrative concepts, such as That is a cactus. The Descriptive Object-Independent Theory (DOIT) of Schiffer (1978), Searle (1983), et. al. analyzes the content of such thoughts in descriptive terms, → *la* Russell, thereby securing content in the "empty case" where no object is singled out, owing to the thinker's suffering from a referential illusion or hallucination. The Object-Dependent Theory (ODT) of Evans (1982) and McDowell (1984, 1986) maintains that demonstrative content is irreducibly singular-demonstrative and that this means that in the empty case no genuine demonstrative content can exist. The Predicative Object-Independent Theory (POIT) of Burge (1977, 2005, 2009) rejects this as counter-intuitive, sharing with DOIT the view that demonstrative content can exist even in the absence of any referent, but maintaining, with ODT, that this content is still irreducibly singular-demonstrative.

Advocates of ODT and POIT criticize DOIT on the grounds that it dramatically and implausibly over-intellectualizes demonstrative thought content. This paper takes off from the debate between ODT and POIT and presents an argument in favour of ODT's claim that in the empty case the subject is not having a genuine demonstrative thought.

All three theories agree that propositional representational content is constitutively truth conditional. The basic argument in favour of ODT is that when there is a failure of demonstrative reference, no truth conditions for the overall thought exist. This is because thoughts containing context-bound demonstrative concepts have *conditional* truth conditions of the (rough) schematic form: 'If a speaker refers to x with "that" in an utterance of "That is F", then "That is F" is true iff x is F'. Since there is no successful contextual-demonstrative reference in the antecedent, the consequent cannot be detached in order to derive any truth conditions. Given the assumption that propositional representational content is constitutively truth conditional, it follows that there is no demonstrative representational content in the empty case. There is no way the world is represented as being by the subject's attempted demonstrative thought. So he is not having a demonstrative

thought with propositional representational content (whatever else might be going on cognitively).

This result seems to have been disguised by the fact that POIT does give the (attempted) thought a semantic evaluation. The empty thought has no *truth* conditions and so is not true. Nevertheless, as Burge (1983) and Braun (1995), and perhaps Kripke (2011, 2013), claim, it plausibly has *falsity conditions* that include the case where it is not true. So for POIT, the empty thought *That is F* is false by dint of being not true (so bivalence is preserved). This general semantical strategy is of a piece with that famously pioneered by Donnellan (1974) for negative existentials. Donnellan made absolutely clear, however, that the strategy crucially divorces semantic evaluation (truth or falsity conditions) from meaning or content and therefore does not provide an analysis of content. Unlike POIT, ODT appreciates the full significance of this dissociation, correctly maintaining that, despite the possibility of a semantic evaluation, an empty demonstrative thought remains devoid of propositional representational content and is in that respect not a genuine demonstrative thought.

Matej Drobňák
Communication in Context: an Inferentialists Approach

University of Hradec Králové

According to the standard view, the content ascribed to utterances often deviates from the standard meaning of sentences. This distinction is well-known as a difference between speakers meaning and linguistic/standard meaning. The deviation is considered to be a result of context dependence of sentence content – the standard meaning is modulated by contextual cues (present in a conversation).

The aim of this talk is to discuss how the framework of normative inferentialism (Brandom 1994, 2000; Peregrin 2006, 2014), could be used to explain and represent the difference between standard meaning and a content of sentence as uttered in context. For this purpose, I will focus on the distinction between the inferential potential (IP) and the inferential significance (IS) of a sentence.

IP of a sentence A can be understood as a set of sets of sentences which can be inferred from A and other premises. In some way, it represents the meaning of a sentence. Moreover, inferentialists believe that the meaning of a sentence is perspectival; i.e. one sentence may have different significance in different contexts. As Peregrin claims, the relation between the context invariant IP and the context dependent IS is straightforward: "the inferential significance of A within the context C is the value of the inferential potential of A for C" (Peregrin, 2014, 51).

One way how to represent the meaning/IP of a sentence A is by a set of sets of sentences:

$$IP(A) = \{\langle X_1, D_1\rangle, \langle X_2, D_2\rangle, \ldots, \langle X_n, D_n\rangle\}$$

where $X_1 = \{M_1, ..., M_n\}$ and D_1 is a sentence which can be inferred from $X_1 + A$; $X_2 = \{N_1, \ldots, N_n\}$ and D_2 is a sentence which can be inferred from $X_2 + A$; $X_n = \{O_1, \ldots, O_n\}$ and D_n is a sentence which can be inferred from $X_n + A$.

According to this approach, context is inherently part of the representation of meaning – instantiated by collateral premises which are members of sets $X_1, X_2 \ldots X_n$. If IS is the value of IP for C, then to understand a sentence as uttered means to eliminate those sets from IP which

do not include premises relevant in particular conversation. For example, if $IP(A) = \{\langle X_1, D_1 \rangle, \langle X_2, D_2 \rangle, \ldots, \langle X_n, D_n \rangle\}$ and a hearer knows that premises from X_1 apply in the conversation and (in the ideal case) none of the premises from X_2 to X_n apply then $IS(A) = \{\langle X_1, D_1 \rangle\}$.

If we assume that sets of premises $X_1, X_2 \ldots X_n$ are socially well-established and well-known among members of communities, then we found an elegant way how to incorporate contextual dependence into an inferentialists framework. According to this approach, the meaning of a sentence includes all contextual values. A contextual modulation is not seen as a deviation, but as a specification of meaning. To fully understand a sentence in general requires knowing all its possible significances. To understand a sentence as uttered requires singling out an appropriate subset of significances from its inferential potential.

Neil Hamilton Fairley
The De Se and Computation

University of Reading

There has been debate between those who maintain that indexical expressions, "I" in particular, are not essential (e.g. Cappelen and Dever 2013) and those who maintain that such indexicals cannot be dispensed with without an important loss of content (Ninan 2016, Recanati 2012). This seems to entail that thoughts must also have indexical elements. Otherwise the indexical would not be necessary to that propositions content. In which case, it seems that we must have some degree of context sensitivity in thought. But context sensitivity in thought appears to be in tension with the computational theory of mind (CTM) (Fodor 2010). In this case we have an apparently inconsistent triad between the following three claims:

(i) De Se thoughts are essential.

(ii) De Se thoughts are indexical, they have a (Kaplanian) character.

(iii) Computations can only take the syntactic *type* into account, they cannot take tokens into account (Ball 2010).

(i) and (ii) together seem to imply that we can have thoughts which are context sensitive and that these thoughts cannot be replaced without a loss of content. The problem comes with claim (iii). If this is correct then we cannot make sense of a thought which uses a character such that its referent could vary from tokening to tokening. This leaves us with a choice of which of these claims to reject.

Rejecting claim (i) seems to be a move towards the inessential indexical view. Rejecting (ii) might leave us with a view akin to Millikan's (1990) on which de se thoughts are present, but make use of non-indexical components. Rejecting either (i) or (ii) will be a move against the essential de se. Rejecting (iii), on the other hand, might require us either to reject the CTM or to explain how computations can be sensitive to character. This paper argues for the second disjunct, showing how computations might be taken to be sensitive to character.

In section 1 I argue that (i)-(iii) are inconsistent, as they commit us to thinking that we both can and cannot have thoughts with a character. In sections 2 – 4 I say something in favour of each of claims (i) - (iii). I defend

(i) on the basis of Ninan's recent defence of essential indexicalism and then argue that (ii) follows from (i). Finally, I suggest that (iii) is plausible by making use of Ball's (2010) argument.

In section 5 I attempt to show how (iii) can be rejected, whilst maintaining the CTM. I claim that computations need not be sensitive to the features of a tokened symbol in the way that character demands. This job may be passed along to a non-modular part of the mind (similar to Borg 2012). Here, (ii) must be respected. So I will also argue that thoughts can be underderdetermined with regard to their content. I suggest that resolving the inconsistent triad in this way provides us with reason to accept the essential indexical view, in which case we should accept some degree of context sensitivity in thought.

Giulia Felappi
Truthmaker Semantics and the Argument from Validity

University of Southampton

0.

Central to truthmaker semantics (TS) is the concept of verification or truth-making: something on the side of the world - a fact or a state of affairs - verifies, or makes true, something on the side of language or thought - a statement or a proposition. One of the alleged advantages of truthmaker semantics is that it allows us to shape the notion of validity in a better way: while validity is usually merely spelled out in terms of necessary preservation of truth, truthmaker semantics allows for two notions of validity, one in terms of preservation of truth, the other in terms of containment of the conclusion in the premises.

1.

The aim of this paper is to show, first, that there are good reasons to think that the notion of validity as containment is unnecessary. I will discuss in turn each reason that defenders of the notion of validity as containment have put forward in support of it: incoveniences; analyticity; imperatives; expressives. I will show that none is compelling, either because they rely on theses that are too controversial, or because a different explanation of the data is immediately available.

2.

Secondly, I will argue that the two notions of validity TS is able to shape are jointly insufficient. I will put forward two cases: propositional attitude sentences and some selfreferential 0-premise arguments. We will see that TS, in having notions of validity that concern truth or content *on the side of the world*, is unable to shape validity for these cases.

Conclusion.

While suggesting 'the replacement of the notion of 'being true" by that of "being made true" (1969: 484–485), and thus laying the basis for the development of TS, van Frassen himself noted that truthmakers would only be interesting if they served to define interesting relations, and then held

that we did have an interesting relation truthmakers were able to shape in logic, i.e. *entailment*. Since then, the idea has been that validity was a point in favour of TS. But, as we saw in §1, there are reasons to think that the notions of validity such a semantics can provide us with are unnecessary and, as we saw in §2, there are also reasons to think that they are insufficient. There might be reasons why we should embrace the TS project, but, if I am right, validity cannot be taken as one of them.

References

Fine, Kit forthcoming, 'A Theory of Truthmaker Content I: Conjunction, Disjunction and Negation', in *Journal of Philosophical Logic*

Fine, Kit manuscript, 'Compliance and Command I - Categorical Imperatives'

van Fraassen, C. Bas 1969, 'Facts and tautological entailment', in *Journal of Philosophy* 66

Williamson, Timothy 2009, 'Reference, Inference and the Semantics of Pejoratives', in J. Almog, P. Leonardi (eds.), *The Philosophy of David Kaplan* (Oxford: Oxford University Press)

Sarah Fisher
Framing as Presupposition-Denial

University of Reading

Framing studies reveal distinct effects of using linguistic expressions that are standardly considered equivalent to one another.[1] For example, if *exactly 80% of people survive* some disease or disaster then *exactly 20% die* from it. Yet the very same outcome is typically viewed more favourably under the 'survive' frame than the 'die' frame.

One 'information leakage' account explains framing effects in terms of pragmatically-conveyed information regarding a reference point.[2] For example, a speaker's use of the 'survive' frame may communicate that fewer people were originally expected to survive (more were expected to *die*). In contrast, the 'die' frame may communicate that fewer were expected to *die* (more were expected to *survive*). Assuming this comparative information guides evaluative judgements, we can explain the 'shifty' preferences observed across framing conditions: if more people survived than expected, the outcome is *relatively good*; if more died than expected, it is *relatively bad*. Thus, the same outcome is evaluated differently, in relation to diverging reference points.

A 'presupposition-denial' account has been developed to explain similar effects associated with positive and negative natural language quantifiers.[3] The use of a positive quantifier like 'a few', as in 'a few people attended' may implicitly presuppose a prior expectation of there being fewer *attendees* (more *absentees*). Meanwhile, the use of a (quantitatively equivalent) negative quantifier like 'few', as in 'few people attended' may implicitly presuppose an expectation of fewer absentees (more attendees). This helps explain shifts in focus that are typically observed, switching from the set of *attendees* in the first ('a few') case, to the set of *absentees* in the second ('few') case. In general, the focus shifts to whichever set is *abundant* in relation to the

---

[1]Beginning with Tversky, A., & Kahneman, D. (1981). The Framing of Decisions and the Psychology of Choice. *Science*, 211(4481), 453-458.

[2]Sher, S., & McKenzie, C. R. M. (2006). Information leakage from logically equivalent frames. *Cognition*, 101(3), 467-494.

[3]Moxey, L. M. (2006). Effects of what is expected on the focussing properties of quantifiers: A test of the presupposition-denial account. *Journal of Memory and Language*, 55, 422–439.

presupposed, expected quantity (and in this way *denies* that expectation). It is simply that 'a few' and 'few' typically presuppose *different, opposing* expected quantities.

Extending this account to framing, we can hypothesise that speakers' uses of frames may presuppose – and deny – divergent expectations. For example, saying that 'exactly 80% survive' may presuppose an expectation of *fewer than 80%* surviving. Denying this expectation by surpassing it, the 'survive' frame renders the outcome relatively good (good relative to what was previously expected). In contrast, saying that 'exactly 20% die' may presuppose an expectation of *fewer than 20%* dying, rendering the outcome *relatively bad.*

I will argue that presupposition-denial through framing enables the subtle addition of expectation information to the conversational common ground.[4] This has potentially widespread and far-reaching social epistemological implications, including for our ability to identify and dispute the presupposed information.

---

[4]Stalnaker, R. C. (1999). *Context and Content.* Oxford: Oxford University Press.

Mara Floris

Language and Conceptual Development: Effects of Labelling on Perception and Cognition

Università degli Studi di Torino

It has been argued, both on the philosophical side and on the psychological one, that abstract concepts may be grounded on concrete concepts and that concrete concepts, in turn, are grounded on perception. A similar distinction has been draw, in philosophy of language, between two aspects of lexical semantic competence: inferential and referential competence (Marconi 1997). The inferential competence is the ability to deal with the network of semantic relations among lexical units. The referential competence is the ability to identify referents of words: either objects in the world or mental images. It is therefore crucial, both for semantic and conceptual knowledge, to be able to categorize the objects we see.

I will investigate the interactions between language and the acquisition of concepts. If we follow the suggestions of the Sapir-Whorf Hypothesis, we may think that the language we speak affects our perception or that, at least, it affects our conceptual organization. It has been proved that very young infants (prelinguistic) are able to form categories on the basis of the perceptual features of objects and it has been claimed that language shapes those categories.

As Althaus and Westermann (2016) point out, two main questions have emerged: whether language enables infants to form categories they would not form without labels and whether labels can override visual perceptual features and change the structure of categories despite the perceptual similarity. Both the possibilities have been explored: when the same noun is applied to a set of visually dissimilar objects infants (and adults as well) are more willing to consider them as members of the same category (e.i. Althaus & Plunkett, 2016; Lupyan, Rakison & McClelland, 2007; Plunkett, Hu, Cohen, 2008), conversely, when different nouns are applied to a set of distinct but similar objects the differences are stressed and the visual stimuli are segregated in different categories (e.i. Dewar & Xu, 2009; Landau & Shipley, 2001; Waxman & Braun, 2005; Xu, Cote & Baker, 2005).

At least fve alternative accounts have been proposed to explain the effects of labels on categorization: the "auditory overshadowing" hypothesis (Robinson & Sloutsky, 2007), the "natural pedagogy" (Csibra & Gergely,

2009), the "referential" (Ferguson & Waxman, 2016; Waxman & Gelman, 2009), the "labels-as-features" view (Gliozzi, Mayor, Hu & Plunkett, 2008; Sloutsky 2010, Sloutsky & Fisher 2012) and the "Labels Feedback Hypothesis" (Lupyan, 2012). In the light of the empirical data I will evaluate the theoretical options and I will argue that there are good reasons to reject the "auditory overshadowing" hypothesis, the "natural pedagogy" view and the "referential" one. Given the evidence it is not possible to reject the "labels-as-features" view and the "Labels Feedback Hypothesis". According to both these accounts there is an impact of labels on visual concepts, either because labels have the weight of other visual features or because there is a penetration of language on early-level perception. Whatever is the case, language plays a role in increasing, or decreasing the similarity between concepts.

Antonella Foligno
Abstractions and Idealizations as Objects for Epistemic Representation

University of Urbino

In mechanics we investigate the velocity of a moving body or the length of the path it covering by neglecting or omitting a series of quantitative or qualitative properties, because we deem them superfluous to the purpose at hand. Population biologists study the evolution of a species which procreates at a constant rate by considering it as an isolated ecosystem. These descriptions are descriptions of a modeled system, and scientists use models to represent aspects of the real world's system.

Models are of central importance in many scientific contexts, they are – broadly speaking – abstract, idealized, approximated or simplified theoretical tools, or, as they have been otherwise defined, they are surrogate structures (or vehicles) that can potentially represent the variety of real-world phenomena.

Since 1960s, models have started to play a crucial role in philosophy of science; the most debated questions are: *In virtue of what is a model to be considered a representation of something else? What is it for an object* (that is not a word or a sentence) *to represent a phenomenon scientifically?* The last century's literature has shown widespread disagreement over what it means to say that a certain model represents a certain state of affairs, as well as over *how* a certain model represents a certain system. As far as philosophers of science are concerned, the increased interest in the subject of scientific representation has disappointingly not been accompanied by an increase in our understanding of how models represent any particular target system of reality.

The aim of this contribution – which will be supported by historical evidences related to a case study model from today mathematical physics – is to show that models (i.e. abstractions and idealizations) are surrogates of the surrounding physical reality. Scientists usually connect the data of physical reality with the abstract features of models by means of a process that requires at least two steps: i) *interpretation*, with the intention that elements of an abstract model are provided with general physical interpretations (i.e. mass, position and velocity; and ii) *identification* or *denotation*, namely of

elements of a representational model which are identified with elements of the real system.

Within this perspective the role fulfilled by the user becomes of central importance; therefore by analysing the scientific practice I will put forward the idea that a model is a mind-dependent object, which exists only in the mind of the scientific community, defined as a social construction and developed within a certain cultural context. A scientific model is an epistemic representation insofar as it gives a representation of any particular set of physical phenomena under observation, even though the knowledge drawn about it does not correspond to what we directly observe in nature. There is, in fact, an epistemic gap between the conclusions that we infer from the model and what we directly observe in nature. Again, the model of a certain system is an epistemic representation of reality, insofar as it allows the user – by following a certain set of rules – to infer *some conclusions* and acquire *some knowledge* concerning that particular system of reality.

Ragnar Francén
Peripheral Concept Variation and Belief Individuation

University of Gothenburg

If A and B both believe that apples are fruit, they believe the same thing: that apples are fruit. According to *propositionalism about belief individuation*, A and B have the same belief in this sense in virtue of believing the same proposition. Relatedly, according to *propositionalism about 'that'-clauses in attitude ascriptions*, 'that'-clauses refer to (or signify) propositions. I will argue that these common and intuitive views fail to make sense of a phenomenon we can call *peripheral concept variation*.

Consider e.g., the ordinary concept of an apple. Most objects either uncontroversially are apples (e.g., an ordinary *Cox Orange*) or uncontroversially are not apples (e.g., ordinary horses). But then there are peripheral cases. Imagine an object exactly like an ordinary apple except that it was created in a laboratory. Is this man-made, apple-like object an apple? For such peripheral cases we can expect people's intuitive classifications to come apart. Two more examples: A chairlike object with a back that is only ten centimetres high – a chair or a stool? Construction barracks that workers temporarily live in, are they houses?

What characterizes peripheral cases in relation to a certain concept is that there has been little or no pressure towards consensus with regard to how to classify these cases/objects (include or exclude in the extension), e.g., because they have not been sufficiently prevalent or interesting to sort out or communicate about. Given this we can expect variation between how different people are disposed to classify peripheral cases. The first crucial step of my argument against propositionalism is that this variation is such that within the group of competent speakers, different people who use e.g., 'apple' (or 'chair', or 'house') express concepts with extensions that overlap in central cases but differ in 'the periphery'.

If this is correct, should we conclude that only some seemingly competent speakers have the concept of an apple, since concepts with different extensions cannot be the same concept, e.g., the concept of an apple? The second step of my argument is to reject this conclusion, since it doesn't make best sense of our second-order classificatory practice. As long as someone's use of 'apple' includes/excludes the central cases of apples/non-apples we are disposed

to think of (classify) her as talking and expressing beliefs about apples, irrespectively of how she classifies peripheral cases. This also makes sense, since, typically, peripheral cases are such that it has not mattered to us how they are classified.

The view that makes best sense of both our peripherally varying first-order classifications, and our second-order classifications, then, is that concepts whose extensions vary only in peripheral cases can count as concepts of the same thing (e.g., apples). Consequently, two people may both believe, say, *that apples are fruit* – and in this sense have the same belief – even if their beliefs have different propositional content; namely if they have different apple concepts. Furthermore, it follows that the 'that-clause' in question doesn't refer to some specific proposition.

Dirk Franken
Rigidity and Direct Reference

University of Münster

Rigidity and direct reference used to be defined as follows:

x is *rigid* =def. *x* has the same referent in all possible worlds in which the referent exists.
x is *directly referential* =def. *x* contributes nothing but its referent to any proposition expressed by a sentence in which *x* occurs.

According to an almost universally accepted assumption, directly referential terms are paradigm examples of rigid expressions. The following argument, however, shows that, if the definition of rigidity is taken at face value, both notions are not even compatible with each other:

(P1) A term's being rigid requires its expressing an intension whose extension is its referent
(P2) If a term expresses an intension whose extension is its referent, this term's contribution to a proposition is this intension and not its referent.
(C) Rigid terms cannot be directly referential.

(P1): According to the definition of rigidity, rigid terms *have referents in possible worlds.* Taken at face value, this means that, for any possible world, the term has a referent which is determined *relative to this possible world.* This, however, is just another way of saying what (P1) says: that the term expresses an intension whose value for any possible world is its referent in this possible world.
(P2): If the referent of a term is determined relative to possible worlds, this referent cannot be what this term contributes to a proposition, for what a term contributes to a proposition is precisely *not* determined relative to possible worlds. Hence, as (P2) says, such a term's contribution to a proposition is its intension, not its referent. The case of rigid terms is special insofar as the relevant intensions are constant functions. But that does not affect the truth of (P2). After all, the extensions of such constant intension are as relative to possible worlds as are the extensions of non-constant

intensions. Rigid terms turn out to be just a special class of descriptive terms.

I will try to show that this insight helps to clarify two central issues concerning the notion of rigidity: The question of the applicability of this notion to property-terms and the question of the role of rigidity in the explanation of the necessary *a posteriori*. Regarding the first question, I will argue that, in the final analysis, the claim of the applicability of the notion of rigidity to property-terms amounts to the claim that (some) such terms are descriptive in character. So, what is really at stake in this debate is whether certain property-terms are descriptive or directly referential. The second question, I will argue, has to be reconsidered in the light of the following: Since there are examples of the necessary *a posteriori* which consist of directly referential terms and since directly referential terms are not rigid, the necessity and *aposteriority* of these sentences has to be explainable without appeal to rigidity. I will propose such an explanation and will, moreover, argue that it is also the proper explanation of examples of the necessary *a posteriori* which contain of rigid expressions.

Grzegorz Gaszczyk
Selfless Assertions and Genuine Assertions

The University of Groningen

A selfless assertion is a speech act in which one states something that one does not believe. In the debate about assertion there is a universal agreement that selfless assertions are – even if somehow defective – genuine assertions. The goal of my talk, however, is to argue that selfless assertions are not – in any sense – "defective assertions." Rather, they are a different kind of speech act altogether. Lackey (2007) proposed selfless assertions as an argument against the knowledge account of assertion, according to which our assertoric practice is governed by the knowledge rule (KR) (e.g. Williamson 1996):

One must: assert that P only if one knows that P.

Consider an example of a selfless assertion. Stella does not believe in evolution, but she is a teacher and a part of her job is to teach the theory of evolution as a scientifically correct one. So, she says to her students "*Homo sapiens* evolved from *Homo erectus*" without believing it.

My argument for differentiating between assertions and selfless assertions will be delivered in the normative framework (the idea that there is an appropriate rule which governs the use of a speech act) supplemented with the common ground framework (see Stalnaker 2002 and the modification of his account made in García-Carpintero 2015). In a nutshell, during a communicational exchange we update the common ground (i.e. informations which the speaker and the hearer share), but not only assertions can update the common ground. Selfless assertions can do it too, but in a different way than assertions, i.e. there is a different rule which governs both practices and an intention which the speaker has (or commitment which she undertakes). My aim is not limited just to adding a new way of capturing selfless assertions, but I also propose a way of enriching the normative account.

In the remaining part of this presentation I indicate a highlight of my proposal. KR is nowadays the most common answer to the question what is the rule of assertion. From KR follows the assumption that a speaker believes what she asserts. This view is widespread in philosophy. Let us call it after Searle (1969) the "sincerity condition." Going further, and using the self-representation language, if we assume that KR is correct, we can

say that in performing an assertion, one is representing oneself as knowing that p, and one knows that p. In other words, to assert that p we must not only express but also possess knowledge towards the content of stated proposition. This symmetry is essential feature of assertion. In case of selfless assertions, however, one represents oneself as knowing that p, but one's mental state is different. Thus, there is an asymmetry between expressed and currently possessed mental state. On the other hand, Stella can sincerely and felicitously assert that she does not believe in evolution outside the school without falling into contradiction or being criticized. Thus, selfless assertions are allowed to make only in certain, specific contexts, and so they are governed by a weaker rule than assertions.

References

García-Carpintero, M. 2015 "Contexts as Shared Commitments", *Front. Psychol.*, 6: 1932, doi: 10.3389/fpsyg.2015.01932.

Lackey, J. 2007 "Norms of Assertion", *Noûs*, 41: 594–626.

Searle, J. 1969 *Speech Acts: An Essay in the Philosophy of Language*, Cambridge University Press.

Stalnaker, R. 2002 "Common Ground", *Linguistics and Philosophy*, 25(5–6): 701–721.

Williamson, T. 1996 "Knowing and Asserting", *Philosophical Review*, 105: 489–523.

Tarun Gidwani
Is Indexical Reference Sensitive to The Communication Technology Used?

University of Tartu

Indexicals are expressions, such as "I", "here" and "now", whose reference varies from one context to another. In Kaplan's semantic theory, indexicals have conventionally determined fixed rules, or characters, which determine their referent relative to a context, without appeal to speaker's intentions (Kaplan 1989). However, Kaplan's theory makes incorrect truth value predictions for indexicals when used over recording devices such as answering machines (Sidelle 1991). Several proposals have since then been made, that seek to modify Kaplan's theory, while giving us correct predictions in a rule-governed, intention-free manner. The most recent proposal is Character Shifting Theory (CST), according to which, the rules determining reference of indexical expressions vary according to the communication technology used, and remains invariant for uses within that communication type (Michaelson 2013). The device shifts only one aspect of indexical meaning; namely, the rules that determine the reference of indexicals, relative to a context. The news rules are established by conventions that arise as solutions to problems that occur when humans first learn to communicate using a new recording technology.

In my paper I argue against this theory by presenting two counterexamples, from two different communication technologies: postcards and internet chats. Using these two examples I show the CST fails in two ways. First, that typing contexts according to communication technology is not possible, as there are no coordination problems that occur when speakers communicate over these technologies using indexicals. In the absence of coordination problems, the character shifting theorist will have no intention-free explanation of the proposed shift in meaning. Secondly, I also show that if we did regiment context-types according to a technology, it would make incorrect predictions. While addressing the objection that CST is "viciously ad hoc" and regiments context-types in a manner that is "willy-nilly so as to make them match our untutored intuitions on what those terms refer to in particular contexts", Michaelson (2013, p.528) explains that CST is responding to our best available data - "competent language users' intuitions on the truth and falsity of well-formed tokens of natural language". I show that CST fails in

successfully explaining the same data. Following these counterexamples, I will consider four ways by which Michaelson (2013) can respond against the conclusions of my counterexamples and defend the CST. I then show that none of these four defenses hold.

CST is the most recent attempt that gives a rule-governed semantic account, where "indexicals remain context-sensitive and intention-insensitive, as Kaplan claimed" (Michaelson 2013, p.523). If I have successfully demonstrated the failure of this theory, it has implications on the larger debate about whether or not we can successfully explain communication using indexicals by assigning fixed rules, relative to a context, and still make correct truth value predictions.

