

# Philosophy of Language and Mind Conference

KEYNOTE SPEAKERS:

ROBYN CARSTON (UCL)

TERESA MARQUES (University of Barcelona)

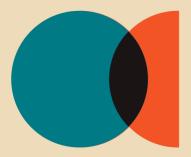
ANDREAS STOKKE (Uppsala University)

AGUSTIN VICENTE

(University of the Basque Country)

**ZSOFIA ZVOLENSZKY** 

(Eotvos Lorand University of Science)



It is our pleasure to welcome You to the Sixth Philosophy of Language and Mind Network Conference (PLM6) that will take place at the University of Warsaw, 15-17 September 2022.

PLM is a European network of centers devoted to the Philosophy of Language and Mind. PLM was founded in 2010 and organizes international conferences, workshops and master classes taught by leading experts in the field.

PLM is a network of the following institutions (with representatives in brackets):

- Arché, St. Andrews (Derek Ball)
- Department of Philosophy, CEU, Budapest (Ferenc Huoranszki)
- ▶ CLLAM, Department of Philosophy, Stockholm University (Kathrin Glüer-Pagin)
  - Department of Analytic Philosophy, Czech Academy of Sciences (Tomas Marvan)
- ▶ Department of Philosophy (KGW), University of Salzburg (Christopher Gauker)
  - ▶ Department of Philosophy, University of Vienna (Max Kölbel)
  - ▶ ILCLI, University of the Basque Country, San Sebastian (María de Ponte)
    - ▶ ILLC, University of Amsterdam (Paul Dekker)
      - ▶ Institut Jean Nicod, Paris (Denis Bühler)
    - ▶ Institut für Philosophie II, Ruhr University Bochum (Markus Werning)
      - ▶ Institute of Philosophy, London (Corine Besson)
  - ▶ Faculty of Philosophy, University of Warsaw (Joanna Odrowąż-Sypniewska)
    - ▶ LanCog, Universidade de Lisboa (Pedros Santos)
    - ▶ LOGOS, University of Barcelona (Manuel García-Carpintero)



The conference is hosted by the Faculty of Philosophy at the University of Warsaw. The local organiser is Joanna Odrowąż-Sypniewska.

Organizing Committee: Natalia Karczewska, Antonina Jamrozik, Zuzanna Jusińska, Kamil Lemanek, Dan Zeman.

Conference is financed by the Faculty of Philosophy UW, the University of Warsaw (IDUB), Polish Academy of Sciences (UiPDN) and the University of Warsaw Foundation.

The Faculty of Philosophy of the University of Warsaw Krakowskie Przedmiescie 3 st. Warszawa

Faculty website

Conference website

PLM Network website















# Keynote speakers:

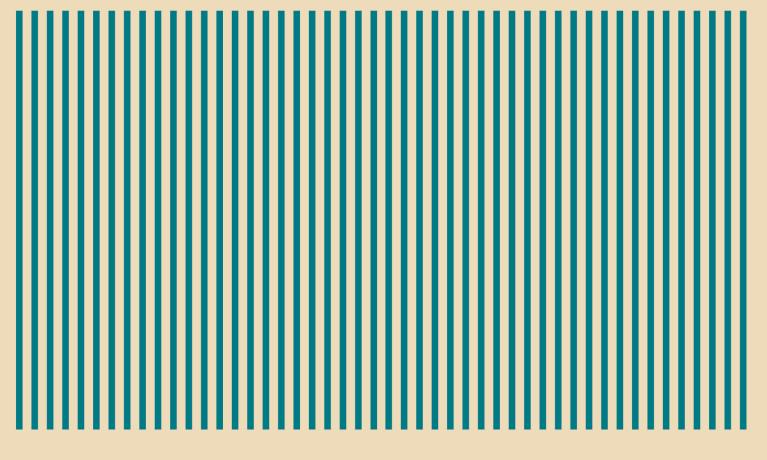
WORDS: PHRASAL STRUCTURES AND PRAGMATIC MEANINGS - Robyn Carston	ς
'SING, DON'T SHOUT' - THE PLACE FOR THE EXPRESSION OF ANGER IN DEMOCRACY - Teresa Margues	10
SPATIAL INDEXICALS - Andreas Stokke	12
LITERALISM AND RIGIDITY IN AUTISM - Agustin Vicente	13
FOCUSING ON FRAMES - Zsofia Zvolenszky	14

# Contributed papers:

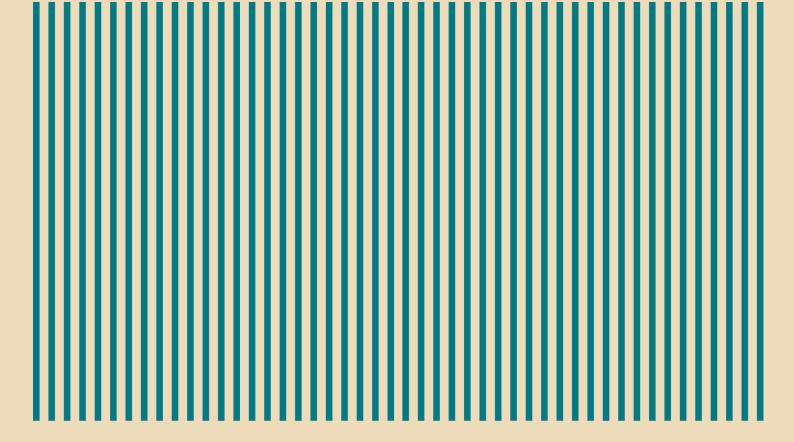
A PUSHMI-PULLYU PSYCHOLOGY - Nimra Asif, Alison Springle	16
A PRAGMATIC APPROACH TO DUAL CHARACTER CLASSIFICATIONS - Lucien Baumgartner	17
RUSSELLIAN PHYSICALISTS GET OUR PHENOMENAL CONCEPTS WRONG - Marcelino Botin	19
(NOT NECESSARILY CREDIBLE) DENIABILITY - Bianca Cepollaro, Dan Lopez de Sa, Andrés Soria Ruiz	22
A NOTE ON THE MEANING OF 'I' - Tadeusz Ciecierski, Jakub Rudnicki	25
A PUZZLE ABOUT MENTAL LEXICONS AND SEMANTIC RELATEDNESS - Alice Damirjian	27
SENTENCE