References:

Kaplan, D. (1989). Demonstratives: An essay on the semantics, logic, metaphysics, and epistemology of demonstratives and other indexicals. In J. Almog, J. Perry, and H. Wettstein, editors, Themes From Kaplan, chapter 17, pages 481–563. Oxford University Press, New York. Originally circulated as a mimeograph from the Department of Philosophy, UCLA.
Michaelson, Eliot (2013). Shifty characters. Philosophical Studies 167 (3):519-540.
Sidelle, A. (1991). The answering machine paradox. Canadian Journal of Philosophy, 21, 525–539.

Zuzanna Gnatek
Co-application Conditions, Criteria of Identity and Fiction

Trinity College Dublin

Whereas neo-Fregean accounts of abstract reference and objecthood, such as Hale and Wright's, employ abstraction principles - the idea being that abstraction principles ensure that a relevant term refers, and thus that a corresponding object exists - some recent developments, an outstanding example being Amie Thomasson's easy ontology, aim at avoiding them. Thomasson claims that her account does not require a term to be associated with an abstraction principle, but only with a simple conditional. The aim of this paper is to question the extent to which this is feasible.

In case of Thomasson, this question has to do with a notion of co-application conditions which plays a crucial role on her approach. Thomasson defines co-application conditions as "rules that (supposing the term to have been successfully applied) specify under what conditions the term would be applied again to one and the same entity" and claims that a term has to be associated with such conditions. I argue that the exact way in which co-application conditions are specified requires them to be stated via an abstraction principle. The idea is, roughly, that their characterization presupposes the use of a so-called two-level criterion of identity which is standardly formalized as an abstraction principle. If this correct, then this raises significant difficulties for the overall view. Reference of many terms that easy ontology is famous for accounting for, such as fictional terms, is threatened - due to a constraint on abstraction principles that requires an equivalence relation involved in their right-hand sides to be grasped intuitively, which, I believe, cannot be met by those terms. Moreover, since the main difference between Thomasson's account and that of neo-Fregeans - that is, that she does not employ abstraction principles - can be questioned, it is not clear how exactly Thomasson's view is to be distinguished from theirs.

Selected References

Dummett, M., 1973. Frege: Philosophy of Language, Duckworth.
Linnebo, Ø., forthcoming. Thin Objects, Oxford University Press.

Thomasson, A., 1999. Fiction and Metaphysics, Cambridge University Press.
    —. 2007. Ordinary Objects, Oxford University Press.
    —. 2015. Ontology Made Easy, Oxford University Press.

Ziemowit Gowin
Chisholmian Objects: on the Controversies over Mereological Essentialism

University of Warsaw

Roderick M. Chisholm's doctrine of mereological essentialism is, by far, one of the most original metaphysical conceptions present in contemporary analytical philosophy. Mereological essentialism holds that wholes have their parts essentially, i.e., if x and y are parts of W, they are necessarily parts of W. This claim violates our intuitions regarding, among others, ordinary objects which seem to be possibly such that they may change their parts, e.g., a car may change one of its tires without ceasing to exist.

This radical view led Chisholm to the belief that most ordinary objects are so-called *entia successiva*, i.e., objects which are composed of "well-grounded" and ontologically independent simple substances – *entia per se*. *Entia successiva*, such as cars or tables, cannot be said to strictly and philosophically persist over time. A table T_a consisting of a, b, c at t_1 isn't (strictly and philosophically) identical to another table Tb if Tb lacks even one of Ta's components.

In the first part of my paper I will present the doctrine of mereological essentialism and the distinction between *entia successiva* and *entia per se* as well as the relations between these two kinds of objects and their distinctive features. I will also explain the ontological implications of mereological essentialism regarding human persons and their nature. Then I will move on to the main part of my paper: the arguments against *entia successiva* – *entia per se* distinction and its ontological consequences: (1) Jim Stone's argument from begging the question against four-dimensionalism, (2) Susan Anderson's inconsistency argument and – the most important one – (3) Lynne R. Baker's criticism of the very idea of the distinction between *entia successiva* and *entia per se*.

Jędrzej Grodniewicz
Linguistic Understanding Across Communicative Scenarios

LOGOS/University of Barcelona

The paper confronts the current orthodoxy in the philosophy of language according to which we should ascribe linguistic understanding of a given utterance to a subject if and only if she is in a mental (or complex—mental plus environmental) state of some specific kind: knowledge (Evans, 1982; Dummett, 1993; Heck, 1995), perception-like grasp (Hunter, 1998; Fricker, 2003; Pettit, 2002), content-entertaining (Longworth, 2016). I argue that, the common assumption of all these accounts is wrong. We will not be able to formulate a satisfactory characterization of linguistic understanding unless we observe that, in different communicative scenarios, different states may be equally aptly characterized as states of understanding.

My discussion proceeds in two steps. First, I argue that what is common to all instances of linguistic understanding is not a kind of hearer's cognitive state, but the specific cognitive process involved. We should, therefore, characterize linguistic understanding primarily as a cognitive process, which, based on a stable disposition to understand sentences in a given language coupled with other cognitive abilities (such as mind-reading, attending to, and remembering the common ground, etc.), gives rise to different kinds of states of utterance understanding.

Second, analyzing arguments present in the debate, I argue that states resulting from this process vary across communicative scenarios. The variance takes place along at least two dimensions. First, examples provided by Hunter (1998), Pettit (2002) show (even though it was not the intention of the authors), that some states of understanding are belief states (*), while some others are mere phenomenal seemings (+). Second, examples provided by Longworth (2016) show (again, *contra* Longworth), that some states of understanding represent in their content the speaker of the utterance (A), while some others do not (B). In result we obtain the following initial typology of states of understanding (where U stands for an utterance, and S stands for a speaker):

(*A) belief that what was said in U by S was that p
(*B) belief that what was said in U was that p

(+A) seeming that what was said in U by S was that p
(+B) seeming that what was said in U was that p

All four are kinds of states of linguistic understanding as long as they originate from an appropriate cognitive process. What introduces the variance is the communicative situation the hearer is in (or her beliefs about the communicative situation). The state of understanding is of type (*A) or (*B) if the hearer does not have reasons to doubt her ability to understand given utterance. If she has reasons to doubt it (as in Pettit's and Hunter's mad scientist scenarios), she ends up in a state of the type (+A) or (+B). If the hearer beliefs that she is in a highly cooperative or highly competitive scenario (in both of them it is important to track the identity of the speaker), she ends up in a state of the type (*A) or (+A). If in a given scenario the identity of the speaker does not matter or if the identity is unknown to the hearer, her understanding state belongs to the type (*B) or (+B).

References

Dummett, M. A. E. (1993). *The Seas of Language*. Oxford University Press.

Evans, G. (1982). *Varieties of Reference*. Oxford University Press.

Fricker, E. (2003). Understanding and knowledge of what is said. In Barber, A., editor, *Epistemology of Language*, pages 325–66. Oxford University Press.

Heck, R. (1995). The sense of communication. *Mind*, 104(413):79–106.

Hunter, D. (1998). Understanding and belief. *Philosophy and Phenomenological Research*, 58(3):559–580.

Longworth, G. (2016). Understanding what was said. *Synthese*, pages 1–20.

Pettit, D. (2002). Why knowledge is unnecessary for understanding language. *Mind*, 111(443):519–550.

Alison Hall
Pragmatic Enrichment and Underdeterminacy of Word Meaning

De Montfort University

Contextualists (e.g. Carston 2002, Recanati 2004) claim that what is explicitly communicated by an utterance (explicature) results from more pragmatic processing than Grice's 'what is said'. As well as mandatory rference assignment and disambiguation, linguistic meaning is freely enriched via 'unarticulated constituents' and lexical modulation. Several recent works challenge the idea that explicatures are mentally represented by the hearer in comprehension, arguing instead for a level of 'said' content that differs minimally from standing meaning (Borg 2016; Jary 2016, among others). Here, I respond to some recent arguments against the idea of explicatures as psychologically real.

Contextualists take explicature to be the result of extensive enrichment of standing meaning. However, Borg claims this enrichment is implausible and unnecessary, and that contextualists have overlooked factors such as social conventions in explaining the move from what was said (in a minimalist sense) to communicated contents (explicatures and implicatures). Jary (2016) provides support for this argument by pointing out that, for some contextualists (in particular relevance theorists), the enriched explicature is not necessarily involved in implicature generation, but only plays a kind of confirmatory role: if the implicature is warranted by a slight pragmatic enhancement to what was said, then the hearer's hypothesis about the implicature is confirmed. Explicatures, then, seem merely part of a rational reconstruction of implicature recovery; it does not follow that they are realized in the hearer's mind.

In response, I argue first that experimental evidence on processing of homonymy versus polysemy provides evidence for enrichment. Frisson (2008) showed that the time course of hearers choosing (provisionally) one sense of a homonymous word differs from the time course of polysemy processing: in the latter, hearers initially operate with an underspecified representation of the word meaning, before choosing one of the conventional senses. All words are *potentially* polysemous – i.e., can convey different concepts on different occasions of use – therefore the experimental evidence on polysemy resolution is also evidence for lexical modulation in general. Second, the objections fail

to fully acknowledge that these utterances take place in existing discourse contexts, where the concepts that contextualists claim are conveyed are likely to be already activated, therefore available for hearers to map linguistic meanings onto: the current 'Question under Discussion', or more generally the hearer's mental model of the context, consists of context-specific concepts of the kind that are claimed to form explicature, rather than of encoded lexical concepts. Third, I consider how far the objections apply across the range of examples marshalled by contextualists, and argue that only in conventionalized cases (e.g. responses to invitations/requests) could factors such as conventions of interpretation plausibly standardly circumvent the inference process from explicature to implicature posited by contextualists.

Borg, Emma. 2016. Exploding explicatures, *Mind & Language* 31: 335-355.
Carston, Robyn. 2002. *Thoughts and Utterances*. Blackwell.
Frisson, Steven. 2008. Semantic underspecification in language processing. *Language and Linguistics Compass* 3: 111-127.
Jary, Mark. 2016. Rethinking explicit utterance content. *Journal of Pragmatics* 102: 24-37.
Recanati, Francois. 2004. *Literal Meaning*. CUP

Rory Harder
Reference and Dynamic Sense

University of Toronto

In this paper I defend a Fregean theory of the meaning of demonstratives from the problem of nonspecificity. A Fregean theory holds that the meaning of a (use of a) demonstrative is a way of thinking of its referent. Ways of thinking are individuated more finely than objects of reference, so on this view the meaning of a demonstrative exceeds its reference. This is opposed to competing direct reference views that hold the meaning of a demonstrative is just its referent.

The problem of nonspecificity is based on observations in King (2017). King notes cases where a (sentence containing a) demonstrative can be used felicitously, but in which there is no fact of the matter which of a group of potential referents is the actual one. Here is one case. Suppose a number of us ordered new tablets all of the same model. They arive together in one box. You open the box and I notice that they are all of some other model. I say

(1) That's not the tablet I ordered!

It seems that nothing about the context, or my mental state, has to determine a specific tablet as the referent of my use of "that" for communication to be successful. There is nothing defective or objectionable about my utterance even if me and my audience members do not fix on the same one, or even any, of the specific potential referents.

The problem of nonspecificity is thus that successful demonstrative communication does not seem to require, in general, that the conversational participants think of the same object as referent, or even that each participant think of a specific object as referent. This is problem for both Fregean and direct reference theories, since they both entail that successful demonstrative communication requires (at least) that the participants fix on the same object.

Here is an outline of the Fregean response that I propose. I take as my starting point the view of Dickie and Rattan (2010), who develop Evans' (1985) notion of dynamic senses. A dynamic sense can include within itself multiple perspectives on an object, which could be had by different individuals at a single time or by a single individual as she is tracking

it over time. I argue that the phenonemon of nonspecificity shows that the dynamic nature of sense is even broader; it includes how multiple individuals can manage and negotiate potential referents through extended courses of conversation. First, I outline a notion of weak reference, which characterizes what it means for an object to be one of many potential referents for a demonstrative. Second, I show, by considering various continuations of discourses containing nonspecific utterances like (1), how the dynamic management of weak referents is a kind of rational engagement, and thus that the Fregean must posit a kind of sense to explain it. Finally, I outline an appropriate notion of sense, and defend it from some immediate objections (for example, that it seems to be one that does not determine reference).

Thomas Hodgson
A Defence of Structured Propositions

University College Dublin

I defend *Structure*: the view that propositions have constituents, against Simplicity the view that they do not. On my view, atomic propositions are complexes of objects and properties, and complex propositions are complexes of propositions and properties of propositions. This is a standard *neo-Russellian* view (Russell 1903).

Merricks (2015) defends *Simplicity*. My strategy in this paper is to accept, for the sake of argument, most of Merricks' view. There are propositions, they represent some things to be thus and so, and there is no explanation of these facts.

Propositions are about some things which are, in an intuitive sense, the subject matter of the proposition. Another way to put this is that these are the things that have to be some way or other for the proposition to be true or false. This suggests claim about aboutness that I will call *Truth* which constrains theories which make use of aboutness.

Truth A proposition P is about the set of things relevant to the truth of P.

Propositions also represent things to be thus and so e.g. a certain marble to be a certain colour. Representation is essential to propositional identity: representing differently is what makes propositions distinct. I will say that of a proposition represents that some things stand in some relation, it *represents of* those things and that relation. This suggests a claim I will call *Grasp*:

Grasp If a proposition P represents of a set of things S, then grasping the members of S is necessary to grasp P.

The motivation for grasp is that, in order to grasp P, one must grasp what it represents.

I will argue that if (i) propositions have constituents and (ii) they are about those constituents, there is a natural explanation of the facts of aboutness. Structure can, therefore explain something that Simplicity cannot.

An obvious response on behalf of Simplicity is that a proposition just is about those things that it represents of. I will argue that this is true of atomic propositions, but it is not always true of complex propositions. In particular, I will argue that an attempt to say that a proposition represents everything it is about is incompatible with Grasp, and that the claim that a proposition is about everything is represents is incompatible with Truth. I conclude that Structure is to be preferred to Simplicity.

References

Merricks, Trenton. 2015. *Propositions*. Oxford: Oxford University Press.
Russell, Bertrand. 1903. *The Principles of Mathematics*. Vol. 1. Cambridge: Cambridge University Press. http://people.umass.edu/klement/pom/.

Yang Hu
Context in Utterance Interpretation

Ecole Normale Supérieure de Lyon & Institut des Sciences Cognitives Marc Jeannerod

Based on Stalnaker's notion of common ground and Roberts's model of context in dynamic interpretation, I will first come up with the essentials of a complete utterance interpretation: 1. ***The explanatory task*** which includes an intended goal and the questions needed to answer for achieving the goal. 2.***The contextual-information-in-use set*** whose elements are pieces of contextual information constrained by the explanatory task. They are constrained in the sense that they are or can be actually used for the explanatory task. 3.***The Stalnakerian context set*** whose elements are possible worlds among which the interlocutor intends to distinguish. It represents the possibilities of the result of utterance interpretation.

I will then introduce two notions which can be deployed for grasping the contribution of contextual information to utterance interpretation.

The first notion is what I call the "***connecting force***" of contextual information. It is derived from the metaphor "context is an onion" proposed by Brézillon, et al (1997, 1999): A piece of contextual information has a connecting force just as each layer of onion skin has a "wrapping force". The connecting force can generate a connection (or connections) between a piece (or pieces) of contextual information and an explanatory task and between the first piece of contextual information and other piece(s) of contextual information. Again, the connecting force typifies the "explanatory force" of a piece of contextual information for a given explanatory task.

The second notion is what I call the "***efficiency***" of using context. It can be defined as this: *given an explanatory task and an interpreter, the efficiency denotes the proportion of a fulfillment of the task to the quantity of pieces of contextual information used by the interpreter for fulfilling the task*. With the notion of efficiency so defined, the use of a piece of contextual information can have an efficiency value.

Based on all above, a model of utterance interpretation can be recapitulated as follows:

Assume an utterance U to be the explanandum of a given explanatory task at a particular time during a talk exchange. The interpreter begins by using a piece of contextual information which is picked up from the common

ground relevant to the explanatory task. The identification and use of the piece of contextual information give birth to its connecting force making a link between itself and the explanatory task, and if it is the unique piece of contextual information that is required in dealing with the explanatory task, it is then the unique element in the contextual-information-in-use set. The context-based cognitive process leads to distinguishing the possible world, say, W_i, among others in the Stalnakerian context set. Given that the proposition P_i is the propositional counterpart of W_i, P_i would be the propositional result. Since the connecting force of the piece of contextual information does not establish any connection between this piece and other possible piece(s) of contextual information for obtaining the propositional result P_i, using this piece of contextual information gains its maximum efficiency value.

Madeleine Hyde
Imagination, Propositional Content and Modal Belief

Stockholm University

There are great benefits to thinking of various mental states as propositional attitudes. Understanding them as having a propositional content, for instance, helps us to see how mental states can represent the world i.e. say something true or false of it. Thinking of e.g. perceptual experiences as propositional attitudes can also help us to explain how these states can pass on their representational information, by having a certain kind of content, namely conceptual content, which they share with language.

This way of thinking of beliefs and perceptual experiences has been much-defended, but when it comes to other mental states like episodes of the imagination, it has often been taken for granted that, perhaps just in being sufficiently similar perceptual experiences, they must have propositional content too. I am going to motivate this understanding of imaginary episodes without appealing to their phenomenal overlap with the character of perceptual experiences.

Instead, I want to look at the epistemological similarities between imaginary episodes and other, nearby kinds of mental state like perceptual experiences. This will also bring us to their differences. At a glance: both perceptual experiences and imaginative episodes can give us reason to believe a certain proposition, namely the one that makes up their content. Where they come apart is in what *kind* of belief these states can provide propositional justification for. Perceptual experiences can give us reason to believe their content in terms of believing a fact about *how things are* with the *actual world, right now*. I call this kind of belief 'empirical belief'. Imaginative episodes, on the other hand, can only provide us with propositional justification for *modal* belief, i.e. belief about how things might, or could be i.e. about the state of affairs in some non-actual world.

This epistemological distinction reveals a difference in the kind of role these states play in our mental lives. Imaginative episodes, unlike perceptual experiences, cannot tell us facts about the actual world; or at least if they happen to do so, this is coincidental, and not to their credit. A successful perceptual experience, by comparison, achieves just this: hence we can see from looking at the kind of information they are fit to provide us with, what

different normative constraints we can place on the propositional content of our different mental states.

Robin Jeshion
Polysemous Expressives and The Problem of Appropriation

University of Southern California

Polysemy raises a host of important problems for any theory of meaning, cognition, and communication. My paper concerns three of the most central:

[1] How are polysemous lexical items represented in cognition?

[2] How are new related lexical meanings generated from old pre-existent meanings?

[3] Why do languages contain so much polysemy?

Two positions (broadly sketched) on polysemy representation and generation remain dominant. One, a lexically rich account treats polysemy as primarily internal, rooted on cognitively encoded generative rules that operate on informationally robust semantic representations, with only scant support from context when default interpretations are exhausted. [Pustejovsky (1995) and Asher (2011)] The other, a radical pragmatic account, treats polysemy as the result of speaker meanings accessed by inference from lean underspecified lexical meanings together with heavy reliance on contextual cues and encyclopedic information. [Sperber and Wilson (1986) and Carston (2002)] Call these the *Generative* and *Pragmatic* accounts. The Generative approach regards polysemy's prevalence primarily as indicative of the tendency toward communicative efficiency, using familiar terms to conveniently abbreviate new ideas, while the Pragmatic approach sees it as due to its ineliminability: speaker-meaning in principle outstrips lexically-encoded meaning.

Much of the debate has focused on certain types of polysemy:

- Inherent polysemy involving related senses of opposing semantic types. [The *book* I read has a thrilling plot. The *book* has a red cover.]

- The systematic polysemy of nouns that patterns with the syntactic count noun/mass noun distinction. [The *rabbit* nibbled my finger. There was *rabbit* all over the highway.]

- The polysemy generated by metaphor. [Adam is *weasel*.]

- The polysemy generated by metonymy. [The *cherry* needs pruning. The *ham sandwich* left without paying.]

Almost all these types of polysemy involve descriptive expressions.

   One area that has largely been neglected within theorizing about polysemy are expressives, including exclamations, expressive modifiers, lauditives and pejoratives, and slurs. Yet expressives are notoriously polysemous. Consider:

Expressive modifiers, initially negative shift polarity (so regularly they are now often taken simply as intensifiers):

- The paper is a bloody mess! → The paper is bloody awesome!

Exclamations:

- 'Ouch!' to indicate one's own pain → 'Ouch!' to note or express sympathy with someone else's pain.

- 'Holy Shit!' [frustration, disappointment] → 'Holy Shit!' [satisfaction]

Lauditives and pejoratives:

- "That bastard Kresge got a raise" [negative] → "Kresge, you old bastard, it's been so long!" [camaraderie]

- 'She's a sweetie' → 'You're a sweetie, spilling your milk all over the floor'

Slurs:

- 'Hillary is a bitch' [negative] →'Hell yeah, she's an angry bitch' [positive]

- 'Stay away, you are all queers' → 'We're here, we're queer, get used to it.' [defiant, proud]

- 'There are no Niggers allowed' → 'Come on, Nigger, let's go' [solidarity]

In this paper I argue that radical pragmatics provides far superior resources for explaining polysemy of expressives. I also argue, however, that to explain the polysemy in many expressives, radical pragmatics needs to widen its analysis of why polysemy is so widespread: as shown in cases where meanings are generated through appropriation, some new meanings are importantly parasitic on old, and, through echo, use of the old term is needed to effectively meta-encode the new meaning.

Asher, N. (2011) *Lexical Meaning in Context* Cambridge: Cambridge University Press.

Carston, R. (2002) *Thoughts and Utterances: The Pragmatics of Explicit Communication* Oxford: Blackwell.

Kaplan, D. [2005] "The Meaning of 'Ouch' and 'Oops'", unpublished manuscript.

Potts, C. [2005] *The Logic of Conventional Implicatures* Oxford: Oxford University Press.

Potts. C. [2007] "The Expressive Dimension" *Theoretical Linguistics*, 165-97.

Pustejovsky, J. (1995) *The Generative Lexicon* Cambridge: MIT Press.

Sperber, D., Wilson, D. (1986) *Relevance: Communication and Cognition* Oxford: Blackwell.

Wilson, D. (2006) The Pragmatics of Verbal Irony: Echo or Pretence?" *Lingua* 116, 1722-1743.

Andreas De Jong
Imagining an Actual Existent Unicorn. On Russellian Worries for Modal Meinongianism

University of Manchester

In this paper, I will address a worry against Modal Meinongianism, put forth in different form by Beall (2006), Sauchelli (2012), and Everett (2013). Their worries involve the ontological explosion of existent entities, which are reminiscent of Russell's (1905) original worries against the Meinongian characterization principle for objects: for any condition A, there is an object o, such that $A(o)$ is true. Since Meinongianism claims that existence is a property, Russell's initial objection was that if $A(o)$ is true, then $A(o) + existence(o)$ is true as well. Hence, for any condition A, there would be an existent o such that $A(o)$, leading to problematic entities, such as the existent Golden Mountain.

Modal Meinongianism is able to deal with Russell's original worries by employing a modal framework and a qualified characterization principle (QCP): for every condition A, there is an object o, such that $A(o)$ is true at some world (Priest 2016 [2005], p. 84). However, Beall (2006), Sauchelli (2012), and Everett (2013) argue that the problem resurfaces once we consider properties such as *actual existence*. QCP entails that for any condition A, there is an object o, such that $A(o) + actual\ existence(o)$ is true at some world. And since actual existence entails existence in the actual world, it would follow that for every condition A, there is an actually existent object that satisfies A.

In response, both Berto (2013)) and Priest (2016 [2005]) have responded to these worries by maintaining that the inferential patterns for expressions such as "actual" do not entail that such entities exist in the actual world. This is due to the fact that expressions such as "actual" need not have uniform operations in non-normal worlds. Those worlds are required to make true claims about impossibility, claims about intentional states not closed under logical consequence, and most notably claims concerning fiction and imagination.

Even though such a response seems formally adequate to steer the worries away, I will assess whether it is semantically adequate. Drawing on explorations by Lewis (1986) and Salmon (1987), I argue that the expression "actual" has a wide scope and narrow scope interpretation. The wide

scope interpretation is a use of the word "actual" that rigidly refers to our world. I argue that the wide scope interpretation of "actual" can be left out of the question. Given that Meinongianism's aim is to capture intentional phenomena, such a restriction is not *ad hoc*. As such, the use of the word "actual" always occurs within the scope of an intentional operator. Its occurrence within that operator blocks the inference that condition A is satisfied by o in the actual world and, hence, does not entail ontological explosion.

References

Beall, J.C. (2006). "Towards Non-Being: The Logic and Metaphysics of Intentionality". In: *Notre Dame Philosophical Reviews.*
URL: http://ndpr.nd.edu/reviewcfm?id=7503.
Berto, Francesco (2013). "Modal Meinongianism and Actuality". In: *Humana Mente* 25, pp. 155– 176.
Everett, Anthony (2013). *The Nonexistent.* Oxford: Oxford University Press.
Lewis, David Kellogg (1986). *On the Plurality of Worlds.* Oxford: Blackwell.
Priest, Graham (2016 [2005]). *Towards Nonbeing.* Oxford: Oxford University Press.
Russell, Bertrand (1905). "Review of A. Meinong, Untersuchungen zur Gegenstandstheory und Psychologie". In: *Mind* 14, pp. 530–538.
Salmon, Nathan (1987). "Existence". In: *Philosophical Perspectives* 1, pp. 49–108.
Sauchelli, Andrea (2012). "Fictional objects, non-existence, and the principle of characterization". In: *Philosophical Studies* 159.

David Kashtan
Dissecting Contexts

Tel-Aviv University

Kaplan's influential (1989)[1] work on indexicals contains both a philosophical thesis and a formal system. The thesis says that indexicals refer *directly*, i.e. through some kind of existential non-cognitive contact between the speaker and the referent. The formal system models this using quantification over contexts in the metalanguage, where contexts are construed as centered worlds. In this paper I do the following: (a) I argue that Kaplan's formal system is the wrong one for his philosophy, and locate the source of trouble in his metaphysics of context; (b) I extract Kaplan's main philosophical insight concerning contexts and use it to characterize several novel semantic notions; (c) I sketch how these notions can be used to reform indexical semantics, along with an application.

(a) I extract from Kaplan two adequacy criteria for an indexical semantic theory. First, such a system should allow the definition of a *pragmatic validity* predicate, which holds of sentences that are true whenever uttered, for example the sentence "I am here now".[2]

Second, the system should disallow the definition of "monsters", which are operators on indexical meanings (*characters*), or what comes to the same, operators defined in terms of quantification over contexts. In Kaplan's own system monsters are straightforwardly definable, and therefore, I contend, it does not constitute an adequate indexical semantics. Moreover, I show that Kaplan's treatment of pragmatic validities is ad-hoc and incomplete.

Kaplan views contexts as centered worlds. The main philosophical insight behind this view is the centrality of the linguistic agent. I argue, based on the conclusions above, that the centered worlds view inadequately captures Kaplan's main philosophical insight, and that any view that makes contexts out to be logically standard objects will be deficient. Rather, a foundational revision is called for, on which linguistic agency is not cashed out in terms of another kind of object, but informs the very way we set up the semantics.

---

[1]"Demonstratives", published in Almog, Perry and Weinstein (1989): Themes from Kaplan; but written in the 1970s and circulated widely before its publication.

[2]"Pragmatic validity" is my term.

(b) Identifying linguistic agency with content generation, I flesh out the insight that contexts are centered around a linguistic agent by developing a semantics for content generation. I revise the Kaplanian notion of *character*, so that it isn't modeled in the metalanguage by a mathematical function, but rather by a rule of content generation. I characterize a new type of expression in the pragmatic metalanguage, which is a sentence with prescriptive rather than descriptive force. I claim that this kind of metalanguage doesn't inherit the ontological commitments of its object-language, and that consequently we can express in it an adequate indexical semantics (one that doesn't give rise to monsters). The kind of reference operative in the metalanguage I dub *abstract reference.*

(c) In closing, I will sketch briefly a fuller formalized definition of *abstract semantics*, along with an application for a theory of abstract objects.

Filip Kawczyński
Analyticity and Meaning-Dependence Reconsidered

University of Warsaw

I put forward a series of arguments in defence of Meaning Holism ("MH"). According to MH — roughly speaking — every meaning depends on every other meaning. Somehow paradoxically, some philosophers have accused MH of leading to analyticity being *omnipresent* in language. Michael Devitt, for instance, ("Coming to Our Senses", 1996) attacks Inferential-Roles MH in such a way. My aim is to demonstrate that in fact MH does not pose any problems concerning analyticity. In particular, it does not entail the account according to which whole language should be considered analytic. My critique of the case against MH consists of the three parts.

(**I**) Devitt makes the assumption:

The thesis of Semantic Localism, according to which some-but-not-all expressions are meaningdependent, entails that some sentences are analytic;

and arrives at the conclusion:

The thesis of MH, according to which all expressions are meaning-dependent, entails that all sentences are analytic.

I argue that the reasoning is incorrect. Devitt erroneously applies to both Localism and MH the same notion of analyticity which I provisionally dub the *essential* analyticity. Whereas the essential analyticity is perfectly at home within Localism it is impossible to 'grow' such an analyticity on the ground of MH. The notion of essential analyticity is basically incoherent with the conceptual framework of MH, because formulating the thesis of MH is possible only provided that the essential analyticity has already been rejected.

(**II**) It can be said that Devitt's argument may involve some weaker notion of analyticity and that the weaker notion would pose a genuine danger to holism. I try to figure out what kind of analyticity would be compatible with MH and as a result I define the notion of relative analyticity. i.e. analyticity with regard to the state of a given language at a given time. I argue that the

relative analyticity being omnipresent is not really inconvenient because it does not pose the problems that the omnipresent *essential* analyticity brings. It is so because thanks to its empirical foundations, the relative analyticity preserves the link between language and the world.

(**III**) Aside from the above issues, I argue that Devitt's criticism is valid only if its target is the radical and 'crude' version of MH that I dub *obsessive-compulsive holism*, which — if fact — has never been endorsed by anyone The other and more reasonable version of holism I consider is *total-pair holism* presented by Peter Pagin ("Meaning Holism", 2006). I show that within this theory we can distinguish between various types of meaning-interdependence relations. Thus, anyone who would like to criticise MH for implying the omnipresent analyticity is supposed to specify which of the types of meaning-interdependence actually generates analyticity. I argue that there is no good answer for that — namely, it is impossible to reduce analyticity to a given type of meaning-interdependence without begging the question.

Katarzyna Kijania-Placek
Deferred Uses of Proper Names

Jagiellonian University

In typical cases of deferred use, an indexical refers to an object (a deferred referent) by way of a different object (an index), which is related in a contextually salient manner to the deferred referent (e.g. referring to an author by pointing at his book). Such uses were first analyzed by Nunberg in (1978) and (1993). In (1992) Nunberg suggested that, in contrast to indexicals, proper names do not have deferred uses. Nunberg's hypothesis has been contested and I will give examples of uses of proper names that arguably satisfy the core features of deferred reference as defined by Nunberg. In those examples names refer to objects other than their default referents and those objects do not bear the name but are in a relevant way related to the default referent (like in "Lassie has come again" uttered by a child looking at a stray Scotch shepherd dog standing in front of their house).

The idea that proper names in their deferred uses can refer to objects that do not bear the name has been opposed, for example by Manuel García-Carpintero (during a seminar at the Institute Jean Nicod). He argued that the property of bearing the relevant name should be considered in analogy to phi-features of pronouns and should thus be a required property of the referent. I argue, however, that we should not consider all phi-features in the same way. My arguments go back to the original formulation of deferred reference by Nunberg, who has divided linguistic features of expressions between deictic, classificatory and relational components. The deictic component is associated with the index and is responsible for its identification, the classificatory component concerns the deferred referent and the relational component constrains the relation between index and referent. For example, the deictic component of "these" dictate that the index is given by demonstration and is relatively proximate to the speaker, while the classificatory component requires plurality (Nunberg 1993).

I suggest that only the features that are possible candidates for the necessary properties of the relevant objects are required classificatory features and thus necessarily constrain the referent, while contingent features may be part of the deictic component. This hypothesis allows me to explain why the property of bearing a particular name - a contingent property - need not be a property of the deferred referent but may be considered part

of the deictic component. In this way only the default referent of a use of a name - the index - should bear the name, while this is not required of the deferred referent. This forms a defense of the deferred analysis of the proposed examples against Carpintero's objection.

References

Nunberg, G. 1978. *The pragmatics of reference.* Bloomington: Indiana Linguistics Club.

Nunberg, G. 1992. Two kinds of indexicality. In C. Barker and D. Dowty (eds.), *Proceedings from the Second Conference on Semantics and Linguistic Theory.* Columbus: The Ohio State University, pp. 283-301.

Nunberg, G. 1993. Indexicality and deixis. *Linguistics and Philosophy* 16: 1-43.

James Kirkpatrick
Generic Conjunctivitis

University of Oxford

There is a standard view of characterising sentences (e.g., 'Ravens are black', 'Ducks lay eggs', etc.) according to which they express quantificational propositions about what properties are shared by a contextually salient majority of the members of a kind. However, phrasal conjunctions embedded under a generic operator give rise to a challenge for this orthodox treatment (Carlson 1977; Schubert and Pelletier 1987 Nickel 2008; Nickel 2016; Liebesman 2011). Consider the following sentences:

(1) **a.** Elephants live in Africa and Asia.

  b. Elephants live in Africa and live in Asia.

  c. Peafowls lay eggs and have colourful tails.

Sentences in (1) involve equally good, but mutually incompatible characteristic properties, none of which are satisfied by the majority of the kind in question. For example, both (1a) and (1b) seem true even though a majority of elephants do not live in both Africa and Asia, and (1c) seems true, even though only female peafowl lay eggs and only male peafowl have colourful tails, and neither of which can consitute a majority. Call these sentences **generic conjunctions**.

Generic conjunctions have been used to undermine support for the orthodoxy and to motivate revisionist approaches to the semantics of generic sentences, such as semantics involving existential quantification over ways of being normal (Nickel 2008; Nickel 2016). This paper provides an explanation of generic con- junctions that vindicates the standard approach to generics sentences and thus undermines a central motivation for these revisionist semantics.

First, I observe that the phenomenon in question is more general than has been previously assumed in the literature. In particular, there are cumulative readings of conjunctions embedded under habituals, definite descriptions, and other quantificational phrases. Given the local nature of their solution, I argue that revisionist semantics are not equipped to deal with the problem of phrasal conjunctions in full generality.

Second, I argue that generic conjunctions can be accommodated from within the standard framework by admitting a more nuanced understanding of the role that conjunction plays in natural language. My particular solution follows Partee and Rooth (1983), Winter (2001), and Champollion (2016) by appealing to an intersective treatment of and in combination with independently-motivated phonologically null type-shifting operators. Using these tools, I show that generic conjunctions, as well as cummulative readings of non-generic conjunctions, can be accommodated from within an orthodox framework. I conclude by discussing the role that context plays in the pragmatics of selecting when cumulative and non-cumulative readings is are available for phrasal conjunction.

References

Carlson, G. N. 1977. A unified analysis of the English bare plural. *Linguistics and Philosophy* 1 (3): 413-457.

Champollion, L. 2016. Ten men and women got married today: Noun coordination and the intersective theory of conjunction. *Journal of Semantics* 33 (3): 561-622.

Liebesman, D. 2011. Simple generics. *Nous* 45 (3): 409-442.

Nickel, B. 2008. Generics and the ways of normality. *Linguistics and Philosophy* 31 (6): 629-648.

Nickel, B. 2016. *Between Logic and the World: An Integrated Theory of Generics.* Oxford: Oxford University Press.

Partee, B. and M. Rooth. 1983. Generalized conjunction and type ambiguity. In *Meaning, Use, and Interpretation of Language*, 361-383. Berlin: De Gruyter.

Schubert, L. K. and F. J. Pelletier. 1987. Problems in the represetnation of the logical forms of generics, plurals, and mass nouns. In *New Directions in Semantics*, 385-451. London: Academic Press.

Winter, Y. 2001. *Flexibility Principles in Boolean Semantics.* Cambridge, MA: The MIT Press.

Jiwon Kim & Paul Conlan
"Who Am 'I'?"– The Ambiguity of Virtual 'Me'

University of St Andrews

As society adopts new technologies, new philosophical puzzles arise. In this paper, we will present "The Virtual 'Me' Puzzle" (VMP), a set of 'monster' cases, which undermine Kaplanian Semantics. We generate the VMP by considering uses of 'I' and 'here' when used like "shifty indexicals" (Michaelson 2014) in the context of massively multiplayer online role-playing games. We argue that there are indexical uses in these contexts which are, at least, semantically ambiguous to produce a challenge to Kaplan.

According to Kaplan (1989: 492), the content of indexicals – 'I', 'here', 'now' – are context-sensitive: the character determines the content – the direct referents – in different contexts. The character of indexical is preserved within any context, but the content always varies depending on the context. Michaelson contends that the nature of indexicals to shift according to character-rules of Kaplanian semantics is a virtue which can be used to solve certain semantic puzzles.

We present a series of eight cases which together form the VMP. We defend the thesis that the VMP is a problem for Kaplanian semantics because the ambiguities in indexicality generate 'monster' cases. 'Monsters' are operators which change the context of interpretation of a term. Kaplan (1977) claims that 'monsters' do not occur in English and could not be added.

Consider the following toy case: As I log in to World of Warcraft, I type in the universal chatroom: "I am here now." Intuitively, 'here' in this case seems to refer to an online space, but, by default, Kaplanian Semantics tells us that "here" refers to the physical location of the player. Both appear to be correct, so the reference of 'here' is ambiguous in this context. Likewise, we present cases in which 'I' is more ambiguous as 'I' refers to the physical 'me', virtual 'me', and virtually physical 'me'.

Using the VMP, we defend the following claims:

1) There are multiple references for the indexicals, which make these indexicals semantically ambiguous.

2) Indexical ambiguity suggests that two propositions can equally be expressed by an indexical sentence.

Regarding 1), we analyse that there are multiple equally valid referents for "I": the real player, the body of the avatar, and the online presence of the player. Similarly, there are multiple locations that each utterance of "here" could refer to: the place that the real player is located, the place in the game-map where the avatar is, and the server where the online player is connected. The former is a location in physical space, whereas the latter two are 'virtual' locations. Regarding 2), we present a specific case in which an utterance appears to express two propositions due to the ambiguity in 1).

We argue that the VMP as a whole is problematic for the standard Kaplanian semantics because 1) and 2) imply that the VMP cases are 'monster' cases, which are barred from existence in Kaplanian Semantics.

Joanna Klimczyk

Normative 'ought', Semantic Disagreement and the Idea of Substantive Semantics for Normative Language

Polish Academy of Sciences

Here is a general definition of semantic disagreement that I assume to be roughly correct.

Semantic disagreement (SD):
A semantic disagreement is a situation where there is a speaker A, a speaker B, and a lexem l, such that:

(i) A and B remain in the same conversational situation, and

(ii) A produces an utterance using l

(iii) B produces an utterance using l (typically in response to A's claim)

(iv) A and B believe correctly that neither of them changes the topic of their conversation while producing the utterance with l

(v) A and B believe that their both use l correctly, i.e. in accordance with l's meaning

(vi) A and B attach different semantic properties to l

(vii) There is no evidence that either of them is guilty of a mistake in assigning to l the very content she assigns.

Now, consider the following mini-dialogue.

Speaker A: Kate ought to go and inform the manager.
Speaker B: *No*, not at all. Why do you think that it is *Kate's* job to do that. She works in a *team*!

In my presentation I will argue that the above piece of conversation can be sensibly interpreted as an illustration of a *semantic* disagreement

about the *normative* sense of the word 'ought'. If my arguments are not irreparably flawed, and the idea of purely semantic disagreement about normative vocabulary is a viable option, there is a lot of interesting points to be learned from such disagreement about the nature of conversational disagreement (for instance, that rejecting is not contradicting), and more importantly about semantics and metasemantics of normative ought-thought and ought discourse. In particular, I will argue that the possibility of being in a genuine dispute about the normative meaning of 'ought' speaks in favour of considering a certain neglected approach in the theory of meaning of normative language, namely one according to which a successful semantics of *normative* vocabulary should be built on certain substantive decisions in *metanormative* theory. I call the proposed approach 'substantive semantics' (SS).

I will sketchily present the central tenets of SS, and explain why I think of it as an attractive alternative to the extant theories.

References

Finlay, Stephen. 2017. Disagreement lost and found. In Russ Shafer-Landau (ed) *Oxford Studies in Metaethics*, 12, 187-205. Oxford: Oxford University Press.

Von Fintel, Kai, and Anthony S. Gillies. 2011. 'Might' made right. In Andy Egan and Brian Weatherson (eds), *Epistemic Modality*, 108-130. Oxford: Oxford University Press.

Klimczyk, J. 'Ought', agents and ambiguity that matters, forthcoming.

Klimczyk, J. *Normativity that matters*, book manuscript.

Plunkett, David, and Timothy Sundell. 2013. Disagreement and the Semantics of Normative and Evaluative Terms. *Philosophers' Imprint*, 13 (23), 1-37.

Silk, Alex. 2016. *Discourse Contextualism. A Framework for Contextualist Semantics and Pragmatics*. Oxford: Oxford University Press.

Robert Knowles
Alethic Appearances: In Conversation with 'The King of France'

University of Leeds

According to a popular view of natural-language semantics, we should try to capture all of our truth-value judgements about the target language semantically, leaving pragmatics to clean up after the party. Through a novel treatment of a familiar case study, I argue that this approach is sorely mistaken. The case study is the puzzling contrast in the truth-value judgements elicited by the following pair of sentences: (1) The King of France is bald; (2) The King of France is sitting in that chair. We are squeamish about assigning a truth-value to (1), while (2) strikes us as immediately false.

I start by identifying methodological shortcomings in extant accounts. These suggest that a semantic explanation of the contrast is likely to be implausible and unmotivated, and that a pragmatic explanation that fails to account for the role that truth-value judgements play in a wider theory of conversation will be ad hoc and unexplanatory. This paves the way for my own pragmatic account, which presumes nothing (controversial) about the semantics of the expressions involved, and locates our truth-value judgements in a wider story about how we achieve fast and effective communication.

On my view, as part of our ordinary utterance-interpretation procedure, we simulate given features of context along with the content of the utterance, and check the former against the latter for incompatibility. While there are extant proposals in this vein, they have failed to accommodate all the data, failed to represent a communicative role that needs filling, and failed to make contact with relevant psycholinguistic data. In contrast, my posited stage in our utterance-interpretation procedure is supported relevant psycholinguistic data, and is crucial for meeting both global and local aims of communication. On my account, as a result of the relative cognitive accessibility of the given features of context, an utterance can appear false, even if, on reflection, we may wish to revise the judgement this appearance elicits. This is what explains the contrast between (1) and (2). Unique among the extant accounts, my explanation predicts the right results for the range of recalcitrant data I consider in this paper.

The success of my pragmatics-first approach here, and the methodological issues prompted by other extant accounts, strongly suggest that explaining phenomena regarding our reactions to sentences, such as our truthvalue judgements, is the remit of empirically-informed theories of utteranceinterpretation, and that semantics should be concerned instead with assigning truth-conditions as part of providing the best systematisation of a broader set of data, including the output of pragmatics. I close by arguing that this way of divvying up labour between disciplines sheds light on seemingly intractable disagreements over the existence and theoretical utility of minimal semantic content, and the seemingly over-ambitious philosophical work sometimes expected of semantic theory.

Palle Leth
Do Intentions Matter for Utterance Interpretation?

Stockholm University

Intentionalists claim that it is the speaker's communicative intention, subject to more or less severe constraints, which determines those aspects of the meaning of an utterance which are not settled by lexicon, syntax and purely objective parameters; for instance, what is referred to, what is asserted, what is implicated [e.g. Grice 1967, Kaplan 1989, Recanati 2005, Korta & Perry 2011, Neale 2016]. In this paper I will attempt to show that a hearer engaged in utterance interpretation does not take the speaker's actual intention into account and that there is thus no evidence from the interpretive interaction between the speaker and the hearer in favour of intentionalism about utterance meaning. There seem to be three scenarios in which the hearer might be thought to take an interest in the speaker's intention and consider it as decisive.

1) The regular course of utterance interpretation. By default, when confronted with an utterance of the speaker's, the hearer sets out, using the linguistic material of the sentence and diverse contextual cues of the utterance as evidence, to discover the speaker's intended meaning. Often, the hearer also explicitly asks the speaker for clarification in case she does not get at the speaker's meaning [Schegloff et al. 1977]. However, the hearer's submitting to the speaker's authority in these cases does not amount to the hearer's taking the speaker's intention to determine the meaning of the utterance. Rather the hearer goes along with whatever meaning the speaker wants the hearer to go along with. The goal of this kind of interpretation is not correspondence with the speaker's intention, but rather obtaining the speaker's acceptance [Clark 1996].

2) Inaccuracy. In some cases, the hearer may suspect that the speaker is, because of oblivion, confusion or insincerity, not accurate about her intention and the hearer may take an apparent interest in the speaker's actual intention, as distinguished from the speaker's alleged intention. Such a suspicion will not, however, result in an inquiry into the speaker's communicative intention, but in a factual inquiry

into how things are in the world or in the speaker's mind [McDowell 1980, Stalnaker 2002] for which the speaker's actual intention is of no consequence.

3) Lying. According to many accounts, the speaker lies if and only if she asserts that p and believes that not-p [e.g. Stokke 2013]. Contrary to what most theorists assume, the notion of assertion at stake in lying does not involve the speaker's intention, but the most reasonable interpretation of the speaker's utterance, for which the speaker's actual intention is irrelevant [Burger 1973, Sperber & Wilson 2015].

Intentionalists might of course argue that interpretive practice does not matter for the metaphysics of meaning constitution. But as long as no positive evidence for intentionalism is presented, there seems to be no reason not to abide by what interpretive practice indicates: the speaker's intention does not determine utterance meaning, but is regulative only for one variety of utterance interpretation and is irrelevant in the remaining candidate scenarios.

References

Burger, C. J., 1973. *Bronston v. United States*, 409 U.S. 352, URL: https://supreme.justia.com/cases/federal/us/- /352/case.html.

Clark, Herbert H., 1996. *Using Language*, Cambridge, CUP.

Grice, Paul, 1967. Logic and Conversation, in *Studies in the Way of Words*, Cambridge, Mass. & London, Harvard University Press, 1991, pp. 1–143.

Kaplan, David, 1989. 'Afterthoughts', in Joseph Almog, John Perry & Howard Wettstein, eds., *Themes from Kaplan*, New York & Oxford, OUP, 1989, pp. 565–614.

Korta, Kepa & Perry, John, 2011. *Critical Pragmatics: An Inquiry into Reference and Communication*, Cambridge, CUP.

McDowell, John, 1980. 'Meaning, Communication and Knowledge', in Zak van Straaten, ed., *Philosophical Subjects: Essays Presented to P. F. Strawson*, Oxford, OUP, pp. 117–39

Neale, Stephen, 2016. 'Silent Reference', in Gary Ostertag, ed., *Meanings and Other Things: Themes from the Work of Stephen Schiffer*, OUP, pp. 229–342.

Recanati, François, 2005. 'Literalism and Contextualism: Some Varieties', in Gerhard Preyer & Georg Peter, eds., *Contextualism in Philosophy: Knowledge, Meaning, and Truth*, Oxford, Clarendon Press, pp. 171–96.

Schegloff, Emanuel A., Jefferson, Gail & Sacks, Harvey, 1977. 'The Preference for Self-Correction in the Organization of Repair in Conversation', *Language*, vol. 53, no. 2, pp. 361–82.

Sperber, Dan & Wilson, Deirdre, 2015. 'Beyond Speaker's Meaning', *Croatian Journal of Philosophy*, vol. 15, no. 44, pp. 117–49.

Stalnaker, Robert, 2002. 'Common Ground', *Linguistics and Philosophy*, vol. 25, pp. 701–21.

Stokke, Andreas, 2013. 'Lying and Asserting', *Journal of Philosophy*, vol. 110, no. 1, pp. 33–60.

Indrek Lobus
Self-Knowledge via Ontological Embezzlement

University of Stirling

Jane Heal (2001) proposes an account of first-person authority on which mental states are constituted by first-person thoughts about those mental states but which purports to avoid any conflict with naturalism about the mind. Her explanation of how the conflict is avoided is that a naturalistic account of the mind provides an *outline* for a mental state that can be *filled in* by a second-order thought. It is far from clear, however, what this means. Building on the works of Charles Travis, I will introduce the idea of shifting ontological standards as a way of understanding Heal's explanation, in the hope of making it not only intelligible but plausible. An account of first-person authority that emerges is one on which the second-order thoughts of subjects fix standards so that some state the subject is in, whatever it is, *counts as* the mental state the second-order thought is about — e.g. in thinking that I am scared, I count as scared. Once set up, the framework of ontological standards does make some parts of Heal's original proposal superfluous but also introduces interesting implications to it, notably the possibility of second- and third-person authority.

Wjatscheslaw Loev
Intuition Experiences as Cognitive Feelings

Institut Jean Nicod

In this talk I outline a feeling-based theory of intuitions. Its central claim is that intuitions are specific cognitive feelings: feelings of rightness or wrongness.

I will proceed as follows: I first identify my target state (so-called "intuition experiences") among the many things we call "intuitions". Since this target state is usually phenomenologically characterized I will subsequently outline and clarify respective phenomenological features. Then I will briefly present extant accounts of intuitions and explain why they are unsatisfactory. One kind of account identifies intuitions as a kind of state (beliefs, judgements, dispositions to believe/judge) that usually doesn't have the bulk of features ascribed to the target state. The other kind of account posits that intuitions belong to a mysterious (sui generis) kind of state (inclinations to believe/judge, seemings, presentations) which exhibits precisely the ascribed features but whose postulation appears otherwise explanatorily uninformative and ad-hoc. I will then advance cognitive feelings as the class among which we find our target state. Cognitive feelings such as feelings of familiarity or feelings of knowing are affective phenomena that share features with emotional and bodily feelings but are distinguished in that their function is to inform us about our ongoing cognitive processing rather than about the world or our body. Importantly, among these phenomena we find the feeling of rightness and the feeling of wrongness which have the function to signal that a certain representation satisfies (not) a normative requirement such as correctness or truth. I argue that intuition experiences just are such feelings of rightness or wrongness that detect the rightness/wrongness of an explicitly intentional representational cognitive state (e.g. propositional imaginings) with which they co-occur. This account fits the feature profile of intuitions and unmysteriously explains why intuitions have these features: because they are specific cognitive feelings (co-occurring with other states).

Marcin Łyczak
On Some Extension of the Logic of Change LC by Epistemic and Temporal Modalities

Cardinal Stefan Wyszyński University in Warsaw

The logic of change LC, was given by K. Świętorzecka in [1]. It is an extension of the classical sentential logic by the primitive operator C to be read: *it changes that ...* . In its epistemic interpretation the correlates of sentences are the convictions of some agent. When the agent obtains new information (the LC language is constantly growing), he classifies it as true or false. His convictions are represented by the conjunction of atomic expressions of level n. This may change at level $n + 1$. This change is restricted by the laws of classical logic and depends on the conjunction the agent has chosen at level n. In the presented approach we consider the changeability of convictions acceptable by agent. The acceptable convictions are the ones which the agent considers possible from the point of view of his current knowledge, i.e. convictions which he has at level n. Now the changeability represented by the operator C may be related to modalities of knowledge at level n. One of such modalities is the necessity represented by $\Box$. The subject of our interest is the description of relationships between $\Box$ and C. Our next step is to enrich the analyses with the functor of temporal necessity F (extension of LC logic by temporal necessity was presented in [2] expressing unchangeability.

References

[1] Świętorzecka, K., (2008), *Classical Conceptions of the Changeability of Situations and Things Represented in Formalized Languages*, Warsaw, CSWU
[2] Świętorzecka, K., Czermak, J., (2011), "A Logic of Change with Modalities", *Logique et Analyse*, 52(253), 511-527.

Piotr Makowski

Collective Intentionality and Automatic Imitation. A Hybrid Account Shared Action

University of Warsaw & Adam Mickiewicz University in Poznan

Classic accounts of shared actions (Bratman, Searle, Gilbert, Tuomela) are sometimes criticized as offering definitions of joint activities which are too semantic (Paternotte, 2014). This may make them inadequate to some empirically salient contexts of shared agency. It also seems that approaching the question of shared actions from the perspective of fully-fledged propositional attitudes as intentions or plans—partially it is an upshot of the semantic approach—may result in overlooking those aspects of doing things together which seem to be possible without persistent presence of propositional attitudes.

In cognitive and social psychology there has been growing interest for automatic information processing and performance (Bargh, 1994, 1996; Cheng, 1985; Moors & De Houwer, 2007). The importance of automatic actions has also been recognized early in organization science (Ashforth & Fried, 1988); they are also receiving more and more attention in the newest philosophy of action (Brownstein, 2014; Di Nucci, 2013; Makowski, 2017). And there is a growing awareness that the context of automatic actions changes the traditional philosophical image of human agency (Di Nucci, 2013). But, so far, it has not been fully considered when it comes to shared actions. One may reasonably expect that some cognitively salient forms of automaticity constitute important contexts of joint actions.

In my presentation, I intend to approach this issue by considering the idea of automatic imitation. I stipulate that the research on imitation and copying behavior which has recently emerged in cognitive psychology (Heyes 2011) may enrich our philosophical understanding of automatic actions. To what extent automatic imitation can condition shared activities? If there are salient cases of shared intentionality which necessarily involve automatic imitation, how such cases affect the picture of collective intentionality?

In my poster presentation, I attempt to respond the above questions by introducing and analyzing an example of the so-called Mexican Wave. I understand it as a case of complex, multi-agent activity, the performance of which involves—as I claim—both full-blooded propositional attitudes and those instances of intentionality which are possible solely as an automatic

(frequently blind) imitation. This leads to a *hybrid account* of shared action.

References

Ashforth, B. E., & Fried, Y. (1988). The mindlessness of organisational behaviors. *Human Relations*, 41(4), 305-329.

Bargh, J. A. (1994). The four horsemen of automaticity: Awareness, intention, efficiency, and control in social cognition. In R. Wyer & T. Srull (Eds.), *Handbook of Social Cognition* (pp. 1-40). Hillsdale, NJ: Lawrence Erlbaum.

Bargh, J. A. (1996). Automaticity in social psychology. In E. E. Higgins & A. Kruglanski (Eds.), *Social Psychology: Handbook of Basic Principles*: Guilford.

Brownstein, M. (2014). Rationalizing flow: agency in skilled unreflective action. *Philosophical Studies*, 168(2), 545-568.

Cheng, P. W. (1985). Restructuring versus automaticity: Alternative accounts of skill acquisition. *Psychological Review*, 92(3), 414-423.

Di Nucci, E. (2013). *Mindlessness*. Newcastle: Cambridge Scholars Press.

Makowski, P. T. (2017). *Tadeusz Kotarbiński's Action Theory - Reinterpretive Studies*. New York: Palgrave-Macmillan.

Moors, A., & De Houwer, J. (2007). What is automaticity? An analysis of its component features and their interrelations. In Bargh, John A. (2007). *Social Psychology and the Unconscious: The Automaticity of Higher Mental Processes*. Frontiers of Social Psychology. (Pp. 11-50). New York: Psychology Press. X, 341 Pp.

Paternotte, C. (2014). Constraints on Joint Action. In M. Gallotti & J. Michael (Eds.), *Perspectives on Social Ontology and Social Cognition* (pp. 103-123). Dordrecht: Springer Netherlands.

Benjamin Marschall & Lukas Lewerentz
Demonstratives and Conflicting Intentions

University of Cambridge
University of Oxford

Speaks (2017) argues that there is a serious problem concerning demonstratives such as 'this' or 'that', which can have different referents in different contexts: In order to understand how competent speakers are able to figure out what the referent of a demonstrative is, we need an account of how that referent is fixed in context. Plausibly, a demonstrative refers to an object o in context c only if the speaker intends that it refers to o in c. But a constraint along those lines leads to a problem with *conflicting intentions*: Suppose that a speaker mistakenly believes that some object A is identical with some other object B, and says 'This is nice', intending to refer to A with 'this' *as well as* intending to refer to B with 'this'. What do demonstratives refer to in cases like this? Speaks tries to show that all existing accounts fail to give a satisfying answer. We will argue against this by defending a largely unexplored strategy for dealing with conflicting intentions according to which demonstratives can have *multiple* referents.

In cases of conflicting intentions, a referential intention i_1 *trumps* another intention i_2 iff i_1 determines a referent for the demonstrative used but i_2 does not. According to *the no tolerance theory* (NTT), no referential intention trumps any other intention. Speaks dismisses NTT because he takes it to entail that demonstratives fail to refer in every imaginable case of conflicting intentions, which is implausible: At least in some of these cases, the demonstrative seems to have a referent. (Speaks 2017, 720-3)

As we will show, NTT entails reference failure only given the additional assumption that if a demonstrative has a referent in a certain case, it has *exactly one* referent. If demonstratives can have multiple referents, NTT is compatible with successful reference. Speaks quickly dismisses the possibility of multiple referents, but without providing any arguments. (Speaks 2017, 720)

In response, we will argue that postulating multiple referents in cases of conflicting intentions can be motivated independently: Sometimes, claims involving demonstratives that are used with conflicting intentions are treated as if they are ambiguous. This is so not only in some cases in which the speaker is aware that her intentions pick out different objects (Egan 2009).

As we want to show, it is also true in certain cases in which the speaker is not aware of the fact that her intentions conflict: If her audience realizes that conflicting intentions are in play, treating the utterance as ambiguous can be a normal reaction. The assumption that demonstratives can have several referents and may thus give rise to ambiguities, helps to explain this.

Admitting multiple referents is an underexplored response to cases of conflicting intentions. Since there are independent reasons supporting it, it is also a promising solution to the problem Speaks has raised.

Egan, Andy (2009). Billboards, bombs and shotgun weddings. *Synthese* 166 (2), 251-279.

Speaks, Jeff (2017). A puzzle about demonstratives and semantic competence. *Philosophical Studies* 174, 709-734.

Lucy McDonald
Sluts, Slurs and Neutral Counterparts

University of Cambridge

The view that all slurs have or could have 'neutral counterparts' is an unexamined orthodoxy in philosophical work on slurs. Most philosophers think that the neutral counterpart of the N-word is something like 'black person', and that equivalent inoffensive counterparts for all other slurs are possible. In this paper I argue that not all slurs have neutral counterparts. I show that although the word 'slut' is clearly a slur, in so far as it behaves in the same way as paradigmatic slurs in important respects, it doesn't have a neutral counterpart, because it contains evaluative semantic content that cannot be disentangled from its descriptive semantic content. The paper frames the problem as an inconsistent triad, consisting of the following three theses:

> T1: All slurs have, or could have, a neutral counterpart
> T2: 'Slut' is a slur
> T3: 'Slut' does not have, and could not have, a neutral counterpart

In Part 1, I motivate Thesis 1. A neutral counterpart is typically defined as an inoffensive, descriptive expression identifying the demographic group the slur targets. For example, the neutral counterpart of what is commonly regarded as the paradigmatic slur, the 'N-word', is something like 'black person'. There are several reasons for believing all slurs could have such a counterpart. Not only does this thesis find widespread support from philosophers and linguists, but also it explains why we might dispute the accuracy of a particular predication of a slur, even if we condemn the slur's derogatory component.

In Part 2, I motivate Thesis 2, with an argument by analogy. I show that 'slut' shares all the linguistic features of the paradigmatic slur, the N-word. Firstly, both exhibit *non-displaceability* with regards to their offensiveness; their offensiveness projects out of most linguistic operators. Both are offensive in a conditional sentence, a negation, a question, an imperative and a modal claim, and when used in past and future tenses. Secondly, both words are regarded by their users as having application conditions which can be disputed. Thirdly, both have numerous loose synonyms of varying offensiveness. Fourthly, both have been subject to attempts at reclamation.

This evidence suggests both words derogate through the same kind of mechanism and are both slurs.

In Part 3, I motivate Thesis 3. I reflect on the variety of contexts in which 'slut' is used, and consider two candidates for its neutral counterpart; 'promiscuous woman', and 'woman who has a disposition to sleep with many people'. I show not only that these candidates are too context-sensitive and too vague, but also that they ultimately fail to capture the referent of 'slut'. I then argue that there is no way of characterising the referent of 'slut' in purely descriptive terms. This is because 'slut' is a thick term and its semantic content is 'nonevaluatively shapeless', which means that there is no term for the non-evaluative, descriptive features of 'slut' that we are sensitive to when applying the word, because the feature that unites all women classified as 'sluts' is itself evaluative. 'Slut' has both descriptive and evaluative content, but this content is entangled; it cannot be separated out.

Theses 1, 2 and 3 cannot all be true. In Part 4 I argue that to resolve the trilemma we must concede, contra philosophical orthodoxy, that Thesis 1 is false. Some slurs do not have neutral counterparts.

Christoph Michel
Conditional Representations as Dependency-License-Models

Universität Regensburg

.

a. Two dimensions of conditional representation

I develop an account to conditional representations which appeals to intuitions about their basic cognitive functions. The core function is to represent "connections" between antecedent and consequent, referred to by "A" and "C". A "*Dependency-License-Model*" (DLM) aims at representing connections between A and C in two dimensions, "license" and "dependency". Although being systematically linked, both notions differ in what they represent and make salient to a cognitive system. Dependency is a *metaphysical* notion. For example, A can "depend on" or "require" C. Statements of license suggest inferences and suggest an extension of knowledge beyond the immediately given. This *epistemological* notion is akin to Ryle's (1950) idea of "inference tickets".

b. Disambiguation

While unmodified conditional phrases possess a syntactically clear, ordered one-dimensional "antecedent"/"consequent"-structure, DLMs represent more complex DL-patterns. One element is standardly fixed: By default, conditional statements offer a license to move to C, given A, $\{A \to C\}$. However, a dependency-structure, $\{A \Rightarrow C\}$ (read "C requires A") or $\neg\{A \Rightarrow C\}$ or $\{A \Leftarrow C\}$ or $\neg\{A \Rightarrow C\}$ remains unfixed and calls for disambiguation.[1] Two typical patterns may illustrate different frequent DL-structures.

(1) "If you know that p, you believe that p."

In a conventional reading, (1) gives a default license $\{A \to C\}$ to infer belief from knowledge and, as a statement of dependency, asserts a depedency of

---

[1]The arrows are not meant to be connectives. They visualize a structure rather than formalize.

knowledge on belief, $\{A \Leftarrow C\}$. From $\{A \Leftarrow C, A \rightarrow C\}$, like from DLMs in general, TA-, FA-, TC-, FC-cases can be derived.

(2) "If wind blows on Sunday mornings, Peter goes sailing."

(2) is a dispositional statement that offers a default license to infer or expect that C, given A. Here, wind is also an important reason for Peter to go out sailing. Moreover, wind is not merely a reason but also a requirement for enjoyable sailing. Therefore, also { Wind $\Rightarrow$ Sailing } seems to hold. This dependency suggests a corresponding license { Wind $\leftarrow$ Sailing }. Then, (1) and (2) gives us two frequent DL-patterns beyond the default license $\{A \rightarrow C\}$: $\{A \leftrightarrow C, A \Rightarrow C\}$ and $\{A \rightarrow C, A \Leftarrow C\}$.

c. Explaining reasoning

Context and background knowledge usually lead to a disambiguation of conditional phrases in terms of dependencies and licenses. A DLM determines reasoning in as far as it determines consequences to TA-, FA-, TC-, FC-cases. However, in many classical reasoning tasks different DL-patterns are not controlled for and sufficient disambiguation between DL-patterns not supported. I will look at different classical problems around conditional thought (the paradoxes of material implication, antecedent strengthening, transitivity, DA, AC, rejection of MP and MT, the Wason Selection Task). I compare DLM-interpretations of reasoning data to accounts from classical logic, Mental Models Theory, mental logic, the conditional probability account of conditionals and relevance theoretical approaches.

Daniel Milne-Plückebaum
On the (Meta-)Semantics of Virtualects

Universität Bielefeld

Pokémon GO (PG) players say things like:

(1) I see a `jigglypuff`!

Such utterances facilitate communication and coordinated action among PG players, and so are plausibly regarded as having content. Yet *prima facie*, "`jigglypuff`" is a species predicate; and a species predicate P's application condition is metasemantically dependent on those members of the species that have originally been picked out as falling under P (Kripke, 1980; Nimtz, 2017). But there are no such original samples to ground "`jigglypuff`"'s application condition. So "`jigglypuff`", and thus (1), is devoid of content.

Faced with this problem, I argue that "`jigglypuff`" isn't a *real* species predicate at all. That is, assuming that, in playing PG, players engage in a game of make-believe (in Walton's (1990) sense), I hold that this involves not only their make-believedly presupposing that certain non-linguistic facts obtain, but also their make-believedly speaking what I call a *Virtualect*; in particular, the expression "`jigglypuff`" is such that it's a species predicate only *within the PG-Virtualect*, as it can be associated with an application condition only make-believedly. Games of make-believe thus effectively play the role of (parts of) utterance contexts in which expressions like "`jigglypuff`" have been originally introduced as predicates denoting all and only members of a particular species, and thus in which certain expressions have contents they don't *really* have.

This picture gives rise to two questions: (i) What, in reality, *grounds* the fact that make-believedly, "`jigglypuff`" has the application condition it has? (ii) What kind of expression is "`jigglypuff`" *really*, and what content does it *really* have, if any?

In reply to (i), I propose that it's the real-world game itself, PG, which provides players with *denotation fixers* for the Pokémon predicates as used within the PG-Virtualect. One's external knowledge of PG *qua* augmented reality game thus metasemantically grounds the Virtualect spoken as part of the game of make-believe. In other words, while the PG-Virtualect's *semantics* can be specified only make-believedly, this "virtual semantics" is

metasemantically grounded in reality. In particular, the denotation fixer of what's make-believedly the PG-Virtualect's species predicate "`jigglypuff`" includes those *jigglypuffy* surface characteristics shown on players' smartphones.

In reply to (ii), I hold that what are *make-believedly* species predicates are really just that: *makebelieve species predicates*. Such expressions form a separate semantic category (*contra* Kripke (2013)), and understanding a make-believe species predicate involves knowing that it falls into this category. In search of a *real* content for the make-believe species predicate "`jigglypuff`", I show how one can externally get hold of its PG-Virtualect-*internal* content—by drawing on a broadly Fregean semantics (Künne, 2007) and by making use of *deflectors* (Kracht, 2015), i.e., devices to quote another language's *contents*. Specifically, I suggest that "`jigglypuff`", outside of the PG-Virtualect, is synonymous to the *deflector phrase* "`what, in the language of a PG player, is a jigglypuff`", where the language of a PG player is the PG-Virtualect. Thus, even non-PG-players can exploit the interplay between real-world descriptions and virtual contents to grasp (1)'s real content.

References

Kracht, Marcus. 2015. Using Each Other's Words. In *The Road to Universal Logic*, ed. Arnold Koslow and Arthur Buchsbaum, 241–249. Birkhäuser.

Kripke, Saul A. 1980. *Naming and Necessity*. Blackwell.

———. 2013. *Reference and Existence: The John Locke Lectures*. Oxford University Press.

Künne, Wolfgang. 2007. Fiktion ohne fiktive Gegenstände: Prolegomenon zu einer Fregeanischen Theorie der Fiktion. In *Fiktion, Wahrheit, Wirklichkeit: Philosophische Grundlagen der Literaturtheorie*, ed. Maria E. Reicher, 54–72. Mentis.

Nimtz, Christian. 2017. Paradigm Terms: The Necessity of Kind Term Identifications Generalized. *Australasian Journal of Philosophy* 95:124–140.

Walton, Kendall L. 1990. *Mimesis as Make-Believe: On the Foundations of the Representational Arts*. Harvard University Press.

Robert Mirski & Arkadiusz Gut & Krystian Macheta
Action in Social Context: the Development Of Social Cognition From an Action-Based Perspective

The John Paul II Catholic University of Lublin

In this paper, we sketch the significance of context for the development of social cognition within an action-based framework. We hold the action-based perspective against the traditional cognitivist approaches to social cognitive development and argue that it better accounts for existing empirical data that demonstrate contextual variance in social cognitive development. We also briefly discuss how emotion comprehension relates to cognitive-state comprehension within the framework.

The action-based frameworks upon which we draw are Carpendale and Lewis' socio-experiential approach (Carpendale & Lewis, 2004, 2006), Katherine Nelson's *Community of Minds* (Nelson, 2005, 2007), and Bickhard and colleagues' interactivism (Allen & Bickhard, 2013; Bickhard, 1998, 2009a, 2009b; Bickhard & Terveen, 1995). We signal convergence with other similar approaches where appropriate.

The central principle of the action-based perspective is primacy of action over representational structure. In traditional cognitivist approaches, all action is underwritten by representational competence; in an action-based framework – the reverse is true – action leads to the emergence of mental representation. This basic theoretical difference ramifies into how we understand development in each of the paradigms.

According to interactivism, mental representation is knowledge of interaction potentials or affordances. This builds on the newborn's initial reactivity over the course of development and her active exploration; mental structure forms based on interactive experiences that the child has lived through.

Social interactive competence will therefore derive from past social experience. Necessary representational capacities will emerge over interactive encounters with other people (Carpendale & Lewis, 2004, 2006). This interactive competence will not, however, involve representing abstract minds. The child is just a competent social agent and the representational abilities involved pertain only to physical or behavioral level of social exchanges. This is consistent with minimal interpretations of infants' social abilities (e.g. Fenici, 2015; Heyes, 2014; Ruffman, 2014).

Abstract reasoning about minds and mental states (including affective states) emerges in *linguistic* social interaction (Nelson, 2005, 2007). Conversational situations constitute a new kind of interactive context that enables representation of unobservable minds. The child through language acquires *an explicit theory of mind.* This is supported by the large body of research showing language's influence on explicit reasoning about the mind (see Astington & Baird, 2005). Moreover, it is consistent with the studies that show contextual variance in folk psychology and explicit psychological reasoning – e.g. the influence of family context (Devine & Hughes, 2016) or the cultural context (e.g. Dixson, Komugabe-Dixson, Dixson, & Low, 2017). This account is also minimally consistent with accounts that posit a disassociation between implicit and explicit social cognitive abilities (e.g. Fenici, 2012; Gallagher & Hutto, 2008; Newen, 2015).

We believe that an action-based perspective offers a more attractive framework to study social cognitive development than the traditional cognitivist theories. It acknowledges the multiply contingent nature of developmental processes and offers a comprehensive account of the cross-contextual variance in social cognitive abilities, which has proven problematic for the cognitivist approaches.

References

Allen, J. W.P., & Bickhard, M. H. (2013). Stepping off the pendulum: Why only an action-based approach can transcend the nativist–empiricist debate. Cognitive Development, 28(2), 96–133.
https://doi.org/10.1016/j.cogdev.2013.01.002
Astington, J. W., & Baird, J. A. (Eds.). (2005). Why language matters for theory of mind. New York, Oxford: Oxford University Press.
Bickhard, M. H. (1998). Levels of representationality. Journal of Experimental & Theoretical Artificial Intelligence, 10(2), 179–215.
https://doi.org/10.1080/095281398146824
Bickhard, M. H. (2009a). Interactivism. In J. Symons & P. Calvo (Eds.), The Routledge companion to philosophy of psychology (pp. 345–359). London, New York: Routledge Taylor & Francis Group.
Bickhard, M. H. (2009b). The interactivist model. Synthese, 166(3), 547–591.
https://doi.org/10.1007/s11229-008-9375-x
Bickhard, M. H., & Terveen, L. (1995). Foundational issues in artificial intelligence and cognitive science: Impasse and solution. Advances in psychology: Vol. 109. Amsterdam, New York: Elsevier.

Carpendale, J. I. M., & Lewis, C. (2004). Constructing an understanding of mind: The development of children's social understanding within social interaction. The Behavioral and brain sciences, 27(1), 79-96; discussion 96-151.

Carpendale, J. I. M., & Lewis, C. (2006). How children develop social understanding. Understanding children's worlds. Malden, Mass., Oxford: Blackwell Publishers.

Devine, R. T., & Hughes, C. (2016). Family Correlates of False Belief Understanding in Early Childhood: A Meta-Analysis. Child Development. Advance online publication.
https://doi.org/10.1111/cdev.12682

Dixson, H. G. W., Komugabe-Dixson, A. F., Dixson, B. J., & Low, J. (2017). Scaling Theory of Mind in a Small-Scale Society: A Case Study From Vanuatu. Child Development. Advance online publication.
https://doi.org/10.1111/cdev.12919

Fenici, M. (2012). Embodied Social Cognition and Embedded Theory of Mind. Biolinguistics, 6(3-4), 276–307.

Fenici, M. (2015). A simple explanation of apparent early mindreading: Infants' sensitivity to goals and gaze direction. Phenomenology and the Cognitive Sciences, 14(3), 497–515.
https://doi.org/10.1007/s11097-014-9345-3

Gallagher, S., & Hutto, D. D. (2008). Understanding others through primary interaction and narrative practice. In J. Zlatev, T. P. Racine, C. Sinha, & E. Itkonen (Eds.), Converging Evidence in Language and Communication Research. The Shared Mind (Vol. 12, pp. 17–38). Amsterdam: John Benjamins Publishing Company.
https://doi.org/10.1075/celcr.12.04gal

Heyes, C. (2014). False belief in infancy: A fresh look. Developmental science, 17(5), 647–659.
https://doi.org/10.1111/desc.12148

Nelson, K. (2005). Language Pathways into the Community of Minds. In J. W. Astington & J. A. Baird (Eds.), Why language matters for theory of mind (pp. 26–49). New York, Oxford: Oxford University Press.

Nelson, K. (2007). Young minds in social worlds: Experience, meaning, and memory / Katherine Nelson. Cambridge, Mass., London: Harvard University Press.

Newen, A. (2015). Understanding Others: The Person Model Theory. Advance online publication.
https://doi.org/10.15502/9783958570320

Ruffman, T. (2014). To belief or not belief: Children's theory of mind. Developmental Review, 34(3),265–293. https://doi.org/10.1016/j.dr.2014.04.001

Andrei Moldovan
An Intensional Theory of Presuppositions

University of Salamanca

It is common to find in the literature on definite descriptions reference to the "Frege-Strawson" theory of definite descriptions, (e.g., Kaplan (1970: 279), Kripke (1977: 269), Garcia-Carpintero (2000: 132), Salmon (2007: 69), Pelletier and Linsky (2005: 203), Elbourne (2013: 2), Schoubye (2013)), or to the "Frege-Strawson" theory of presuppositions (e.g., Beaver and Geurts (2013: §4.1)). But the differences between the Fregean and the Strawsonian theories are sufficiently important so as to motivate a separate treatment of the two theories. Although both theories introduce a presupposition of existence and uniqueness, these presuppositions are of different kinds. While a Strawsonian theory of definite descriptions is one in which the presupposition is a condition on the *character* (to use Kaplan's 1989 terminology) of the definite description, the Fregean theory should be understood as one in which the presupposition is a condition on the *content* of the utterance. In other words, a Strawsonian theory invites us to conceive of the character as a partial function, and of the content as a complete function. The Fregean theory conceives of the character as a complete function, and of the content as a partial one. While the Strawsonian theory is the usual theory of presuppositions that one finds in the literature, the Fregean intensional theory of presuppositions has received much less attention.

In the first part of the talk I present the differences between the two theories of presuppositions – with special consideration of definite descriptions – in the standard framework for truth-conditional semantics developed in Heim and Kratzer (1998) for extensional contexts, and Fintel and Heim (2011) for intensional contexts. The difference has no interesting effect in normal extensional contexts, but it becomes interesting in intensional contexts. The theory has interesting applications for the analysis of sentences containing definite descriptions in the scope of propositional attitude verbs.

In the second part of the paper I focus on such contexts, which make visible the intensional properties of expressions. In particular, I look at the interaction between definite descriptions and other presuppositional expressions with conditionals. As Stanley and Szabo (2000: 252) argue, the restriction of the domain of quantifiers embedded in the consequents of conditionals needs to be achieved by providing a contextually determined

property (that is, function from worlds to sets) and not simply a set. Intensional contexts, such as those introduced by conditionals, show that the domain of quantifiers needs to vary in ways that a simple extensional theory of domain restriction would not allow for. This idea offers indirect support for an intensional theory of presuppositions as well. In the case of definite descriptions, although they are not quantifiers on the Fregen presuppositional view, an account of contextual domain restriction is still needed for the case of the so-called 'incomplete' descriptions. I argue that a correct account of incomplete descriptions needs to consider a intensional theory of the presupposition of existence and uniqueness of definite descriptions.

Joanna Odrowąż-Sypniewska
The Problem of the Many, Supervaluation, and Faultless Disagreement

University of Warsaw

Recently is has been argued the problem of the many is a problem in semantics rather than in metaphysics. Moreover, Lopez de Sa (2014) claimed that Lewis was wrong in accepting his combined solution consisting of almost-identity and supervaluation, because almost-identity is a sufficient solution on its own, whereas supervaluationism faces serious difficulties as an answer to the problem of the many. According to him precisifications postulated by supervaluationism are not admissible: they do not preserve penumbral connections (because the cat-candidate that counts as the cat on any given precisification is arbitrarily chosen and is only marginally different from other cat-candidates) and do not preserve clear cases (since there is no cat-candidate which is the cat on all precisifications, there is nothing which is definitely the cat).

My talk will be divided into three parts. In the first part I'll try to defend supervaluationism from the above objection. I'll argue – following Quine (1960) – that two types of vagueness should be distinguished: the sorites-series type (which leads to the sorites paradox) and the fuzzy-boundaries type (which leads to the problem of the many). The authors who criticise supervaluationism for not preserving clear cases and penumbral connections confuse those two types and wrongly accuse precisifications postulated as a way to deal with the problem of the many of not satisfying requirements that are imposed on precisifications postulated to solve the sorites paradox.

In the second part I'll criticise Merlo's recent proposal to replace supervaluationists' indeterminate reference with multiple reference. On Merlo's view „Kilimanjaro" multiply refers to K1, K2, K3..., which are precise mountain-like things located roughly where Kilimanjaro is located (2017: 2646). The main problem for the view that regards "Kilimanjaro" as a multiple reference term referring to indefinitely many precise objects is that it has a consequence that no names for composite objects come out as singular.

In the third part I'll point out to a surprising consequence of Lewis's combined solution. Namely, his solution leads to the acceptance of faultless disagreements concerning the number of objects of the same type. He argues that there are two kinds of intended interpretations: those that regard all

cat-candidates as cats and those according to which the cat is just one and writes: "Context will favour one sort of interpretation or the other, though not every context will settle the matter" (1993: 179). If indeed there are contexts which do not favour one interpretation over the other, and if we imagine that in just such a context A and B say:

A: There is (exactly) one cat on the mat.

B: That's not true, there is more than one cat on the mat,

then they clearly disagree and moreover their disagreement appears to be faultless. In such a context the two ways of counting (by identity and by almost-identity) are equally right, so neither A nor B makes a mistake.

References

Lewis, D. 1993. Many but Almost One. In D. Lewis, *Papers in Metaphysics and Epistemology*, CUP.

Lopez de Sa, D. 2008. Is the Problem of the Many a Problem in Metaphysics?, *Nous* 42: 746-752.

Lopez de Sa, D. 2014. Lewis vs. Lewis on the Problem of the Many, *Synthese* 191: 1105-1117.

Merlo, G. 2017. Multiple Reference and Vague Objects, *Synthese* 194: 2645-2666.

Quine, W. 1960. *Word and Object*, MIT Press.

Shivam Patel
Inner Speech Is not Speech

University of Pittsburgh

Inner speech is a common phenomenon among the general population, and is often reported to occur during thought and silent reading (Uttl, et al., 2012). Philosophers and psychologists often assume that inner speech is the representation of speech: that it involves the representation of vocal tract manipulation and/or speech sounds (Langland-Hassan, 2008). Call this view 'concretism' about inner speech. I challenge this consensus position by arguing for 'abstractionism' about inner speech: that inner speech involves neither the representation of vocal tract manipulation nor that of speech sounds, but only abstract representations such as lexical, syntactic, and semantic representations. I argue against concretism and for abstractionism on both empirical and theoretical grounds.

The empirical data often cited in favor of concretism include EMG data showing minute vocal tract activation during inner speech (e.g., McGuigan and Dollins, 1989); neuroimaging data showing activation in motor/auditory areas of the brain during inner speech (e.g., Huang et al., 2002); and behavioral data implicating the representation of articulatory features during inner speech (e.g., Scott, et al., 2013). I argue that this data fails to support concretism. In the case of EMG data, I argue that vocal tract activation is not necessary for inner speech. Moreover, I argue that both neuroimaging and behavioral data fail to distinguish between inner speech, on the one hand, and the perception of inner speech, on the other hand (see also Gauker, forthcoming).

I next turn to challenge concretist models of inner speech. A number of psychologists and philosophers have used the comparator model to understand inner speech (e.g., Wolpert et al., 2011). There are two versions of the comparator model: according to the direct model, inner speech is identified with the execution of a command to produce a representation of articulatory/auditory information (e.g., Seal, et al. 2004), while according to the indirect model, inner speech is identified with the predicted sensory consequence of producing outer speech (e.g., Swiney and Sousa, 2014). The direct model is problematic, since the predicted sensory consequence and the execution of the command will be identical in both representational content and phenomenology, thereby making it arbitrary which to count

as identical to inner speech. However, if inner speech is identified with the predicted sensory consequence of producing outer speech, as on the indirect model, I argue that there is no plausible way to summate over excitatory and inhibitory activations. The comparator model is therefore not a plausible way to account for inner speech in a concretist framework.

I close by presenting empirical evidence in favor of abstractionism. One piece of evidence comes from McGuire et al. (1997), who has shown that inner speech and inner signing – a possible correlate of inner speech in the deaf population – activate similar brain regions, a result that is inconsistent with concretism. I conclude, on both empirical and theoretical grounds, that abstractionism is more plausible than concretism. I close by suggesting that the concrete aspects associated with inner speech are constructed by introspection on inner speech.

Denis Perrin
Deferentialist Account of Indexicals

Université Grenoble Alpes

There exist two main approaches to the semiotic profile of indexical expressions, that is: to the way they carry out their referential task. On the classical approach, indexical expressions draw on relations that pre-exist to the applying of their lexical meaning. For instance, on standard token-reflexivism (Reichenbach, 1947; Garcia-Carpintero, 1998), the relation of any token of "I" to the agent of its utterance/use is required by the lexical meaning to determine the relevant referent. But crucially this relation is merely factual, thus devoid of any semiotic significance, which is wholly provided by the lexical meaning. Though overwhelmingly dominant, the classical approach can be opposed an alternative. On the dissident view, the context of utterance of an indexical expression is not devoid of semiotic significance. Instead, it includes semiotic relations – perceptual, in particular – that pre-exist to the application of the linguistic meaning of the expression. For instance, consider a situation in which you know that a friend of yours stands in your apartment, but you don't know where. Imagine you ask where she is, to which she replies: "Here!". As you hear the token, your perceptual experience points toward where the sound is coming from. In other terms, the acoustic properties of the token of "here" designate in a pre-linguistic way, within the perceptual space of interlocution, the referent of the token. The same holds for the individual acoustic character of a token of "I" and for the temporal feature of a token of "now". In Bühler's terms (1934), who initiates the view, the semiotic profile of indexicals is (also) that of signals, who are endowed with their semiotic significance through perceptual features, rather than (only) that of symbols, whose semiotic significance is provided by their conventional lexical meaning. Their semantics is (sometimes, at least) embodied.

Arguing for this view, I hold that indexicals have a *deferential* lexical meaning. Their meaning includes a component that defers the referential task to semiotic features of the context, rather than it carries out this task by itself – borrowing Nunberg's terms (1993), they have a deferential deictic component. There are good reasons for buying this view. In particular, it provides an explanation of indexical reference that is more economic (thus more plausible) than the classical approach. On the latter, one would have

firstly to find out the context of application for the token-reflexive rule through the mentioned pre-semantic perceptual features, then to apply the rule to get a referent (Korta and Perry, 2011). Instead, the dissident view says, the perceptual determination of the relevant context is the referential operation, which is more economic, thus more in line with the requirements of face-to-face communication. The view faces objections, nonetheless. While it accounts for deictic uses of indexicals, one could observe, it cannot do so for their shifted uses (Recanati, 2004). Shifted uses, I reply, are anaphorical uses (Hunter, 2013) in which indexical meaning defers the referential task to pre-existing – linguistic, in this case – semiotic relations provided by the context, namely the anaphorical anchor of the indexical.

Olga Poller
The Challenges of Neo-Davidsonian Semantics

Pedagogical University of Cracow

Almost fifty years ago Donald Davidson (1967) proposed events semantics which was inspired by uses of sentences with adverbs and prepositional phrases. Since then his semantics has become the main tool for analyzing such sentences. Davidson's main idea was to add an additional argument place to action-predicates for an event-variable which could be bound by a quantifier. Later on, one version of Davidsonian semantics (enriched with states and mereology) was proposed by Zoltán Szabó (2003) to analyze uses of a particular prepositional phrases (as-phrases, as in 'John is rational as a chess-player'). Consider the following sentence:

(1) Michael is a brave chess-player.

(1) is ambiguous between intersective ('Michael is a brave person *and* is a chess-player') and non-intersective readings ('Michael is *brave as a chess-player*'). 'To be brave' and 'to be a chess-player' are state predicates and if you want to use Davidsonian semantics of events to analyze state predicates you have to extend this semantics and add states. The non-intersective reading of 'a brave chess-player' is analyzed as an intersective conjunction of predicates predicated of states (so it is a state (not a person) which is both a state of bravery and a state of being a chess-player). Now consider (2):

(2) Michael is a brave chess-player and a sneaky investor.

Analyzing Q as P as a simple conjunction of Q and P predicated of states, you will get a problematic consequence that Michael is simultaneously in brave and sneaky states. In order to avoid such a dilemma, Szabó (2003) additionally assumed a mereology of events and states: if someone is in a state Q and in a state P then he is in a state Q as P *only if* the state Q is a part of the state P. Now you can say that at the same time Michael was in two different states, a-chess-player state and an-investor state and the to-be-brave state was a part of only one of them.

In my talk I will attempt to demonstrate that adding mereology to events semantics could lead to challenges: for example mereology antisymmetry

axiom leads to an unwanted identity of events and the axiom of transition forces unwelcome entailments.

Tim Pritchard
Analogical Cognition and the Role of Abstraction in Language

University College London

Many words have extensions whose members are not related by perceptual resemblance (*obstacle, cause, holiday, carnivore*, etc.). Neither is there compelling reason to suppose we can relate the members by 'family resemblances'.

Over the past 40 years, psychologists (in particular D. Gentner and colleagues) have extensively studied 'analogical cognition', namely, reasoning based on understanding of higher order relations between items or events. GG is perceptually dissimilar to XX, but both display the relation of identity between their parts. A puddle is perceptually very different from a mountain, but both can enter into the relation of preventing the movement of something else (ant; helicopter). This common relation can be marked by the term *barrier*. Awareness of this kind of relation is characteristic of human cognition but is almost absent from animal cognition.

There is strong motivation for claiming that many (arguably most) words are used on the basis of an appreciation of higher-order relational commonalities. Without difficulty, we apply (very many) words on the basis of analogical relations. Empirical research shows that this ability requires explanation. Even for humans, there is a tendency to default to more obvious grounds of similarity unless certain conditions are in place. In particular, the research suggests that grasp of an abstracted schema, which represents the relational commonality but omits details specific to particular manifestations, is required to facilitate this kind of cognition.

From this, we can propose the hypothesis that an important aspect to our understanding of a word is grasp of a highly abstract schema. Without this, we lack adequate explanation for how we use words in the way we do.

This proposal has several consequences and challenges.

First, it indicates that grasp of a word is distinct both from general knowledge and the abilities that underpin our ability to categorize. By hypothesis, our understanding of *barrier* relates to a schema for prevention, but does not include our general knowledge about barriers, nor the abilities that enable us to identify barriers in the world.

Second, it indicates that grasp of a word should not be modelled along the lines of a label for some encyclopaedic conceptual entry ('pointer to a

concept'). Instead, our response to a word primarily involves the activation of a schema, with this schema playing a direct role in our thought. Further reflection on a particular manifestation of the schema is not part of linguistic understanding.

Third, it supports the claim that language helps to augment our cognitive capacities, rather than merely acting to 'translate' a thought into a form that enables it to be communicated.

Fourth, the proposal appears to run directly into Wittgenstein's challenge: is there 'anything in common' between different games. We may consider whether Wittgenstein's reflections apply also to higher-order relations.

Fifth, the use of many words (e.g. names for species, or terms that classify items in fine- grained ways) are not obviously based on higher-order relations.

Sixth, polysemy will need to be accounted for.

David Rey
An Intensional Account of English Embedded Tenses

University of Barcelona

In this paper I present a new semantic account of embedded tenses in English. Although the account that I want to recommend involves various theoretical assumptions, its central insight may be roughly summarized as follows:

The basic job of a tense is to locate an eventuality in time with respect to a parameter that I call *perspective time*[1] . In some grammatical environments, the time of utterance is the perspective time of a finite clause.

Consider, for example, the sentence

(1) Hillary married a man who became president of the U. S.

Here the perspective time of the matrix and embedded clauses is the time of utterance of (1). The semantic contribution of the past tense consists in locating the two events—the marrying event and the becoming-president event—before the perspective time, which happens to be identical to the time of utterance.

However, when a tense appears under the scope of an appropriate grammatical trigger, a switch of temporal perspective is activated and the eventuality is temporally located in relation to a shifted perspective time. Intensional verbs are triggers of perspective-time shifts. Modals can trigger perspectival shifts, too, but their shifting action is optional and is subject to certain restrictions in the idiolects of some speakers.

To illustrate, consider the sentence

(2) Mary claimed that John robbed a bank.

Here the intensional verb claim shifts the perspective time of the embedded clause to the time at which Mary makes her claim. As a result, the past tense of the matrix clause locates Mary's speech act before the time of utterance—which is the perspective time of the matrix clause— and the past tense of the embedded clause locates the event of robbing a bank before the time of Mary's speech act — which is the perspective time of the embedded clause.

Thus, the perspective-shifting effect of the verb claim explains the seemingly different behaviors of the embedded tenses of (1) and (2).

---

[1]My notion of perspective time is similar to that of Kamp & Reyle 1993, sec. 5.4.

I will argue that by combining these ideas with a principle of feature deletion and a sequence-of-tense (SOT) rule, we can correctly predict the different readings that have been discussed in the literature on embedded tenses, including the so-called simultaneous readings, double-access readings, and later-than-matrix readings[2]. I will also argue that my account predicts the differences between SOT languages such as English and non-SOT languages such as Polish regarding the interpretation of embedded tenses.

Although the account proposed in this paper can conceivably be implemented using different formal frameworks, I will adopt an intensional framework. I want to offer a plausible alternative to the referential approach that has been dominant in the literature since the eighties. The account of embedded tenses presented in this paper uses syntactic structures that are in line with modern generative syntax, but it does not posit time-denoting variables at the LF level.

References

Enç, Mürvet (1987) "Anchoring Conditions for Tense", *Linguistic Enquiry*, Vol. 18, No. 4, pp. 633-657.

Kamp, Hans, and Reyle, Uwe (1996) *From Discourse to Logic.* New York and Oxford: Oxford University Press.

Ogihara, Toshiyuki (1996) *Tense, Attitudes, and Scope.* Dordrecht: Kluwer Academic Publishers.

Kusumoto, Kiyomi (2005) "On the Quantification over Times in Natural Language", *Natural Language Semantics*, Vol. 17, Issue 4, pp. 317-357.

---

[2]For discussion of these readings, see Enç 1987, Ogihara 1996, and Kusumoto 2005.

Robert van Rooij, Katrin Schulz
A Causal Power Semantics of Generic Sentences

University of Amsterdam

Descriptive analysis of generics. Generic sentences (which we take to have the form '*G*s are *f*') are sentences that, by their very nature, express useful generalizations. But they express generalizations that allow for exceptions: although not all birds fly (Penguins don't), 'Birds fly' is a good generic. Generics like 'Birds lay eggs' show that it also need not be the case that almost all, or *most G*s have feature *f* in order for the generic '*G*s are *f*' to be true. Moreover, even if most *G*s are (or are taken to be) *f*, the corresponding generic sentence still doesn't have to be true, as exemplified by '*Italians are right-handed'. According to a natural alternative quantificational proposal, the generic is true exactly if all, or most, *normal G*s are *f*. But without an independent analysis of what it is to be a normal *G*, such an analysis hardly makes any empirical predictions. Second, such an analysis cannot account for the intuition that both 'Ducks lay eggs' and 'Ducks have colorfull feathers' are *both* true, because a normal duck cannot be both female *and* male. Van Rooij (2017) argued that all of these problems can be accounted for by demanding that for a generic '*G*s are *f*' to be true, the measure $\Delta^* P_g^f$ has to be significantly higher than 0, or significantly higher than $\Delta^* P_g^h$, with *h* any (contextually) relevant alternative feature to *f*, and with $\Delta^* P_g^f =_{df} \frac{P(f/g) - P(f/\neg g)}{1 - P(f/\neg g)}$ with $\neg g = \bigcup Alt(g)$: the natural alternatives to *g*, and $g \notin Alt(g)$. Notice that (i) $P(f/g) - P(f/\neg g) = \Delta P_g^f$ has to be positive for '*G*s are *f*' to be true (which explains why 'Italians are right-handed' is a bad generic), which means that the generic has to be true on what Cohen (1996) calls the 'relative reading' of generics, and (ii) that in contrast to ΔP_g^f, for $\Delta^* P_g^f$ the value of $P(f/g)$ counts for more than the value of $P(f/\neg g)$.

    An Explanation using Causal Powers. Now assume (with Cheng, 1997) that objects of type *g* have unobservable causal powers to produce features of type *f*, denoted by p_{gf}. It is the probability with which *g* produces *f* in the absence of any alternative cause. We denote by *a* the union of alternative potential causes of *f*, and by p_{af} the causal power of *a* to produce *f*. We will assume that p_{gf} is independent of p_{af}, and that both are independent of $P(g)$ and $P(a)$. The latter independence

assumptions are crucial: by making them we can explain the stability and context-independence of generic statements.

To derive p_{gf}, we will first define $P(f)$ assuming that f does not occur without a cause and that there are only two potential causes, g and a and that g and a are independent: $P(f) = P(g) \times p_{gf} + P(a) \times p_{af} - P(g) \times p_{gf} \times P(a) \times p_{af}$. Then we can derive $P(f/g) = p_{gf} + (P(a/g) \times p_{ag}) - p_{gf} \times P(a/g) \times p_{af}$ and $P(f/\neg g) = P(a/\neg g) \times p_{af}$. As a result, $\Delta P_g^f = P(f/g) - P(f/\neg g) = p_{gf} + (P(a/g) \times p_{af}) - (p_{gf} \times P(a/g) \times p_{af}) - (P(a/\neg g) \times p_{af})$ and thus $\Delta P_g^f = [1 - (P(a/g) \times p_{af})] \times p_{gf} + [P(a/g) - P(a/\neg g)] \times p_{af}$. From this last formula we can derive $p_{gf} = \frac{\Delta P_g^f - [P(a/g) - P(a/\neg g)] \times p_{af}}{1 - P(a/g) \times p_{af}}$.

Because a is taken to be probabilistically independent of g, $P(a/g) - P(a/\neg g) = 0$. Moreover, $P(a/g) \times p_{af} = P(f/\neg g)$. As a result, p_{gf} comes down to $\frac{\Delta P_g^f}{1 - P(f/\neg g)} = \Delta^* P_g^f$. We have **explained** the above descriptive analysis of generics, and **grounded** it by providing truth-makers for it: the causal powers.

The pros and cons of the causal power analysis of generics. As for the pros, (i) the causal power theory is **explanatory** where the frequency analysis using $\Delta^* P$ is not; (ii) the theory can **distinguish** good *generics* from *accidental generalizations*, (iii) having truth-makers means that generics have stable, context-independent meanings and can thus **express propositions** as well when embedded, (iv) the analysis is **more general**. For the latter, notice that using p_{gf} one is not required to assume independence of a from g. For instance, if g and a (the assumed alternative cause of f) are *incompatible*, one can show that $p_{gf} = P(f/g) \neq \Delta^* P_g^f$. This explains, we will argue, why people assume generics to come with a high conditional probability (cf. the similarity with Adam's thesis for the analysis of conditionals). We'll argue that this is relevant especially once powers are associated with *essences* of kinds. As for the cons, in contrast to an analysis using $\Delta^* P_g^f$, the causal power analysis has to assume that even if all and only all Gs are F, at most one of the two generics 'Gs are F' and 'Fs are G' can be true on the standard reading, because in contrast to frequency measure $\Delta^* P_g^f$, the causal measure p_{gf} is essentially **asymmetric**.

Giorgio Sbardolini
Towards a Model of Conceptual Change

The Ohio State University

Concepts may change. In time, a word expressing a concept c may come to express a concept c' . Examples come from the history of science (newtonian *mass* and relativistic *mass*) and historical linguistics (*deer* used to mean "animal" in Old English). In philosophy, it is sometimes argued that concepts should be changed if defective (cf. Haslanger (2012) on socially significant concepts like *woman* or *marriage* , and Scharp (2013) on paradoxical concepts like *truth*). In this paper I sketch a model of conceptual change. The goal is to develop a plausible diachronic account of lexical semantics, and to assess philosophical projects (such as ameliorative inquiry or conceptual engineering) that rely on conceptual change. Intuitively, conceptual change is the outcome of an aggregation problem in a competitive game.

Preliminaries: (1) Despite some "platonistic" accounts of the metaphysics of words and meanings, conceptual change seems to occur, at least at some level of description. I assume metaphysical accounts of meaning (Richard, 2016) and words (Kaplan, 1990), on which the meanings of words may change. (2) I'm talking about concepts as *meanings*, in a broad (not-only-semantical) sense. This is not an equation, of course. Strictly speaking and more carefully, in this paper I sketch a model for the study of semantic change, assuming that the meaning of words is determined at least in part by how speakers tend to use them.

Models for conceptual change may be adapted from models for social change from Social Choice theory (Arrow 1963). These models abstract away for the reasons speakers may have for change. Idealizing, we represent meanings as sets of features, and let each speaker i of L assign a set of conceptual preferences to a word w of L, i.e. an ordered set of alternative features such that the meaning of w according to i is characterized by i's preferred alternatives:

$$P_i(marriage) = \{[\text{male/female}] \leqslant [\text{genderless}], [\text{polygamous}] \leqslant [\text{monogamous}], \ldots\}$$

With respect to marriage, speaker i (weakly) prefers it to apply regardless of gender rather than to male/female relations only, and to apply to monogamous relations rather than polygamous, etc. Different speakers of L have

different preference sets. Given a distribution of conceptual preferences for each speaker (with respect to at least a word), an aggregation rule is a method to calculate a social preference, i.e. the preferred alternatives of the population as a whole. A model of change is given by different assignments of preferences at different times.

The simplest aggregation rule is majority (as in voting systems), but it's doubtful that majority leads to a plausible account of the phenomena. Evidence comes from the role played in conceptual change by a distinguished set of speakers, who seem to have some priority over the use of certain words. Examples are target groups in reappropriated slurs, and scientific communities for the meaning of scientific terms (Putnam, 1975). Different hypotheses about the aggregation of preferences can be studied rigorously by making different assumptions about the model.

References

Arrow 1963, *Social Choice and Individual Values*, New York: Wiley, 2nd ed.
Haslanger 2012, Resisting Reality, Oxford: Oxford University Press.
Kaplan 1990, 'Words', *Proceedings of the Aristotelian Society*, 64, 1, 93–120.
Putnam 1975, 'The Meaning of "Meaning"', *Minnesota Studies in the Philosophy of Science*, 7, 131–193.
Scharp 2013, *Replacing Truth*, Oxford: Oxford University Press.

Oscar Salvatore Scarpello
Perceptual-Working-Concept

University of Genoa and University of Turin

I will argue that the content of the early vision is primarily protopropositional (Peacocke 1992), that the perceptual concepts are based on such contents, and that it is possible that some visual perceptive experiences are influenced by higher-order cognitive states, therefore that perception is cognitively penetrated.

Protopropositions are an additional layer of non-conceptual content (ie: mental representations of the positioned scenario content). Our experiences of properties such as color, length, temperature, height and volume that are all continuous and analogical, will be segmented and digitalized, in order to understand the scene that is being observed. These representations are now conceptual, but with the same fineness of grain of non-conceptual representations. The result of this process is the Perceptual-Working-Concept (PWC). Once the perceptive experience is over, the PWC loses a large part of the fineness of grain it possessed during the perceptive experience, and it is stored in memory, becoming a persistent perceptual concept or going to corroborate pre-existing perceptual conceptual categories.

To defend my assumptions, I will assume that the PWC are formed through the language entry rules (Putnam 1981), which allow us to introduce a perceptual concept for an object on the basis of a protoproposition. As a result, I will relate cognitive models on phenomenal similarity spaces (models of perceptive experiences from non-conceptual content) with the conceptual spaces of Gardenfors (2014), (Brossel 2015). The dimensions in the phenomenal similarity space represent the area within which the objects can be experienced as similar or dissimilar, and the distances between points in space represent the level of phenomenal similarity of objects within the space. Gardenfors suggests that concepts are associated with certain regions in a conceptual space. Conceptual spaces are geometric spaces and their geometric structure is used to capture properties and semantic relationships. The idea is to model the properties and semantic relationships of pure perceptual concepts and PWC, belonging to the conceptual space, as regions in the phenomenal similarity space.

Furthermore, I support the idea that top-down influences on perception allow an observer to integrate relevant information with sensory inputs at

an even more global level , in order to optimize the information processing. This also makes it possible for future perceptions to be influenced by these concepts. Indeed, compatibly with the Cognitive Penetration Lite, there will be two Tokens of different perceptual experiences belonging to the same Type (Macpherson 2015): an experience with non-conceptual content and an experience with perceptual representations penetrated by concepts. The condition for the existence of the second Token is that the observer has already been in appropriate causal relations with what would make that content correct, i.e. the observer has to have had a previous experience belonging to the first Token that has been subsequently conceptualized.

Affirming the PWC's existence on the basis of non-conceptual contents allows us to justify the perceptive beliefs: on the one hand without falling into coherentism, on the other hand saving the perceptual beliefs from the myth of the sense-data. Furthermore, it is also compatible with the idea that there is no clear separation between perception and cognition. (Nanay 2017; Neven, Vetter 2014).

Henry Ian Schiller
Context and Implicit Content

University of Texas at Austin

It is standardly claimed that the semantic content of an utterance can be determined by that which a speaker's audience must entertain in order to have understood what was literally meant by that utterance. This picture of semantic theorizing— often called 'moderate contextualism—holds that at least some semantic content must be supplied by the 'wide' context of the utterance (facts about speaker intentions, facts about what is mutually assumed by speaker and hearer, etc.).

If semantic content is to be determinable (in part) by what a hearer must entertain in order to understand what a speaker has said, then there is a puzzle that must be solved. As Buchanan (2010, 2013) notes, what it takes to understand utterances of sentences like,

(1) All the beers are in the fridge.

(2) Simone is ready.

(3) The dog ate your clothes.

is in some sense indeterminate. These are all sentences that are typically used to express what we might call *implicit* literal meaning — there is some sense in which their explicit meaning needs to be completed or expanded in order for utterances of these sentences to express a full proposition (or to express a proposition that makes sense) (Bach 1994). But for just about any utterance of a sentence like (1), there will be multiple intensional contents we can identify as being what was literally said (that all the beers *in the house* are in the fridge in the kitchen, that all the beers *we bought for this party* are in the fridge *in this house*, etc.). If what it takes to understand a speaker in these situations (qua entertaining a content) is to be a guide to semantic content, then this semantic content must reflect this indeterminacy.

Rather than focus on understanding, I will focus on the variety of ways in which speakers may misunderstand one another, and what kind of explanations hearers need to be given so that these misunderstandings can be resolved. I argue in this paper that the resolution of misunderstanding in cases with *implicit* meaning, like (1), (2), and (3), typically involve two features:

i. It is always indeterminate exactly what needs to be said in order to resolve the misunderstanding,

ii. Resolution of the misunderstanding involves expressing an explicit content which always seems to shift the meaning of what was said. That is: it is not possible to further specify what was said without saying something *slightly* different.

We can contrast this with explicit communication, where specification of what was said (the resolution of misunderstanding) involves a kind of coordination on word meaning, and does not require any sort of shift from one meaning to another. This observation leads me to make two claims, which are jointly the core theses of this paper.

Thesis 1 There is a difference in the kind of content expressed by explicit and implicit aspects of speech. Explicit communication requires that a speaker understand intensional content, whereas implicit communication is extensional. Any resolution of misunderstanding involving implicit communication which involves intensional content invariably shifts the meaning of the initial utterance.

Thesis 2 What determines the domain in cases of quantifier domain restriction has to do with mutual assumptions held by speaker and hearer, more than speaker intentions. This is a position which has recently been defended for general semantic theorizing by Armstrong (2016), King (2014), Plunkett & Sundell (2013), and involves a kind of pragmatic basis for metasemantic theorizing.

References

Armstrong, Josh. 2016. The problem of lexical innovation. *Linguistics and Philosophy* 39. 87–118.

Bach, Kent. 1994. Conversational impliciture. Mind & Language 9(2). 124–162.

Buchanan, Ray. 2010. A puzzle about meaning and communication. Noûs 44(2).

Buchanan, Ray. 2013. Conversational implicature, communicative intentions, and content. Canadian Journal of Philosophy 43(5-6). 720–740.

Camp, Elisabeth. 2006. Contextualism, metaphor, and what is said. Mind & Language 21(3). 280–309.

Cappelen, Herman & Ernie Lepore. 2005. *Insensitive semantics*. Oxford: Wiley-Blackwell.

King, Jeffrey C. 2014. Speaker intentions in context. Noûs 48(2). 219–237.

Plunkett, David & Tim Sundell. 2013. Disagreement and the semantics of normative and evaluative terms. Philosophers' Imprint 13(23). 1–37.

Wedgwood, Daniel. 2007. Shared assumptions: Semantic minimalism and Relevance Theory. Journal of Linguistics 43. 647–681.

Maciej Sendłak
On the Pragmatic Approach to Counterpossibles

University of Warsaw

The subject of this paper is counterpossibles; i.e., subjunctive conditionals of the form 'If it were the case that A, then it would be the case that C' ('A>C'), where 'A' expresses impossibility. The problem of counterpossibles is the problem of whether the proper semantic account of counterfactuals should consider every counterpossible to be true. The advocates of what Timothy Williamson called the 'orthodox view', and what was developed in the classic works of Robert Stalnaker (1968) and David Lewis (1973), give a positive answer to this question and argue that every counterpossible is vacuously true (Williamson 2016). The unorthodox opposition argues in favor of a modified account, according to which some counterpossibles are true and others are false (Yagisawa 1988, Nolan 1997).

Recently, Nina Emery and Christopher Hill (2017) developed a pragmatic approach toward counterpossibles. Their central aim is to provide analysis that on the one hand takes into consideration our pre-theoretical intuition about the non-vacuous truth value of counterpossibles and, on the other hand, explains these away without introducing changes to the orthodox semantics of counterfactuals. This assumes shifting the burden of the problem from the question of truth value to the question of assertion.

Following this line of thought, we may say that one merely *tends to assert* a counterpossible (1) 'A>C' rather than to assert (2) 'A> ¬C'. Nevertheless, this should not lead us to the conclusion that (1) and (2) have different truth value. After all—as Emery and Hill pointed out—'[t]here are lots of substantive and true propositions that are never asserted, because speakers regard them as uninteresting' (2017, 138). It seems that (2), is assumed to be uninteresting, nevertheless true, proposition.

The aim of my paper is to argue in favor of the claim that Emery and Hill's proposal is affected by the problem of *tu quoque*. Consequently, it allows one to claim that every counterfactual is vacuously true. This seems to be an unwelcome consequence of Emery and Hill's approach for two reasons. The main reason is that this fails to explain the fact that we do not assert (in the same context) counterfactuals of the same antecedent and opposite consequences. The second one is that critiques of unorthodoxy are often committed to the standard possible worlds' approach toward coun-

terfactuals, according to which some of them are true and others are false. The exposition of this problem will help to indicate a close bond between counterpossibles and counterfactuals with merely possible antecedents. We believe that this bond, along with the assumption that some counterfactuals are false, supports the unorthodox view.

Emery, N. and Hill, Ch. 2017. Impossible Worlds and Metaphysical Explanation: Comments on Kment's *Modality and Explanatory Reasoning. Analysis* Volume 1, Issue 1, January 2017: 134–148.

Lewis, D. 1973. *Counterfactuals.* Oxford: Blackwell.

Nolan, D. 1997. Impossible Worlds: Modest Approach. *Notre Dame Journal of Formal Logic* Vol. 38. No.4: 535–572.

Stalnaker, R. 1968. A Theory of Conditionals. In, *Studies in Logical Theory*, ed. N. Rescher, 98–112, Oxford: Blackwell.

Williamson, T. 2016. Counterpossibles. *Topoi* [online first].

Yagisawa, T. 1988. Beyond Possible Worlds. *Philosophical Studies* 53: 175–204.

Joan Gimeno Simó
De Se Attitudes, Bound Variables and the Logic of Indexicals

University of Valencia

David Kaplan's highly influential theory of indexicals (1989a, 1989b) is committed to a certain view on the interactions between intensionality and indexicality: namely, that indexicals are directly referential expressions, and as such they are not affected by the behavior of intensional (modal or temporal) operators. A consequence of Kaplan's views on indexicality and content is that the so-called essential indexicality (Perry 1979) or *de se* nature (Lewis 1979) of self-ascriptive attitudes is systematically lost when such attitudes are reported: in Kaplan's terminology, what matters in attitude reports is the original utterance's content, and not its character.

It is usual to present two distinct sets of data against Kaplan's ideas: on the one hand, the existence of indexical expressions that happen to be sensitive to the original utterance's character – namely, logophoric pronouns and shiftable indexicals (Schlenker 2002, 2003, 2017; Anand & Nevins 2004) – and, on the other hand, the functioning of English expressions involving PRO (Chierchia 1989). A common strategy for handling these cases is to postulate monster-friendly semantics in which context-shifting operators are allowed. In this paper, however, we will rely on a different strategy that, we hope, will provide a unified account of all these issues: we will be claiming that all the cases in which Kaplan's theory is allegedly incapable of properly reflecting *de se* readings of these sentences can be traced to the behavior of variables at the level of logical form. Namely, we will be arguing that *de se* and *de re* readings of a sentence stem, respectively, from bound and free uses of variables. Our account requires no context-shifting, but it is still monstrous, since it requires a different, more interesting kind of monsters, one dubbed "content monsters" (Rabern & Ball 2017).

Our proposal is structured in three parts. First, we will summarize Kaplan's views on indexicality and intensionality, as well as the consequences of his commitment to direct reference. Next, we will discuss the above mentioned data and show that it all boils down to a single pattern: all de se readings of sentences display Quantifier Raising and bound uses of variables, whereas *de re* readings stem from free uses of variables. We will argue that this idea provides evidence against thinking of bound and free

uses of variables as the semantic counterpart of anaphoric and deictic uses of pronouns, since it can be shown that anaphora may occur even in free uses of a variable. Finally, we will discuss the consequences that this idea bears on a recent discussion on the nature of content and the alleged existence of monster operators in natural language (Rabern 2012, 2013; Rabern & Ball 2017), and thus challenge Kaplan's commitments to direct reference.

References

- ANAND, P & NEVINS, A. (2004) "Shifty Operators in Changing Contexts"., in Young, R. (ed), *SALT XIV*, 20-37, Ithaca, NY: Cornell University

– CHIERCHIA, G. (1989), "Anaphora and Attitudes de se", in Bartsch, R. & van Benthem, J. & van Emde, B. (eds.) *Language in context.* Dordrecht: Foris.

– HEIM, I. & KRATZER, A. (1998), *Semantics in generative grammar.* Oxford: Blackwell Publishers.

– KAPLAN, D. (1989a), "Demonstratives", in Almog, J. et al. (eds.), *Themes from Kaplan*, Oxford, Oxford University Press, pp. 481-564.

– KAPLAN, D. (1989b), "Afterthoughts", in Almog, J. et al. (eds.), *Themes from Kaplan*, Oxford, Oxford University Press, pp. 565-614.

– LEWIS, D. (1979), "Attitudes de dicto and de se", *The Philosophical Review*, vol. 88 (4), pp. 513-43.

– PERRY, J. (1979), "The Problem of the Essential Indexical", *Nous*, vol. 13 (1), 3-21.

– PREDELLI, S. (1996), "Never Put off until tomorrow what you Can Do today", *Analysis*, Vol. 56(N°2), pp. 85-91

– PREDELLI, S. (2014), "Kaplan's three Monsters", *Analysis*, vol. 74, pp. 389-93.

– RABERN, B. (2012), "Against the identification of assertotic content with compositional value", *Synthese*, 189: 75-96.

– RABERN, B. (2013), "Monsters in Kaplan's LD", *Philosophical Studies*, 164: 393-404.

– RABERN, B. & BALL, D. (2017) "Monsters and the theoretical role of context", *Philosophy and Phenomenological Research*, pp. 1-25.

– SCHLENKER, P. (2002), "Indexicality, Logophoricity, and Plural Pronouns", In Lecarme, J. (ed.), *Afroasiatic Grammar Ii: Selected Papers From the Fifth Conference on Afroasiatic Languages*, Paris, 2000. John Benjamins. pp. 409-428.

– SCHLENKER, P. (2003), "A Plea for Monsters", *Linguistics and philosophy*, vol. 26 (1), pp. 29-120.

– SCHLENKER, P. (2017), "Supermonsters: attitude and action role shift in sign language", *Semantics and pragmatics* 10(9).

Daniel Skibra
Ought and Agentivity

Northwestern University

Normative *ought*-statements are often said to come in two sorts. Consider (1)

(1) The president ought to re-join the Paris climate accord.

On a fairly standard way of construing the meaning of *ought*, (1) means something like *the state of affairs wherein the president rejoins the Paris climate accord is a better one than one where he does not.* Evidently, a different kind of meaning is exhibited by (2).

(2) John ought to return the lost money.

It's not simply that the world where John returns the lost money is better than one where he doesn't; (2) enjoins John to perform a particular action return the money. Philosophers moved by this distinction point out that *ought*-sentences like (2) play a special kind of role in our practical deliberations. (1) and (2) exemplify the distinction between evaluative and deliberative *ought*, respectively. (The distinctions between ought-to-be/ ought-to-do (Castaneda [1975], Feldman [1986]), and relativized/ unrelativized *oughts* (Grice [2001]) are in the conceptual neighborhood.)

There is no consensus on how to account for this distinction in meaning. Horty [2001] and Horty and Belnap [1995] gives a logic purported to model deliberative *ought*, but it has been criticized as unable to provide a meaning for the expression *ought* for lack of linguistic motivation. Schroeder [2011] offers an ambiguity thesis, contending that the difference in meaning is due to the syntactic difference between a raising *ought* and a control *ought*. On his proposal, deliberative ought is a control verb that takes two arguments, an agent and a property, as opposed to a single, propositional argument. This provides a principled basis for the distinction, but the claim that *ought* is a control verb has been disputed on empirical grounds. Another proposal, due to Chrisman [2012], suggests that deliberative *ought* can be given a non-ambiguity-theoretic analysis with Kratzer's modal semantics if the conversational background against which *ought* is interpreted makes

explicit the relevant agent carrying out the action. However, this proposal fails to give us a precise recipe for distinguishing the two senses.

I will give my own account of the distinction between deliberative and evaluative ought without positing a lexical ambiguity. My contention is that it owes to a difference in the logical form of *ought* sentences. The analysis will make use of two independently needed semantic and syntactic resources for seemingly unrelated problems in the semantics of modals. One resource pertains to the future-orientation of deontic modals with non-stative prejacents. These have an obligatory future-orientation, but the reason for this remains a mystery on standard modal semantics. The first ingredient of my account is the mechanism whereby this future-orientation is obtained. Next, I take a leaf from verbal semantics. It is well-known that the subject ("external") argument of a verb has a different relation to the verb than the internal arguments. On this basis, Kratzer [1996] suggests that the external argument is not a true argument of the verb, and composes with verbal expressions by means of a special composition principle known as **Event Identification**, which gives us our second ingredient. With this we have a robust, uniform, and linguistically motivated distinction between deliberative and evaluative *ought*.

References

H. Castaneda. *Thinking and Doing: The Philosophical Foundations of Institutions.* D. Reidel Pub. Co., 1975.

M. Chrisman. 'ought' and control. *Australasian Journal of Philosophy*, 90(3):433-451, 2012.

F. Feldman. *Doing the Best We Can: An Essay in Informal Deontic Logic.* D. Reidel Publishing Company, 1986.

P. Grice. *Aspects of Reason.* Clarendon Press, 2001.

J. F. Horty. *Agency and deontic logic.* Oxford University Press, 2001.

J. F. Horty and N. Belnap. The deliberative stit: A study of action, omission, ability, and obligation. *Journal of Philosophical Logic*, 24(6):583-644, 1995.

A. Kratzer. Severing the external argument from its verb. In L. Z. Johan Rooryck, editor, *Phrase structure and the lexicon*, pages 109-137. Springer, 1996.

M. Schroeder. Ought, agents, and actions. *Philosophical Review*, 120(1):1-41, 2011.

Izabela Skoczeń & Aleksander Smywiński-Pohl
Quantifier Scope Ambiguities of Scalar Terms in Incomplete Knowledge Contexts

Jagiellonian University
AGH University of Science and Technology

N. Goodman and A. Stuhlmüller have made an experiment to confirm a model that was supposed to predict the probability of scalar implicature formation in contexts where speaker and hearer do not have full knowledge of the situation. (Goodman and Stuhlmüller, 2013) During the experiment the speaker saw for example three apples. However, he had information concerning the color of only two of the apples (red). The hearer was aware of the speaker's state of knowledge. The utterance 'some of the apples are red' cancelled the implicature 'not all' since hearers inferred a similar probability of two and three apples being red. If the speaker saw two out of three apples and said 'one of the apples is red', then the hearer inferred a similar probability of one and two of the apples being red and an extremely low probability of three apples being red. Thus, the implicature 'not more than one' was cancelled. During the experiment various scenarios were used and the speaker had a varied access to objects and their features, while the hearer was aware of the speaker's state of knowledge. The experiment confirmed a fine-grained interaction between the state of knowledge about the world and pragmatic inference. It also provided a strong argument against modular theories of mind since the data showed that the language faculty and the inference about world state were strongly connected. However, there remained the question whether the effects were cross linguistic or rather confined to the English language.

In our paper we present the replication of the above-described experiment in the Polish language. We used the same scenarios and the same quantifiers and numerals as Goodman and Stuhlmüller. Namely, the Polish words 'niektóre' for 'some' and 'wszystkie' for 'all' as well as the numerals 'jeden, dwa, trzy' for 'one, two, three' were used. The data for numerals replicated neatly in an experiment made with the participation of students. However, with a group of varied age, some differences occurred pertaining to scalar implicatures. Similar differences were found in a replication of the experiment with quantifiers. Our hypothesis is that the differences reveal a quantifier scope ambiguity that appears in the experiment when the Polish

translation is employed. We investigate this hypothesis in order to reveal the nature of cross-linguistic differences of human reasoning in incomplete knowledge contexts.

Selected references:

1. Clark, R.L., 2012. *Meaningful games: exploring language with game theory.* MIT Press, Cambridge, Mass.

2. Bergen, L., Goodman, N., Levy, R. Pragmatic reasoning through sematic inference, *Semantics and Pragmatics*, 2016, vol. 9

3. Franke, M., 2009. *Signal to act: game theory in pragmatics.* Institute for Logic, Language and Computation, Amsterdam.

4. Frank, M.C., Goodman, N.D., 2012. *Predicting Pragmatic Reasoning in Language Games.* Science 336, 998–998. doi:10.1126/science.1218633

5. Goodman, N.D., Stuhlmüller, A., 2013. Knowledge and Implicature: Modeling Language Understanding as Social Cognition. *Top. Cogn. Sci.* 5, 173–184. doi:10.1111/tops.12007

6. Grice, P., 1975. Logic and Conversation, in: Cole, P., Morgan, J.. (Eds.), *Syntax and Semantics.* Academic Press, New York.

7. Horn, L.R., 2006. Implicature, in: *The Handbook of Pragmatics.* Blackwell, Malden, MA; Oxford.

Adrianna Smurzyńska
How Understanding Ourselves Is Related to Understanding Others?

Jagiellonian University & Polish Academy of Sciences

The classical other minds problem was the question: 'how can we know that there are other minds apart from our own?' One of the solutions of that problem was based on the analogical inference from the direct self-knowledge and similarities between me and other to the knowledge about other minds (Hyslop, 2016).

Nowadays, philosophers and cognitive scientists are rather poised to ask: 'how do we ascribe mental states to others?' Mental states attribution to others can be modelled in different ways. The most popular accounts are Theory Theory, Simulation Theory, or hybrids of both (e.g., Goldman, 2012; Newen, 2015), but alternative conceptions also exist (see: Gallagher, & Hutto, 2008). In contrast to classical theories of other minds, in most contemporary theories, there is no place for the privileged and direct access to one's own mind. Consequently, the selfknowledge cannot be regarded as the base for inferences about other minds.

The relation between the abilities to understand ourselves and to understand others is not well understood. Some research insists that understanding ourselves is prior (c.f. Goldman, 2006), some suggests that understanding others is more basic (c.f. Carruthers, 2009), and still other pieces of evidence suggests that these capabilities are independent of each other (c.f. Boyer, Robbins, & Jack, 2005).

One of the ways of exploring this relation is to analyse mental disorders in which mentalization is disturbed. This relation has been analysed mostly in autism and schizophrenia (c.f. Carruthers, 2009, et.al.). Nevertheless, in my opinion, it should be more deeply explored in personality disorders, which are mostly related to social functioning, self-image and other-image, therefore are undoubtedly connected to the processes of understanding ourselves and others. In my presentation, I will more deeply present the disturbances of mentalization in Borderline Personality Disorder (BPD). After presenting this disorder (DSM-V, 2013), I will appeal to the conception of mentalization introduced by Fonagy and Luyten (2009), in which this capability is presented on four dimensions: implicit-explicit, internal-external, cognitive-affective and self-other. I will describe mostly the relation between

self-mentalization and othermentalization in BPD (Fonagy, & Lyuten, 2010). I will also indicate why, in my opinion, this conception is more appropriate in exploring the relation between understanding ourselves and others than theory theory and simulation theory. I will also propose how this conception may be integrated with an enactive approach (especially De Jaegher, & Di Paolo) and may lead to the better description of mentalization.

In the conclusion I will propose the answer of the question: 'how understanding ourselves and others are related?' which may be stated as: 'these processes are interdependent and reciprocal (see: Damon, & Hart, 1982; Hodges, 2005), but neither of them can be treated as the underlying mechanism for the other.'

Kordula Świętorzecka
Bolzano's Metaphysics of Substances and Adherences as an Extension of the Theory of Abstract Objects

Cardinal Stefan Wyszyński University in Warsaw

Our research concerns a formal representation of Bolzano's original concepts of *Substanz* and *Adhärenz*. The formalism is built as an extension of Zalta's theory of abstract objects describing two types of predication, viz. *attribution* (to individuals) and *representation* (of properties). Bolzano was awared about this distinction. We focus on the consistency of this formalism and the description of its semantics. Firstly, we explore the possibility to reconstruct a Russellian antinomy based on the concept of the Bolzano's *Inbegriff* of all adherences. (Bolzano's theory of ideas is often suspected of antinomial consequences.) Our aim is to show limitations of his theory that prevent a contradiction when the *Inbegriff* consists of non-self-referential adherences. Next, we discuss two competing semantics for the proposed theory:: Scott's and Aczel's semantics. The first one yields a problematic result, that there are no models for the considered theory, containing a non-empty collection of all adherences. This is due to the fact that Scott's structures verify the formula on *reloading abstracts* in extensional contexts. We show that Aczel's semantics does not contain this difficulty. There are described Aczel's models with a non-empty set of all adherences. The self-referentiality of such a collection becomes irrelevant here. Finally, we show that there are Aczel's structures verifying the formula on reloading abstracts and we exclude them from the class of models intended for our theory.

Antonin Thuns
Polemical Concepts: Multiple Grounding and Virtual Grounding

Université Libre de Bruxelles

Polemical concepts are concepts whose content is still so heavily debated in the epistemic community, without any consensus emerging, that they suffer, as they stand, from a lack of semantic determinacy. Most philosophical and moral concepts are essentially polemical (concepts expressed by terms like 'subject', 'person', 'justice'). Concepts debated within the social sciences and linguistics are also often polemical (concepts expressed by terms like 'science', 'philosophy', 'religion', 'language'). Consider, to take but one example, the issue of the distinction between what counts as 'semantic' and what counts as 'pragmatic' in the philosophy of language. The amount of conflicting views available to draw the line is such that it could be argued, by a sceptical observer, that these alleged terms of the art do not, taken by themselves, correspond to any kind of determinate concepts.

A rather natural conclusion is that polemical concepts are simply not concepts, for they lack univocity. The idea that terms like 'science' or 'philosophy' designate genuine concepts would just be a lexical illusion. Since speakers in general do not converge on a determinate semantic content for these terms, each polemical term would involve a multiplicity of concepts under one (misleading) label. The problem with this view is that it seems to reduce polemicality to a form of semantic ambiguity. Yet a concept is polemical only inasmuch as it is thought that it should not be ambiguous. Speakers often disagree about what they take to be one and the same concept. Polemicality reflects a situation in which, on the one hand, there is some implicit agreement that there should be a stable concept behind a given term and, on the other hand, the epistemic community is unable to provide stable grounds for its use.

I argue that it is possible to make sense of the idea that polemical terms do express bona fide concepts by combining Devitt's notion of "multiple grounding" (Devitt 1981) with Schroeter & Schroeter's "tradition-based determination theory" of concepts (Schroeter & Schroeter 2014). One the one hand, polemical concepts are constantly regrounded in diverging ways so that they are subject to referential drift and their semantic import is generally unstable. On the other hand, the fact that speakers are collectively

engaged in using a term with the presumption that it corresponds to a valid concept gives plausibility to Schroeter and Schroeter's claim that it is the "entire representational tradition (*i.e.*, the entire set of token thought elements bound together by relations of apparent *de jure* sameness)" (2014: 12) that should be taken as the default semantic value for a polemical concept, rather than particular families of conflicting uses. Representational traditions are extended in the past, but they also point toward the future: future epistemic states of the community might mitigate or even eliminate the polemicality of some concepts. I conclude that it is possible to vindicate the idea of polemical concepts by introducing, alongside the notion of multiple grounding, a complementary notion of virtual grounding.

References

Devitt, M. (1981). *Designation.* New York: Columbia University Press.
Schroeter, L., Schroeter, F. (2014). Normative concepts: a connectedness model. *Philosophers' Imprint*, 14(25), 1-26.

Miroslava Trajkovski
The Constitution of Abstract Objects

University of Belgrade

Dummett claimed that the philosophical notion of an 'object' was introduced by Frege. The object is the reference of a singular term, so an abstract object would be the reference of an abstract singular term. The abstract singular term is a term that is built in a particular way which is given by the Principle of abstraction. Given that F and G are variables, and N is a termforming operator, E is an equivalence relation, the principle states:

(F)(G)(NF = NG iff E(F,G)).

Let us call the left-hand side abstrahendum and the right-hand side abstrahens. Hale and Wright argued that neither referential nor conceptual extension in the abstrahendum is conservative. The former is not relevant if the quantification is unrestricted. As to the conceptual non-conservative extension, Frege says that we „carve up the content in a way different from the original one [in the abstrahens], and this yields us a new concept [in the abstrahendum]." Namely, as he says, we replace a particular E with the more generic relation "=". For example, by such a replacement from the notion of equinumerousity we get the notion of the cardinal number of a set, or from the notion of parallelism of lines we get the notion of their direction.

I shall follow an approach that differs from Fregean, instead of taking an object to be defined as the referent of a term constructed as a singular term, I'll take that an object is the referent of a constitutional act. I start from Husserl's idea that logical objects are constituted by categorial acts. He distinguishes among three kinds of categorial constitutions: generalization, idealization and formalization. I will argue that new notions that lead to new names in the principle of abstraction are created through two steps of abstraction (not exactly Cantorian). I will argue that the fact that abstraction principles are both conceptually and referentially non-conservative is implied by these two steps.

Piotr Wilkin
Context Update, Dynamic Semantics and Continuations

Centre for Philosophical Research

One of the most prominent approaches to semantics that encompasses contexts is *dynamic semantics*, an approach originally formulated for the natural language by Groenendijk and Stokhof and closely mirrored by Kamp's *Discourse Representation Theory* and Heim's *file change semantics*. This approach borrows many elements from classic theories of programming language semantics, especially the denotational style of state-change semantics.

A common extension of the denotational semantics for programming language has been *continuation semantics* – an approach which allows one to model processing effects such as side effects of computations and divergent processing paths. This approach has been adopted for natural language much later, mostly thanks to works by Barker and Shan. However, the natural language usage of continuation semantics has been mostly limited to sentence-based effects typical for formal semantics. In classic Montague tradition, the end effect of a sentence's computation is considered to be a *truth value*, as opposed to a state-change function, or, as it's more often understood within the scope of dynamic semantics, a *context-change function*.

We will attempt to converge the two traditions stemming from computer science by attempting to adopt continuation semantics for natural language. For that, we will also consider a simplified notation for continuations as well as an alternative way of handling parallel computations of natural-language sentences, by considering the notion of an execution path. We will try to answer the question whether the added complexity level of continuations is worth the benefits in modeling side effects and non-linear processing paths.

Apart from trying to tackle the various problems associated with the dynamic-semantic way of using context-change functions as semantic values, we will try to propose a new approach that uses continuations (in the form of the abovementioned execution paths) to introduce *parse traces* – subsequences of full execution paths that deal with a select property of the context being changed. We will try to reconstruct some features of classical, static semantics (such as truth conditions) with the help of parse traces.

Bibliography

Heim, Irene, 1983a, "File Change Semantics and the Familiarity Theory of Definiteness", in *Meaning, Use and Interpretation of Language*, Rainer Bäuerle, Christoph Schwarze, and Arnim von Stechow (eds.), Berlin: De Gruyter, 164–189

Kamp, Hans and Uwe Reyle, 1993, *From Discourse to Logic*, Dordrecht: Kluwer

Barker, Chris and Shan, Chung-chieh, 2014, *Continuations and Natural Language*, Oxford University Press

Barker, Chris, 2002, "Continuations and the Nature of Quantification" in *Natural Language Semantics*, September 2002, Volume 10, Issue 3, Springer, 211-242

Yalcin, Seth, 2012, "Dynamic Semantics" in *The Routledge Companion to Philosophy of Language*, Routledge

Groenendijk, Jeroen and Stokhof, Martin, 1991, "Dynamic Predicate Logic" in *Linguistics and Philosophy*, 14(1), Springer, 39-100

Stoy, Joseph, 1981, *Denotational Semantics: The Scott-Strachey Approach to Programming Language Theory*, The MIT Press

Mateusz Włodarczyk
Processing Presuppositional Content of Utterances. Experimental Report

University of Szczecin

In my talk I will discuss the results of an experiment on the processing of presuppositions and conversational implicatures by the hearer. Test subject were presented with 36 short, written dialogues, each consisting of one sentence by person A and one sentence by person B. Each dialogue contained either presupposition trigger or carried a conversational implicature. After reading the dialogue test subjects were asked about information communicated on the implicit level. The results of the experiment show significant differences in both access time and availability of information communicated via presuppositions associated with the use of various triggers. I propose a model of processing presupposed content of utterances based on work of Craige Roberts (2015), in which two stages of presupposition processing - recognition and retrieval - are distinguished. I also suggest, following Mandy Simons (2013), that to retrieve the presuppositional content of an utterance the hearer must make inferences about beliefs of the speaker. However, I put forth a hypothesis that the hearer in the process of interpretation must also consider the context of an utterance – specifically, he should limit his inferences about a speaker's beliefs to those that are relevant to the topic of a current conversation. On this account, presuppositions associated with the use of implicative verbs are more context dependant than presuppositions triggered by different lexical elements – they are available to the hearer only in specific context. And although presuppositions associated with the use of definite descriptions, factive verbs and change of state verbs are more accessible to the hearer in the null context that was used in experimental material, we still observe variations of accessibility among those triggers. This suggest that processing of presuppositional content of an utterance is more complicated than it was originally assumed (e.g. Domaneschi 2016, p. 28).

References

Domaneschi, F. (2016). *Presuppositions and Cognitive Processes: Understanding the Information Taken for Granted.* Springer.

Roberts, C. (2015). Accommodation in a language game. *A Companion to David Lewis*, 345-366.

Simons, M. (2013). On the conversational basis of some presuppositions. In *Perspectives on linguistic pragmatics* (pp. 329-348). Springer International Publishing.

Jonathan Wright
Conceptually Permissive Attitudes

University of Southern California

Jan and Jim are at the party, trying to find Kat. They commonly know that Kit is unacquainted with Kat and that Kat is a musician. Jan knows that Kit believes that every musician at the party is in the pool, so she says to Jim: 'Kit believes that Kat is in the pool.' What she asserts is true even though Kit lacks a concept of Kat. Thus, one can truly ascribe an attitude in a proposition to a subject despite their lacking concepts of constituents of that proposition. I call reports of this kind "conceptually permissive".[1]

It is tempting to give an extensional specification of the conditions in which permissive reports occur in terms of conversational assumptions. Jan and Jim's common assumptions, together with the contents of Kit's loaded attitudes, entail the thought content Jan ascribes to Kit. Accordingly, we might say that one can truly ascribe a permissive attitude in proposition p to subject S in context c if the contents of S's corresponding loaded attitudes "contextually entail" p in c.[2]

Such a view faces an overgeneration problem. There are cases in which a subject's loaded attitudes contextually entail a proposition (on different restricted regimentations of the notion of contextual entailment), yet the speaker cannot truly ascribe a corresponding attitude in that proposition to the subject. As I show, to handle such cases an account of permissiveness must further take into account the *questions under discussion* (QUDs) in the context: whether a sentence can serve to make a permissive report depends on how its complement clause relates to the presuppositions of the context's QUDs.[3]

Adopting a question-based constraint on which contextual entailments ground permissive reports is plausible in its own right. For instance, if we treat transparent readings of noun phrases as a special case of permissive reporting, then the question-based account provides a unified explanation of various generalizations concerning transparency found in the literature.[4] The question-based account both predicts that these generalizations will

---

[1]Blumberg and Holguín (2017) call these "ultra-liberal" reports.
[2]This kind of theory has been suggested by van Frassen (1979) and Parikh (1998).
[3]Roberts (2017), Simons et al (2017).
[4]See Percus (2000), Keshet (2008).

hold when they do and also directs our attention to interesting exceptions to them.

Even if the question-based account is correct, a residual foundational issue concerns the *metaphysical* difference between permissive and non-permissive reports. Given the importance of QUD presuppositions, it is tempting to understand permissiveness as funda- mentally a presuppositional phenomenon.[5] I argue, however, that permissiveness is due to context-sensitivity in which relations are contributed by attitude verbs.[6] The relevance of QUDs to permissive reports lies in their systematic impact on pragmatic enrichment.[7] On my picture, then, attitude verbs are semantically underspecified along a particular dimension: they are conventionally constrained to contribute elements from a class of relations differing in conceptual permissiveness. Which relation is contributed on an occasion by an attitude verb is constitutively constrained by QUDs. This view further connects to long-standing debates in the theory of content. Contrary to certain enrichment-based defenses of externalist theories of content, pragmatic enrichment in uences not what one is said to believe, but which belief relation one is said to bear[8].

References

[1] Blumberg, K., & Holguín, B. (2017). Ultra-liberal attitude reports. *Philosophical Studies*, 1-20.

[2] Carston, R. (2008). Linguistic communication and the semantics/pragmatics distinction. *Synthese*, 165(3), 321-345.

[3] Dorr, C. (2014). Transparency and the context-sensitivity of attitude reports. *Empty representations: reference and non-existence*, 25-66.

[4] Keshet, E. (2008). Only the strong: Restricting situation variables. In *Semantics and Linguistic Theory*, 18, 483-495.

[5] Parikh, R. (1998). Propositions, Propositional Attitudes and Belief Revision. *Advances in modal logic*, 2, 381-400.

[6] Percus, O. (2000). Constraints on some other variables in syntax. *Natural Language Semantics*, 8(3), 173-229.

---

[5]Thus the QUD view would be an elaboration of Romoli and Sudo (2009)'s presuppositional account.

[6]So my QUD view is better seen as an elaboration of Dorr (2014)'s theory of transparent names.

[7]Carston (2008), Recanati (2010).

[8]For such a defenses, see Soames (2005).

[7] Recanati, F. (2010). *Truth-conditional pragmatics*. Oxford: Clarendon Press.

[8] Roberts, C. (2017). Speech acts in discourse context. In Fogal D., Harris D. & Moss M. (eds.) *New Work on Speech Acts*. Oxford University Press.

[9] Romoli, J., & Sudo, Y. (2009). De re/de dicto ambiguity and presupposition projection. In *Proceedings of Sinn und Bedeutung*,13, 425-438.

[10] Simons, M., Beaver, D., Roberts, C., & Tonhauser, J. (2017) The Best Question: Explaining the Projection Behavior of Factives, *Discourse Processes*, 54:3, 187-206

[11] Soames, S. (2005). Naming and asserting In Szabó, Z. (ed) *Semantics versus Pragmatics*, Oxford University Press.

[12] van Fraassen Bas, C. (1979). Propositional Attitudes in Weak Pragmatics. *Studia Logica: An International Journal for Symbolic Logic*, 38, 365-374.

Victor Verdejo
A Thought for You and I

University of Barcelona

Recent philosophical developments on personal indexicals reveal a disagreement between those who defend (Heal 2014, Haase 2014, Lea Salje 2017 or Lucy O'Brien 2007, 82-83) and those who deny (Heck 2002, Peacocke 2014a, 2014b, Chap. 10, Rödl 2007, 2014, Longworth 2013, 2014) the existence of a distinctive class of second person thoughts (distinctivists and non-distinctivists for short). In this piece, I tackle this controversy by highlighting two crucial constraints based on paradigmatic felicitous singular uses of the second person pronoun. On the one hand, the Addressing Constraint is brought out by the awareness and action capabilities displayed in successfully addressing another (cf. Peacocke 2014a, 2014b, Martin 2014 or Salje 2017):

> **(Addressing Constraint)**
> The thinker expressing a 'you'-thought must intend to refer to another thinker while standing and being aware of standing in a relation of mutual awareness in ways that enable behavioural interaction between them.

On the other hand, the Merging Constraint arises, among other things, from the fact that 'I'/'you'-exchanges ground intersubjective disagreement (cf. Bermúdez 2017, Chap. 4):

> **(Merging Constraint)**
> In central communicative exchanges, utterances of 'I'/'you' express the very same thought.

Once these constraints are fully in view, I go on to show that they pose a challenging dilemma for any account of the second person and that the chasm between friends and foes of the distinctness of second person thought is better seen as endorsements of one of the horns of the dilemma. While some versions of non-distinctivism may offer limited accounts of the Merging Constraint, they fail to live up to the Addressing Constraint. Correspondingly, whereas distinctivists devise their approaches to adequately deal with the

Addressing Constraint, they definitely founder when trying to accommodate the Merging Constraint.

In reaction to this, I outline a way of accommodating both constraints in terms of 'perspectives', i.e. ways of thinking of a reference that do not individuate a thought (Verdejo (2017), Verdejo (manuscript)). On the recommended approach, the second person perspective is cognitively distinctive but does not itself signal the existence of a distinctive second person type of thought. By contrast, a single type of thought about selves expressible with personal indexicals—the self-thought type—can be shown to comprise the third, the second and the first person perspective. Perspectives so understood involve drawing on information about the relation between the ongoing event of thinking and the event of thinking whose thinker is the reference of the thought. In addition, the proposed analysis deals satisfactorily with the pretended- and failed addressee problems originally pointed out by David Kaplan (1989) given that it allows preservation of reference to a subject with 'you' even in cases in which the Addressing Constraint is not met.

The initial clash between distinctivists and non-distinctivists can thus be seen as nourishing a unifying and explanatorily rich account of thought expressible with personal indexicals. The resulting picture is one in which a distinctive second person type of thought is ruled out but so is, perhaps surprisingly, a distinctive type of first person thought.

References

Bermúdez, J.L. (2017) Understanding "I". Oxford: Oxford University Press.

Haase, M. (2014) 'Am I You?', Philosophical Explorations 17, 358-371

Heal, J. (2014) 'Second Person Thought', Philosophical Explorations 17, 317-331.

Heck, R.G. (2002) 'Do Demonstratives Have Senses?', Philosopher's Imprint 2, 1-33.

Kaplan, D. (1989) 'Afterthoughts' in J. Almog, J. Perry, and H. Wettstein (eds.), Themes from Kaplan. Oxford: Oxford University Press, 565–614.

Longworth, G. (2013) 'Sharing Thoughts about Oneself ', Proceedings of the Aristotelian Society 113, 57-81.

Longworth, G. (2014). 'You and Me', Philosophical Explorations 17, 289–303.

Martin, M.G.F. (2014) 'In the Eye of Another: Comments on Christopher Peacocke's "Interpersonal Self-consciousness"', Philosophical Studies 170, 25-38.

O'Brien, L.F. (2007) Self-Knowing Agents. Oxford: Oxford University Press.

Peacocke, C. (2008) Truly Understood. Oxford: Oxford University Press.

Peacocke, C. (2014a) 'Interpersonal Self-Consciousness', Philosophical Studies 170, 1-24.

Peacocke, C. (2014b) The Mirror of the World. Oxford: Oxford University Press.

Rödl, S. (2007) Self-Consciousness. Cambridge, MA: Harvard University Press.

Rödl, S. (2014) 'Intentional Transaction', Philosophical Explorations 17, 304–316.

Salje, L. (2017) 'Thinking About You', Mind 126, 817-840.

Verdejo, V.M. (2017) "De Se Contents and Action Generalisation" Philosophical Papers 46, 315-344.

Verdejo, V.M. (manuscript) "Thought Sharing, Communication, and Perspectives about the Self".

Alejandro Vesga
Is Testimony Strictly Semantic?

Cornell University

According to Elizabeth Fricker (2012) and Sanford Goldberg (2007; 2009), testimonial knowledge can only be transmitted through the semantically defined content of the speaker's uttered sentence. Messages inferentially communicated beyond this content fail to be testimony, in the sense that they fail to enable knowledge transmission. In this paper, I argue against this view, which I call the Semantic Primacy Thesis (SPT). Against Fricker and Goldberg, I claim that speakers can testify for the truth of non-semantically defined contents.

I begin by sketching what I take to be Fricker's and Goldberg's main argument for SPT. This argument relies on the notions of testimonial responsibility and blame. As the argument goes, speakers have testimonial responsibility for the truth of what they testify about and, therefore, are liable to testimonial blame if they offer defective testimony. According to Fricker and Goldberg, our practices involving testimonial responsibility are justified only when testimony is restricted to semantically defined contents. For non-semantically defined contents cannot be objectively "pinned down" in order to warrant blame judgments for their falsity. In particular, non-semantically defined contents are excessively dependent on the speaker's subjective psychological states and can thus be incontrovertibly denied by her. Consequently, speakers can always validly avoid testimonial blame for the falsity of the non-semantically defined contents they convey.

Having presented the main argument for SPT, I proceed to argue against this view. My argument has two parts, one negative and one positive. In the negative part I argue that SPT fails to capture our testimonial blaming practices. As we shall see, neither Fricker nor Goldberg offers a precise definition of the semantic restrictions that determine the contents for which speakers are testimonially responsible. For this reason, SPT can be formulated in a variety of ways. I consider three plausible interpretations of SPT, each less restrictive than the last, and offer cases in which the contents for which speakers are responsible exceed the semantic restrictions posited by each of these interpretations. The basic motivation for SPT fails: even assuming a rather permissive notion of semantically-defined content, speakers

can—and should—be blamed for the falsity of kinds of non-semantically defined contents.

In the positive part of my argument I suggest an alternative justification standard for testimonial responsibility. Given the practical constraints entailed by the interpersonal and cooperative nature of conversation, engaged speakers are publicly committed to conversational success. Every cooperative play in a conversational game is subject to public and interpersonal evaluability. Therefore, in a conversation aimed to inform, testimonial blame can be justified by appealing to the cooperative role the speaker's utterance played, given the conversational context. Speakers can be blamed for not being sufficiently cooperative as the conversational state requires or for communicating false contents through cooperative utterances. In the latter case, even if such play made use of a non-literal mechanism to convey a specific content, it played a testimonial role in conversation for which the speaker is responsible.

Further, as Elizabeth Camp argues, cooperation in conversation need not be of constant value. Not every conversation is acknowledged as fully cooperative. Cooperation in conversation can be compromised by specific conditions surrounding it—it may be embedded in competitive contexts or tainted by socially embedded power dynamics, for instance. In contexts where full-on cooperation is not evident, there may be reason to semantically restrict the utterances that can be deemed as valid testimony. Salient cooperation can thus be considered to *modulate* the extent to which certain non-semantically defined messages can be considered testimony. Clearly cooperative environments may enable testimonial communication beyond literal means, while suspicion of less-than-perfect cooperation can make us move our threshold of acceptable testimony towards literal contents.

José Manuel Viejo
How to (Dis)solve Frege's Paradox About Singular Terms and Belief Ascriptions

Universidad Autónoma de Madrid

I shall do four things in this paper. First, I will present the problem of belief reports (and other attitude attributions) that has puzzled philosophers of language for over a century. The difficulty (which, following current practice, I will call 'Frege's Paradox') is that it seems impossible to accommodate the thesis of Opacity of Belief Reports while respecting the following three plausible and widely held semantic principles: Direct Reference, Compositionality, and Semantic Innocence. Against the classical line of response to the problem of belief reports, I will defend that there is no logical incompatibility between these four theses. While these principles seem inconsistent, this is because we are presupposing a Traditional Analysis of Belief Reports, which has in effect two parts: a dyadic relational conception of propositional attitude verbs and a standard account of that-clauses.

Second, I will examine the so-called hidden indexical theory, which allows a subtle analysis of propositional attitude ascriptions that seems to show how the four thesis mentioned above can be compatible. However, this promising analysis is subject to outstanding challenges: the candidate problem, the logical form problem and the meaning-intention problem.

Third, I will show how a certain revision of the hidden-indexical theory which accepts the dyadic conception of the belief relation but denies the standard analysis for that-clauses, avoids these problems. In this model, that-clauses are indexical terms, having a rule that describes the referent as being similar (in contextually appropriate ways) to the proposition expressed by the embedded sentence. I shall close proposing this theory as being the best theory of the semantics of belief reports relative to the assumption that there is no inconsistency between Opacity and the three semantic principles and suggesting the standard analysis for that-clauses is the source of the paradoxes about belief ascriptions.

Finally, I will argue that some enduring problems in the philosophy of mind and language appear to be recalcitrant in consequence of ignoring the context-sensitivity of that-clauses, problems that are otherwise resolvable. I will focus here in the de dicto/ de re distinction and Kripke's puzzle about belief (ascription). I will suggest that traditional attempts to solve these

problems are unsuccessful whereas the indexical theory succeeds. This should be considered as additional motivation for the theory.

Julia Zakkou
Presupposing Counterfactuality

Universität Hamburg & Freie Universität Berlin

Current consensus holds that counterfactuals do not presuppose the falsity of their antecedents; if the speaker of a counterfactual conveys the antecedent's falsity at all, she conversationally implicates it. Most philosophers and linguistics take this to be obvious for *non-past subjunctive conditionals* (NPSP); but many of them also think this is the case for *past subjunctive conditionals* (PSC). In this paper, I'll argue that the the most prominent argument for the claim that not even PSCs presuppose the falsity of their antecedents (due to Anderson, 1951) is inconclusive.

Consider the following utterances as made by a doctor, A, at the autopsy of Jones's body:

(1) If Jones had taken arsenic, he would have shown the same symptoms he actually shows. So he did not take arsenic.

(2) If Jones had taken arsenic, he would have shown the same symptoms he actually shows. So he took arsenic.

Intuitively, (1) is bad and (2) is fine. Given the presupposition hypothesis (PH), though, it should be the other way around – or so Anderson assumes.

I'll object that for (1) and (2) to be fine, two conditions have to be met: (i) A is not voicing beliefs that are in tension. (ii) A is providing a good reason for the claim made by the second sentence. After all, by introducing the second sentence with 'so,' A triggers the expectation that the second sentence is the conclusion of some kind of argument and that, what she said before, provides good reasons for the conclusion.

Condition (ii) is met in (2) but not in (1). With her utterance of the PSC in question, A is conveying that a certain relevant, necessary condition for Jones's having taken arsenic is met. This, though, only provides a good reason for the claim that Jones *did* take arsenic. It does not provide a good reason for the claim that Jones *did not* take arsenic. So PH can explain why (1) sounds bad after all.

Condition (i) is met in (2) (as well as in (1)). It is well known that even though people normally presuppose things they hold true they sometimes also presuppose things they know to be false. Trying to deceive is one reason

for this, but there are others (see Stalnaker, 2002: 176). I contend that this is what we find in (2): by presupposing that Jones did not take arsenic, A signals that she is accepting this proposition for the time being; by asserting that Jones took arsenic, she is making clear that she herself believes the exact opposite. Given this, our condition (ii) is met: in case (2), A is not voicing beliefs that are in tension. PH can thus explain why (2) sounds fine.

Note that it should not come as a surprise that such cases – cases in which a speaker felicitously first presupposes that p and then immediately afterwards asserts that not-p – are possible. If speakers can presuppose things they do not actually believe, why shouldn't they be able to disclose to their audience that this is what they did? Of course, sometimes, revealing that they do not believe what they just presupposed will be counter-productive But, often enough, disclosing their true opinion on the matter in question will a perfectly natural move (think of cases where they grant something for the purposes of an argument).

I'll defend my contention by presenting cases that work analogously (containing definite descriptions and factive/veridical verbs like 'knows' and 'sees') and by pointing out that my suggestion is anything but *ad hoc*, more concretely that it is very similar to what Karttunen and Peters (1979: 5) and more recently Leahy (2017: 17) have said in defense of the claim that indicative conditionals presuppose the epistemic possibility of their antecedents.

Dan Zeman
A Rich-Lexicon Approach to Evaluative Expressions

University of Vienna

Theories of lexical meaning can be divided into two main types: rich-lexicon theories and sparse-lexicon theories. According to the latter, the lexical entries of natural language expressions encode "thin" concepts that are enriched during the compositional process; according to the former, they encode "rich" concepts, with several dimensions of meaning that get selected during the compositional process.

My aim in this presentation is not to enter the debate between the two types of theories of lexical meaning described, but to present one way in which a rich-lexicon approach can be applied to evaluative expressions (by which I mean expressions that contain an evaluative component). While there are many rich-lexicon theories in the literature, I will use the one put forward by Del Pinal (2017) and Del Pinal & Reuter (2016), according to which the meaning of words can be represented by a tuple consisting in an extension and a conceptual structure. The conceptual structure comprises a multitude of meaning dimensions; in the case of common nouns, these include perceptual information about the object referred to (PERCEPTUAL), information about what it is are made of or its parts (CONSTITUTIVE), how it came to being or the purpose of its creation (AGENTIVE), its typical function (TELIC) etc. While the full (literal) meaning of a word comprises the conceptual representation with all its dimensions, particular uses of a word highlight certain dimensions in the detriment of others. Thus, in a sentence like "This book is heavy", the CONSTITUTIVE dimension is foregrounded, while in a sentence like "This book is interesting", a different dimension is foregrounded. This approach explains how expressions can get different interpretations in different contexts, while preserving their full (literal) meaning.

With this sketch in hand, I contend that this theory has good chances to apply to a wide range of evaluative expressions and to handle a variety of ways in which such expressions are used. Take, first, slurs: expressions like "faggot", "n*****", "gypsy" etc. As widely acknowledged, slurs are linguistic means to derogate individuals by conveying a negative attitude towards the group they belong to. The main challenge for theories of slurs has been to capture their derogatory (i.e., evaluative) character. A rich-lexicon theory of

the type sketched above can capture that by introducing an EVALUATIVE dimension in the concepts such expressions encode, alongside the ordinary dimensions mentioned above. However, the theory also accounts for various non-derogatory uses of slurs (e.g., didactical, appropriated, self-identificatory) by allowing the EVALUATIVE dimension to be backgrounded and the other dimensions to be foregrounded.

A second class of evaluative terms that a rich-lexicon theory is well suited to apply to are normative expressions such as moral terms ("good", "bad"), aesthetic adjectives ("beautiful", "ugly"), predicate of taste ("tasty", "disgusting"), epistemic terms ("known", "justified") and terms that involve other kinds of normativity. Presumably, by using such terms one expresses a certain attitude, but also conveys to the audience that they should adopt that attitude too. A rich-lexicon theory would claim that besides the ordinary dimensions, such expressions encode concepts that have both an EVALUATIVE dimension and a NORMATIVE dimension. However, the theory also accounts for various non-normative uses of such expressions by allowing the NORMATIVE dimension to be backgrounded and the other dimensions to be foregrounded.

A nice feature of the theory proposed is that an account of a different type of evaluative expressions falls out automatically: that of expressives such as "damn", "shit" etc. The main characteristic of such expressions is that they are used to simply express a certain attitude. This is captured in the rich-lexicon theory by making the EVALUATIVE dimension the central one in the concept encoded by such expressions, while leaving very little or no role to play for the other dimensions. A rich-lexicon approach to evaluative expressions thus yields a unitary and flexible framework able to handle a wide variety of uses of a great range of natural language expressions.

Adrian Ziółkowski
Survey Pragmatics: Communication Between the Researcher and the Participant in Experimental Philosophy

University of Warsaw

Experimental philosophy is a young movement within analytic tradition that introduces novel, empirical methods to the philosophers' toolbox. Experimental philosophers analyze philosophical intuitions of the folk and try to draw philosophically significant conclusions from the collected data (cf. Knobe & Nichols, 2008). The most popular method of data gathering is known as "the method of cases", which consists in presenting subjects with hypothetical scenarios based on famous philosophical thought experiments and asking subjects questions similar to those addressed by the authors of these thought experiments.

Experimental philosophers have collected various data suggesting that intuitions expressed by non-philosophers are often unstable and sensitive to insignificant factors, such as cultural background, the order of scenario presentation, or subtle differences in phrasing of crucial questions. In my talk I will focus on that latter phenomena and present a theoretical account that aims at providing a systematic explanation of them by referring to pragmatics and context-dependence.

Take the famous Gettier cases in epistemology, which are believed to describe situations in which a certain agent has a justified belief that p without knowing that p. Cullen (2010) found that non-philosophers react to Gettier cases differently depending on subtle differences in phrasing of the knowledge-question: laypeople are more likely to attribute knowledge when they choose between options "A knows that p" and "A doesn't know that p" than when the answer options are "A really knows that p" and "A only believes that p". Similarly, Swain et al. (2008) found that non-philosophers' intuitions about knowledge for another kind of cases are affected by priming: when we precede the presentation of the target scenario with a clear case of knowledge, laypersons will be less happy to attribute knowledge in the target scenario than when the prime is a clear case of non-knowledge. Basing on such results, many philosophers concluded that folk philosophical intuitions are unreliable and cannot be trusted.

Survey Pragmatics (e.g. Schwartz 1996) is a theoretical approach that offers an alternative perspective on phenomena such as those described above by suggesting that taking a survey is a situation of bilateral communication between the researcher(s) and the participant(s), and as such, is constrained by the pragmatic rules we impose on other kinds of conversations. In this vein, one might explain many apparently problematic results as instances of pragmatic interactions, such as adherence to conversational maxims or seeking for conversational implicatures (cf. Grice, 1975). For instance, the above mentioned priming effect might result from the participants' adherence to the maxim of relevance - they might have assumed that, if the researchers preceded the presentation of the crucial case with some other scenario, the contrast must be relevant and should be taken into account when assessing the following case.

In my talk, I will present various data on folk intuitions in the fields of epistemology and philosophy of language (collected by me and other researchers) that seem to show the instability and unreliability of folk intuitions. I will argue that in many cases, the data is only apparently problematic, since it can be explained by referring to systematic, pragmatic phenomena in the researcher-participant communication. I will also argue that neglecting context-dependence of survey materials is often the source of misinterpreting the data collected by experimental philosophers.

References

Cullen, S. (2010) Survey-Driven Romanticism, *Review of Philosophy and Psychology* 1 (2): 275-296.

Grice H.P. (1975) Logic and Conversation [in:] P. Cole, J. Morgan (ed.) *Syntax and Semantics*, vol.3, New York: Academic Press.

Knobe, J., Nichols, S. (2008) An Experimental Philosophy Manifesto [in:] J. Knobe, S. Nichols (ed.) *Experimental Philosophy*, vol. 1, Oxford: Oxford University Press.

Schwarz, N. (1996) *Cognition and communication: Judgmental biases, research methods, and the logic of conversation.* Mahwah: Erlbaum.

Swain, S., Alexander, J., Weinberg, J. M. (2008) The instability of philosophical intuitions: Running hot and cold on truetemp, *Philosophy and Phenomenological Research* 76 (1), s. 138–55.

Tomasz Zyglewicz & Filip Kawczyński
'Just More Theory' Revisited

University of Warsaw

In his recent book Simchen (2017) argues against Sider's (2011) claim that reference magnetism should be taken to be *the* metametasemantic theory, in the sense that it incorporates all possible metasemantic accounts. In our paper we criticize Simchen's point.

In the first part of the paper we reveal flaws in Simchen's argument by demonstrating that he equivocates two meanings of the predicate 'is true' that we provisionally dubbed the *factual* and the *semantic* meaning. He argues that incorporating a given metasemantics into magnetism results in diminishing the explanatory value of that theory, because in such a case the metasemantic account is nothing more than 'just more theory' added to the corpus of magnetism. To support his point he shows that a metasemantic sentence S and its 'semantic ascent', i.e. S *is true* are not explanatorily equivalent, while they should be if the metasemantic theory that includes S is supposed to be 'just more theory'. We argue that if 'is true' is taken with its *factual* meaning, then S and S *is true* are explanatorily equivalent indeed. On the other hand, when 'is true' is taken with its *semantic* meaning, then if fact they are not explanatorily equivalent, but it still does not pose a problem, because S *is true* is the metametasemantic sentence that has passed through the 'semantic ascent' and thus the notion of the truth it involves is not the same as in the sentence S which is the metasemantic statement.

In the second part we put forward an alternative way of capturing the explanatory value of 'just more theory' that provides additional reasons for embracing Sider's claim. Following Dorr (2016) we assume that identifications of the form "To be F is to be G" constitute a distinctive mode of metaphysical explanation and that they express an equivalence relation. Granting the truth of causal theory of reference, its contribution to our general metasemantic picture can be captured by claims along the following lines:

(1) For "X" to be joint-carving is for "X" to be caused by X.

Now, due to the symmetry of the identification, the following is also true:

(2) For "X" to be caused by X is for "X" to be joint-carving.

However, we argue that (2), as opposed to (1), is not a good explanation. Our reasons have to do with pragmatics of philosophical discourse about metasemantics. In particular, (2) does not seem to be an adequate answer to:

(3) What is it for "X" to be caused by X?

Moreover, the generalization of (2) cannot be squared with there being ways of joint-carving other than causation (e.g. being selected for by evolutionary processes). If our argument is correct, it preserves the elevated position of reference magnetism, while recognizing explanatory work to be done by other metasemantic theories.

References

Dorr, C. (2016). To be F is to be G. *Philosophical Perspectives* 30 (1): 39-134.
Sider, T. (2011). *Writing the Book of the World*: OUP.
Simchen, O. (2017). *Semantics, Metasemantics, Aboutness*: OUP.

POST-CONFERENCE WORKSHOP

Edward N. Zalta
Typed Object Theory

Stanford University

In this talk, I first review some history of relational type theory and functional type theory, and focus on the fact that most historical formulations of relational type theory interpret relations as functions. Typed object theory (Zalta 1983, 1988) is different: it assumes that the domains of higher-order entities are genuine (primitive) relations.

The main features of type-theoretic object theory are: (a) the domain of each type t divides up into ordinary objects of type t and abstract objects of type t, and (b) when representing natural language, (i) the denotation of a natural language term of type t may be either an ordinary or an abstract object of that type, whereas (ii) the sense of a natural language term of type t is an abstract object of that type. As a result, the intension of a natural language term has the very same logical type as its extension – there is no type-raising involved.

Given these features, we review the simple solutions to the following issues and problems:

- the informativeness of "woodchuck = groundhog" identity claims and the paradox of analysis
- modeling Fregean senses and the analysis of attitude reports
- the analysis of fictional properties (hobbits, unicorns, etc.) and fictional relations (absolute simultaneity)
- the analysis of mathematical objects and relations

As part of the discussion, comparisons to other theories are developed, by first noting that type-theoretic object theory doesn't require a primitive type for truth-values (as in functional type theory), doesn't require a primitive type for possible worlds (as in Montague's intensional logic), doesn't collapse the types for individuals and propositions (as suggested by Partee, Liefke, and Liefke & Werning); and uses only one set of types for both its syntax and semantics (unlike Williamson 2013, which adds to the semantics a primitive type for worlds and and an infinite type for finite expressions). Other interesting differences with these and other theories are explored as well.

Arvid Båve
Structured Propositions and Alternative Analyses

Stockholm University

The phenomenon of alternative analyses is, very roughly, the fact that "Rab" can be analysed as the result of (i) saturating "Rax" with "b", (ii) saturating "Rxb" with "a" and (iii) saturating "Rxy" with "a" and "b". Drawing on an argument by Dummett, I argue that one must posit alternative analyses if one is to state adequately general inference rules. I exemplify with the law of substitution of identicals. But the phenomenon of alternative analyses raises serious difficulties for many accounts of structured propositions. For instance, it dooms the act type theory of propositions, which holds that propositions are act types directed toward propositional constituents. I present a general objection, targeting at once Russellian and Fregean act-type theories, which merely assumes that some sentences have alternative analyses plus certain intuitive identity conditions on act types, namely:

(A1) If the act type of ϕ-ing x_1, ..., x_n = the act type of ϕ-ing y_1, ..., y_n, then at least one of x_1, ..., x_n is identical to one of y_1, ..., y_n.

The upshot is that the theory entails that any sentence with alternative analyses (in fact, any atomic sentence with a polyadic predicate) is ambiguous, an unacceptable conclusion. To concretize, a Russellian act-type theorist like Soames is committed to saying that "Mary loves John" expresses both the act type of predicating the property of loving John of Mary and the act type of predicating the property of being loved by Mary of John. But by (A1), these are distinct, hence, the sentence expresses one proposition relative to one "analysis" and a distinct proposition relative to another. I also discuss various possible retreat positions for act-type theorists and argue that they all come with considerable problems and are all very different from extant act-type theories.

Jeffrey King's theory of propositions is similarly shown to have this consequence, although by a somewhat different route. On his view, the proposition that Fa is the fact that a stands in a certain complex relation REL with F-ness. He must thus take "Rab" to express both F1 = the fact that a stands in REL to PROPx(Rxb) and F2 = the fact that b stands in REL to PROPx(Rax). However, I argue that, given his definition of REL, he

is committed to saying that F1 $\neq$ F2, which means that he must take "Rab" to be ambiguous, expressing different things depending on the analysis at hand. This result follows both on King's actual definition and on a definition on which "language" is taken in the Lewisian, abstract sense, rather than in King's sense of "actual, spoken language".

Finally, I state a general constraint on theories of structured propositions and discuss its more general implications. The constraints reads,

(GC) $[F(a)] = X([F(\xi)], [a])$ only if X does not satisfy an (A1)-principle.

A particular difficulty is pointed out, which is that of avoiding the objection by retreating to the claim that propositions are equivalence classes while defining equivalence in a sufficiently fine-grained way, yet without making sentences with alternative analyses come out as ambiguous.

To apply the condition to Zalta's theory, which refers to propositions using a kind of gerund, we can see that it meets the constraint only if Mary's loving John is identical with John's being loved by Mary (and so on). Since this seems plausible enough, his theory survives the objection (although other aspects of it are arguably more worrisome).

References

Dummett, M. (1981), *The Interpretation of Frege's Philosophy*, Cambridge, MA., Harvard University Press.

Geach, P. (1975), "Names and Identity", in S. Guttenplan (ed.), *Mind and Language*, Oxford, Clarendon Press, 139-158.

Hanks, P. (2011), 'Structured Propositions as Types', Mind 120, 11-52.

King, J. (2007), *The Nature and Structure of Content*, Oxford, OUP.

King, J., S. Soames, and J. Speaks (eds.) (2013), *New Thinking About Propositions*, Oxford, OUP.

Soames, S. (2010), *What is Meaning?*, Princeton, NJ., Princeton University Press.

Soames, S. (2015), *Rethinking Language, Mind, and Meaning*, Princeton, NJ., Princeton University Press.

Zalta, E. (1988), *Intensional Logic and the Metaphysics of Intentionality*, Cambridge, MA., MIT Press.

Arkadiusz Chrudzimski
Encoding and Varieties of Intentionality

University of Szczecin

In the history of philosophy one can find many theories of intentionality introducing special entities under the label "intentional objects". Franz Brentano's immanent objects, Alexius Meinong's entities "beyond being and non-being", or Roman Ingarden's purely intentional objects can serve as examples of such entities. What they all have in common is that they have been introduced in order to extensionalise the so called "intentional contexts". But not all entities which function this way deserve the name of intentional objects. In particular, neither Frege's senses nor mental contents of the early Husserl are to be classified as intentional objects in this sense. Roughly speaking, to be properly called "an intentional object" a postulated entity must be supposed to function as a quasi-target of the subject's intention. In other words: intentional objects are supposed to stand "before the subject's mind", so that they, in a sense, "replace" the common sense objects of reference.

Some of the mentioned above philosophers (like e.g. Brentano and Ingarden) saw intentional objects as entities involving a peculiar non-standard kind of exemplification. The contemporary incarnation of this idea is Zalta's theory of abstract objects and his concept of encoding. From the point of view of ontological parsimony this seems to be the best way of construing intentional objects.

Nowadays philosophers who are sympathetic to intentional objects are accused of planting an ontological jungle. All the problems of the theory of intentionality, it is claimed, can be resolved within the framework of a theory assuming a much more parsimonious ontology, like the theory of mental content proposed by the early Husserl or the so-called "adverbial" theory of intentionality. However, it turns out that the competitors of the theory of intentional objects face serious difficulties, the most important being that within their framework the relation between the representing entity (mental content or "adverbially specified" mental property of the subject) and the external target object has to be construed as primitive, while in the theory of intentional objects it can be easily defined. The consequence is that the partisans of mental contents and adverbialists are forced to require

a distinguished kind of epistemic access not only to the representing entity but also to this "representing relation".

Moreover, it turns out that we gain certain significant benefits from the introduction of intentional objects. It seems that with help of them we are able to explain some traditional philosophical puzzles that are hard to treat without them, like e.g. the problem of qualia and the problem of ontological negativity.

Andreas De Jong
"There is no Santa Claus": On the Meinongian Treatment of Negative Existentials and Negative There Be-Sentences

University of Manchester

In *The Nonexistent* (2013, pp. 144-148), Anthony Everett points out that any fictional realist account has to explain the intuition that negative existentials are true, even though fictional realists claim that there are such things. A Meinongian account of negative existentials has a straightforward answer to that question: we have those intuitions because negative existentials are true. More specifically, negative existential claim *de re* of the objects that they do not exist. The Meinongian line of defence is that this account of negative existentials is descriptive of the semantics of ordinary people, as is defended by Reimer (2001; 2001a). Everett levels two criticisms against this line of defence, which he says has "a rather ad hoc flavour and would need careful justification and defense" (2013, p. 144).

The first criticism is that the Meinongian cannot account for why explicit anti-Meinongians do not feel pressure to deny negative existentials pertaining to X, whilst they do feel this pressure in the case of first-order truths about X. The second criticism is that even if the Meinongian can accept negative existentials, she cannot explain why we also use Negative Existential There-be sentences (NETBs, of the form "There is (are) no X(s)") to express negative existentials. To support the second criticism, Everett maintains that we should interpret the quantification in NETBs as unrestricted in the absence of obvious contextual cues. Everett dismisses the Meinongian response that NETBs without contextual cues are restricted to nonexistent objects. We take him to issue two challenges. The *Why Restricted Challenge* presses us to answer why NETBs are restricted to existent objects as opposed to being unrestricted. The *Why Existence Challenge* presses us to explain why NETBs are restricted to existent objects as opposed to another restriction, say *concrete objects*.

Against the first criticism, we will argue that the disparity that Everett claims is exaggerated. The existence of antirealist strategies (expressivism and fictionalism) shows that not all antirealists about a subject matter have felt the need to become error theorists about that subject matter. However,

Everett could argue that those attitudes result from pressure to adopt an alternative semantic construal of first-order truths about X. Therefore, we showcase respectable philosophers that assert first-order truths about X while being antirealist about Xs. Therefore, we maintain that former claim of the disparity is not as severe as Everett claims. In addition, the historical treatment of negative existentials shows that anti-Meinongians feel the pressure to deny negative existentials on their original semantic treatment. Everett may want the Meinongian to explain why the Anti-Meinongians have opted to revise the semantic treatment of negative existentials, rather than become Meinongian. The only way to do that is to attribute a mistake to the Anti-Meinongian or to consider the Anti-Meinongian irrational.

Against the second criticism, we will argue that both challenges can be met. For meeting the *Why Restricted Challenge*, we cite NETBs from Moltmann (2013; 2015) that quantify over nonexistent objects, such as "there are books in this catalogue that do not exist". These examples suggest that the context free NETB "there are no buildings that do not exist" should be interpreted as restricted.

For meeting the *Why Existence Challenge*, we point out that, given Meinongian characterization principles (e.g. Priest 2016 [2005], p. 83), unrestricted NETBs are trivially false. Hence, we need a restriction that makes NETBs nontrivial. A speculative answer to the contrastive question posed by this challenge, we cite Meinong's prejudice for the actual (1960 [1904], p. 79). Pressed for a reason, we maintain that a restriction to real objects, as opposed to objects of make-believe, seems to be a reasonable. We use a pretence-theoretical characterization of nonexistent objects: i.e. as generated by stipulated rules of a game of make-believe (Walton; Evans 1982). Such objects are subject to the aims of those that produce the will and their behavior can be dictated. As opposed to those objects, existent objects do not adapt as easily to the will of sentient beings. This is because we cannot dictate the behavior of these objects, for they behave independently from our stipulations. That makes them worthy of study, because we have to adapt to their behavior.

References

Evans, Garreth (1982). *The Varieties of Reference*. Oxford: Clarendon.
Everett, Anthony (2013). *The Nonexistent*. Oxford: Oxford University Press.
Meinong, Alexius (1960 [1904]). "The Theory of Objects". In: *Realism and the Background of Phenomenology*. Ed. by R.M. Chisholm. Trans. by I. Levi,

D.B. Terrell, and R.M. Chisholm. Glencoe: Free Press.

Moltmann, Friederike (2013). "The Semantics of Existence". In: Linguistics and Philosophy 36.1, pp. 31– 63.

— (2015). "Quantification with Intentional and Intensional Verbs". In: *Quantifiers, Quantifiers, Quantifiers: Themes in Logic, Metaphysics, and Language.* Ed. by Alessandro Torza. Dordrecht: Springer, pp. 141–168.

Priest, Graham (2016 [2005]). *Towards Nonbeing.* Oxford: Oxford University Press.

Reimer, Marga (2001). "A "Meinongian" solution to a Millian Problem". In: American Philosophical Quarterly 38, pp. 233–248.

— (2001a) "The problem of empty names" In: Australasian Journal of Philosophy 79, pp. 491–506.

Daniel Milne-Plückebaum
Meinongian Modal Meinongianism

Universität Bielefeld

One can distinguish between three kinds of realism about intentionalia (Sainsbury, 2010, 23). According to *Nonactualism* about intentionalia, intentionalia are existent and concrete, but nonactual; according to *Abstractualism* about intentionalia, intentionalia are existent and actual, but nonconcrete; and according to contemporary *Meinongianism* about intentionalia (inspired by Meinong, 1960), intentionalia are concrete and actual, but nonexistent.

One version of Meinongianism about intentionalia is *Modal Meinongianism* (MM) (Priest, 2016). According to MM, any intentional object *o* instantiates the properties *o* is characterised as having—perhaps not in the actual world (@), though, but in those worlds that realise one's imagination about *o*. For example, according to MM, the golden mountain instantiates goldenness and mountainhood—thought not in @, but only in those worlds that realise one's imagination about the golden mountain. But then, *prima facie*, MM is in conflict with the fact that, according to Meinongianism, intentionalia are supposed to be *actual* (Barz, 2016).

Given this apparent conflict, I argue that while MM does have it that intentionalia are just as characterised in worlds distinct from @, this doesn't amount to their being just as characterised in nonactual worlds, where *nonactual worlds* are Lewisian worlds (Lewis, 1986)— *even if* Lewisian worlds exist; for the worlds *as invoked by* MM are themselves to be construed as nonexistent, independently of Lewisian worlds' existential status. To reach this position, I examine the relationship between Meinongianism and Lewisian Nonactualism and point out a theoretical possibility that Linsky and Zalta (1991), in asking whether Lewis is committed to Meinongianism, sidestep. For while Meinongianism and Nonactualism are undoubtedly structurally *very* similar, even if it turned out that *Meinongianism* and *Nonactualism* are really just two labels for a single metaphysical theory, one would still need to settle whether it is, ultimately, a Meinongian or Nonactualist one—given that Meinongianism and Nonactualism are designed to fulfil different metaphysical job descriptions. But then it's neither incoherent nor patently unreasonable to regard Meinongianism and Nonactualism not as rivalling, but as complementary theories, and thus to adhere to *both* theories at once.

Ultimately, I argue that postulating nonactual worlds and their inhabitants and postulating intentional objects and whatever is needed to individuate them are independent metaphysical projects. While the former is pursued to shed light on the (part of) reality to which *we* belong (call this *non-intentional* [part of] reality $\Re$), the latter is pursued to shed light on the (part of) reality that serves as the target range of our non-$\Re$-related thoughts and imaginative endeavours (call this *intentional* [part of] reality I); and arguably, $\Re$ and I are disjoint. In effect, then, Meinongianism is meant to keep I and $\Re$ neatly apart, *whatever $\Re$ turns out to be*, e.g., even if it contains existent nonactual worlds. But then it isn't surprising that the worlds held to ground aspects of $\Re$ aren't to be held to ground aspects of I as well; instead, if aspects of I are also held to be grounded by objects that play a world-like role, then these objects had better be construed as belonging exclusively to I.

Bibliography

Barz,Wolfgang. 2016. Two-Dimensional Modal Meinongianism. Ratio 29:249–267.

Lewis, David. 1986. *On the Plurality of Worlds*. Blackwell.

Linsky, Bernard, and Edward N. Zalta. 1991. Is Lewis a Meinongian? Australasian Journal of Philosophy 69:438–53.

Meinong, Alexius. 1960. *The Theory of Objects. In Realism and the Background of Phenomenology*, ed. Isaac Levi, D. B. Terrell, and Roderick M. Chisholm. Free Press.

Priest, Graham. 2016. *Towards Non-Being* (Second Edition). Oxford University Press.

Sainsbury, Richard Mark. 2010. *Fiction and Fictionalism*. Routledge.

Andrea Raimondi
Does Intentionality Suffice for Mentality?

Northwestern Italian Philosophy Consortium (University of Turin, University of Pavia, University of Genoa, University of Western Piedmont)

According to a familiar suggestion, intentionality is the mark of the mental. This suggestion is supported in terms of the *intentionalist thesis* (IT), according to which intentionality is a necessary and sufficient condition for mentality. IT has been rejected by many philosophers, who argue that there are non-intentional mental states, such as emotions and moods (Bordini 2017, Dretske 1995, McGinn 1996, Searle 1983) and purely qualitative states, like bodily sensations (Antony 1997, Voltolini 2014). Crane (1998, 2001), Mendelovici (2013) and Tye (1995, 2000) have sought to rebut these objections.

Another option - less discussed - for the opponent to IT is rejecting the thesis that intentionality suffices for mentality. Nes (2008) tries to reject this thesis, moving from Crane's (1998; 2001) necessary and conjunctively sufficient conditions for intentionality, i.e. (a) directedness upon an object and (b) aspectual shape. The idea of Nes's argument is that on a simpleminded but natural reading of (a) and (b), it comes out that certain non-mental states are intentional.

Firstly, I shall reconstruct Nes's argument in a new, clearer form. Nes redescribes (a) and (b) in terms of two features of the intensionality of reports: (a*) a state exhibits the relevant feature of directedness only if the corresponding report suffers from failure of existential generalization; (b*) a state exhibits the relevant feature of aspectual shape only if the corresponding report suffers from failure of substitutivity. Then, Nes shows that there are some non-mental states whose reports suffer from failure of existential generalization and from failure of substitutivity, which makes them intentional states.

I shall then consider and reject two different responses: (i) a *weak response*, according to which one should spell out intentionality in explicitly mental terms; (ii) a *strong response*, according to which one should ratchet up the necessary and conjunctively sufficient conditions for intentionality. After presenting a variety of strengthenings, I shall show that either fail to exclude all non-mental states or are so strong that they ground new challenges to the thesis that intentionality is necessary for mentality.

I shall consider a third response, Crane's (2008) reply, trying to reconstruct it rigorously. Crane argues that the relevant sense of intentionality is that of intentionality-as-*representation* and, hence, (a) and (b) should be reformulated as follows: (a') representativeness (the capacity of an intentional state to represent things) and (b') perspectuality (the fact that an intentional state represents things in a certain way, from a certain perspective). I shall show how Crane's proposal seems to provide a new form of intentionalism, which seems to respond to Nes's objection (Nes's examples of non-mental states cannot be said to represent anything at all). I shall argue that Crane's representationalist proposal is not adequate: although it is true that it is not correct to attribute to non-mental states any representative feature, this is not enough to provide a satisfying reply to Nes, for the new account of intentionality-as-representation plays no relevant explanatory role, if compared with the "old" account of intentionality.

Bibliography

Antony, L. (1997), "What It's Like to Smell a Gardenia?". In *Times Literary Supplement* 4897: 25.

Bordini, D. (2017), "Not in the Mood for Intentionalism". In *Midwest Studies in Philosophy* 41 (1): 60-81.

Crane, T. (1998), "Intentionality as the Mark of the Mental". In *Aspects of Psychologism*, Crane, T. (2014). Harvard University Press.

— (2001), *Elements of Mind*. Oxford University Press.

— (2008), "Reply to Nes". In *Analysis* 68 (3): 215—18.

Dretske, F. (1995), *Naturalizing the Mind*. The MIT Press.

Mendelovici, A. (2013), "Intentionalism About Moods". In *Thought* 2: 126-36.

McGinn, C. (1996), *The Character of Mind. An Introduction to the Philosophy of Mind*. Oxford University Press.

Nes, A. (2008), "Are Only Mental Phenomena Intentional?". In *Analysis* 68 (3): 205-15.

Searle, J. (1983), *Intentionality. An Essay in the Philosophy of Mind*. Cambridge University Press.

Tye, M. (1995), *Ten Problems of Consciousness*. The MIT Press.

— (2000), *Consciousness, Color and Content*. The MIT Press.

— (2006), "The Thesis of Nonconceptual Content". In *European Review of Philosophy* 6: 7-31.

Voltolini, A. (2014), "The Mark of the Mental". In *Phenomenology and Mind*: 169-86.

Henry Ian Schiller
Belief in Fiction

University of Texas at Austin

Brief Summary If you believe what a sentence says, then how many of that sentence's metaphysical commitments do you thereby endorse? A plausible claim is that in order to believe what a sentence says you must endorse all of its metaphysical commitments. Yet, thinking about something does not commit you to the existence of that thing (unicorns, Sherlock Holmes) even when talk about such thoughts does sometimes seem to carry ontological commitment. This paper challenges received wisdom and argues that beliefs do not inherit the ontological commitments of the sentences we use to express them. Specifically, I argue against two views of how to treat beliefs in what is said by fictional sentences and propose an alternative explanation.

Abstract We might appeal to a very plausible principle governing the relationship between the ontological commitments made by a sentence S and those made by people who believe S is true:

> *COMMITMENT PRINCIPLE*
> If a sentence S is committed to the existence of O, and a person a believes that S is true, then a is thereby committed to the existence of O.

In spite of the intuitive plausibility of the Commitment Principle, it runs into trouble when it comes to sentences involving well-known fictional characters and kinds. Consider the following set of sentences about fictional entities.

(1) a. Unicorns have horns.
    b. Apollo has a twin sister.
    c. Sherlock Holmes is a detective.

We have three claims which stand in tension:

i. The sentences in (1) are ontologically committing.

ii. It is rational (acceptable, etc.) to believe that these sentences are true.

iii. The Commitment Principle is true.

It naturally follows that we can resolve this tension by denying one of these claims. 1 Some motivation for denying (i) and (ii) is already present in work in the philosophy of fiction, philosophy of language, and metaphysics. In general, arguments for denying (i) and (ii) look as follows: there is typically taken to be something akin to an 'in the fiction of X' operator present in either the semantics for the sentences (i) or in the associated beliefs (ii). In this paper, I give an account of how we might deny (iii).

The view I endorse is that believing what a sentence says involves being disposed to make predications of representational content in a particular way. When I imagine a parrot, I represent that parrot. There might not be any particular parrot I am imagining – that is, I am not thinking about a real object. Nevertheless, I can imagine that this parrot is red; that is, I can predicate redness of it.

Belief, then, comes when I am committed to representing this parrot in this particular way (as red). But this obviously does not commit me to the existence of a parrot; just to a certain way of representing things. To believe what is said by a sentence 'a is F' is nothing over and above being committed to representing a as F. This does not involve a commitment to the existence of a denotation of 'a'; belief that something is real is therefore a further belief (perhaps it involves beliefs about a having certain location properties, or temporal properties).

The notion that we form representations of fictional entities (i.e., represent the 'referents' of names of fictional characters) is relatively uncontroversial. Further, there does not seem to be any reason for holding that these representations must be sui generis distinct from representations of nonfictional entities. Bach (1987) and Jeshion (2002) both note that we might conceive of mental reference in terms of something like a 'file' associated with a name into which predicates are sorted. A name like 'Aristotle' will be associated with a file under which certain predicates (like 'student of Plato', 'teacher of Alexander', 'lived in Greece', but also 'occupied a spatial location') are sorted. The same is true of a name like 'Sherlock Holmes'. It is thus fairly easy to establish how representations of nonexistent entities are formed. Even philosophers who are broadly opposed to a mental files conception of beliefs can agree with the notion of fictional entities being associated

with mental representations. The mental files conception simply provides an attractive framework for typically adopt something like this approach.

References

Bach, Kent. 1987. *Thought and reference*. Oxford University Press.
Jeshion, Robin. 2002. Acquaintanceless de re belief. In Joseph Keim Campbell, Michael O'Rourke & David Shier (eds.), *Meaning and truth*, New York: Seven Bridges Press.

Merel Semeijn
Metafictional Anaphora. An Abstract Object Account

University of Groningen

1. The workspace account

In Stalnaker's [7] widely adopted pragmatic framework, assertions are modelled as proposals to update the 'common ground' (i.e. the set of mutually presupposed propositions between speaker and addressee). The 'workspace account' [6] is a Stalnakerian approach to modelling, in Recanati's [5] terminology, 'fictional' and 'parafictional' statements. Fictional statements are statements taken directly from some fictional work (e.g. (1) from *Lord of the Rings* (*LOTR*)):

(1)   Frodo had a very trying time that afternoon.

Parafictional statements are statements about the content of some fictional work and contain a (covert or overt) fiction-operator (e.g. (2) as found in a discussion on *Lord of the Rings*):

(2)   In *Lord of the Rings*, Frodo is a hobbit.

Fictional statements are, like regular assertions, modelled as proposals to update a temporal common ground: the workspace. What differentiates non-fiction from fiction is whether, at the end of the possibly multi-sentence discourse, 'assertive' or 'fictive closure' is performed: Whether the content of the updated workspace is added to the 'official' common ground as belief (for non-fiction) or as parafictional belief (for fiction) under the relevant fiction-operator. I present a simplified representation of fictive closure of (1) in the box notation of DRT [2] (in which NP's in a discourse are mapped to 'discourse referents' placed under several conditions):
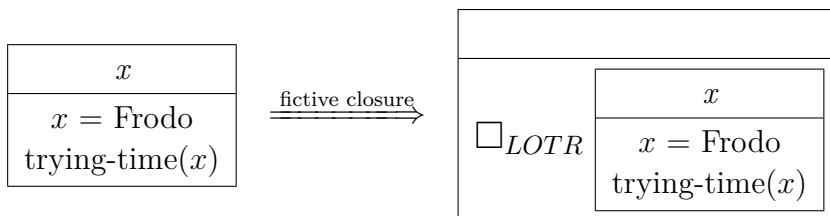
As soon as I stop entertaining the propositions of *Lord of the Rings* (e.g. as I stop reading), I stop updating the workspace and instead update the common ground with parafictional beliefs concerning *Lord of the Rings*. Parafictional statements such as (2) are based and update on such parafictional beliefs.

2. The challenge of metafictional anaphora

A third type of statement about fictions, the so-called 'metafictional statement', poses a problem for the workspace account. Consider the following felicitous discourse:

(3) In *Lord of the Rings*, Frodo$_i$ goes through an immense mental struggle to save his$_i$ friends. Ah yes, he$_i$ is an intriguing fictional character!

The last sentence in (3) is a metafictional statement because it is about Frodo *as a fictional entity*. As Lewis [3] argued, metafictional statements cannot be understood as covertly embedded under fiction-operators. Instead, the metafictional statement in (3) is a proposal to directly update the common ground (in bold):
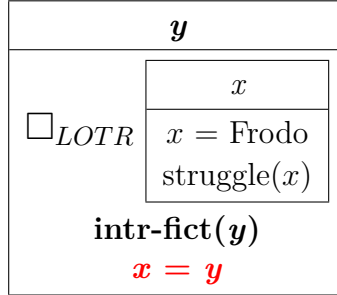
The metafictional statement in (3) contains a pronoun 'he' that is anaphoric on the name 'Frodo' introduced in the preceding *parafictional* statement. We represent this by equating their discourse referents (i.e. adding $x = y$ to the common ground). However, following standard DRT-rules, x is not accessible outside of the *LOTR* fiction-operator. Hence, it is unclear how we can interpret metafictional statements such as the one found in (3).

3. An abstract object account

There are different strategies available to meet the described challenge. First, one can adjust the accessibility relations (e.g. Maier [4]). Second, one can claim $x \neq y$ by arguing that metafictional statements are, unlike

parafictional statements, about abstract objects (e.g. Currie [1]; Recanati [5] to some extent).

The third available strategy can be found in Zalta's application of his 'object theory' [8, 9] to fiction. A story (e.g. *Lord of the Rings*) is an abstract object that encodes only vacuous properties. A fictional character is an abstract object that is *native* to a story (e.g. Frodo, but not Napoleon). Parafictional statements (with overt fiction-operator) are about specific encoding and exemplifying relations between stories and characters. For instance (2) expresses that *Lord of the Rings* encodes the property of being such that Frodo exemplifies being a hobbit:

$$\Sigma_{LOTR}\text{hobbit}(f)$$

From this we can derive that the abstract object Frodo encodes being a hobbit $((f)\text{hobbit})$ apart from exemplifying properties such as being an intriguing fictional character $(\text{intr-fict}(f))$.

Incorporating these ideas into the workspace account suggests a modification of the fictive closure operation. Below is a (simplified) representation of fictive closure* of (1):
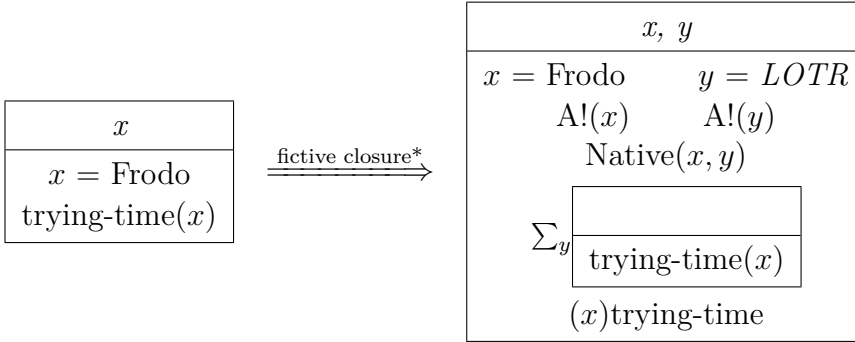
As soon as we stop entertaining the propositions of *Lord of the Rings* (and realize that *Lord of the Rings* is about fictional characters), we update the common ground with discourse referents for the relevant abstract objects (e.g. Frodo) and with parafictional beliefs (e.g. concerning the relations between the abstract objects *Lord of the Rings* and Frodo). This means that the discourse referent for 'Frodo' is accessible outside of the fiction-operator. Hence, returning to (3), when we update a similar common ground (based on the parafictional statement in (3)) with the metafictional statement in (3), we can equate the discourse referents for 'Frodo' and 'he' and hence interpret the statement:

$$\begin{array}{|c|}
\hline
x,\ y,\ \boldsymbol{z} \\
\hline
x = \text{Frodo} \qquad y = LOTR \\
\text{A!}(x) \qquad \text{A!}(y) \\
\text{Native}(x,y) \\
\\
\Sigma_y \begin{array}{|c|} \hline \\ \hline \text{struggle}(x) \\ \hline \end{array} \\
(x)\text{struggle} \\
\textbf{intr-fict}(\boldsymbol{z}) \\
\boldsymbol{x = z} \\
\hline
\end{array}$$

References

[1] Gregory Currie. *The Nature of Fiction*. Cambridge University Press, 1990.

[2] Hans Kamp. A theory of truth and semantic representation. In Jeroen A. G. Groenendijk, Theo M. V. Janssen, and Martin B. J. Stokhof, editors, *Formal Methods in the Study of Language*, Part 1, pages 277-322. Blackwell Publishers Ltd, Oxford, UK, 1981.

[3] David Lewis. Truth in Fiction. *American Philosophical Quarterly*, 15(1):37-46, 1978.

[4] Emar Maier. Fictional Names in Psychologistic Semantics. *Theoretical Linguistics*, 43(1-2):1-45, 2017.

[5] Francois Recanati. Fictional, metafictional, parafictional. (to appear).

[6] Merel Semeijn. A Stalnakerian analysis of metafictive statements. *The proceedings of the 21st Amsterdam Colloquium*, 2017.

[7] Robert C. Stalnaker. Pragmatics. *Synthese*, 22(1-2):272-289, 1970.

[8] Edward N. Zalta. *Abstract Objects: An Introduction to Axiomatic Metaphysics*. Springer, 1983.

[9] Edward N. Zalta. *Intensional Logic and the Metaphysics of Intentionality*. MIT Press, 1988.

Alberto Voltolini
Intentionality as Constitution

Università degli Studi di Torino

In this paper, I want to hold the following theses. i) Intentionality, conceived primarily as intentionality of *reference*, i.e., *aboutness*, is an *internal* thought-object relation. ii) It is not only a modal relation of existential dependence of a thought on its object. It is also a *constitutive* relation, a relation of individuative dependence of a thought on its object, whose predication to something is true in virtue of its nature. So, it is an essential property (Fine 1994) rather than a merely necessary property (Simchen 2012) of its bearer, the thought. iii) It is *the* constitutive relation, i.e., the property of *being constituted by* something.

As to i), as a venerable tradition tracing back to Husserl (1970) and Wittgenstein (1964) has maintained, the best explanation of the nature of intentionality is to claim that it is an *internal* relation: the best forms of externalism presuppose that claim, internalism does not deny it, the individuation conditions of an intentional state suggest it. Yet, as to ii), this claim cannot only mean that a thought depends on the object it is about in an existential sense (Mulligan-Smith 1986): necessarily, if that thought exists, its object also exists. For a) being an internal relation of existential dependence is a *necessary*, but not a *sufficient*, condition of intentionality, and b) being an internal relation of existential dependence at most vacuously accounts for the *individuative* factor of intentionality, i.e., for the individuative role the object a thought is about plays with respect to that thought. Rather, intentionality is a relation of individuative dependence of a thought on its object (content), so that it is not only a necessary, but also an *essential* property for that thought, i.e., a property that affects its own nature. This can be argued for in the following steps: 1) linguistic expressions and other public representations (pictures, etc.) have derived intentionality, i.e., intentionality as borrowed by something else; 2) to have *derived* intentionality *means* to be 'transformed into' (actually, to be paired with) a an entity-*cum*-meaning, i.e., a hybrid entity made by a physical element plus its semantic interpretation; 3) that entity has intentionality essentially; 4) a thought lends its intentionality to what has it only derivatively; hence, 5) a thought has intentionality essentially as well. Finally as to iii), intentionality must be more than a mere relation of

individuation, the relation for the thought of *being individuated* in terms of the object it is about. Since it must be the very same relation holding between an entity-*cum*-meaning and its constituents and this relation is a relation of constituency, intentionality itself is the relation for the thought of being *constituted by* the object it is about; the latter is an *essential part* of the former.

This accounts for the necessary and sufficient conditions for intentionality of reference, that is, the fact that a thought has aboutness iff it is had *directedness* (being such that its object, the *intentional* object of that thought, may not exist) and *aspectuality* (being such that its object appears to be an aspect of another entity) (Crane 2013). On the one hand, a thought is constituted by its object both when it actually exists and when it exists merely possibly. What indeed counts for that constitution is that such an object figures in the overall ontological domain: this holds for all *possible* objects, whether actual or merely possible (if this were not the case, the thought would have no object, it would rather involve an intentional *content*). On the other hand, two thoughts differing in their aspectuality are constituted by two different *intentionalia* that are two mere *possibilia*. Such mere *possibilia* appear however to be aspects of a further yet actually existent object, as the solution of the Frege puzzle for thoughts shows. The former objects are indeed *none other than* the latter object, insofar as being none other than another object is a contingent relation of *ontological contraction* that both differs from and is weaker than identity. This solution of the Frege puzzle reminds some Meinongian solutions to it, according to which the sameness relation weaker than identity that is involved there is a relation either of consubstantiation (Castañeda 1989) or of coexemplification (Zalta 1988) obtaining between different *intentionalia* that are Meinongian entities. Yet unlike such solutions, it does not presuppose that *intentionalia* are entities of a peculiar metaphysical kind. For the different *intentionalia* involved in a Frege puzzle are just mere *possibilia* that contingently contract themselves into another yet actually existent object.

Bibliography

Castañeda, H-N. (1989). Thinking, Language, and Experience. Minneapolis: University of Minnesota Press.
Crane, T. (2013). The Objects of Thought. Oxford: OUP.
Fine, K. (1995). "Ontological Dependence". Proceedings of the Aristotelian Society 95: 269-289.

Husserl, E. (1970). Logical Investigations. London: Routledge.

Mulligan, K., Smith, B. (1986). "A Relational Theory of the Act". Topoi 5: 115-130.

Simchen, O. (2012). Necessary Intentionality. Oxford: OUP.

Wittgenstein, L. (1964). Philosophical Remarks. Oxford: Blackwell.

Zalta, E.N. (1988). Intensional Logic and the Metaphysics of Intentionality. Cambridge MA: The MIT Press.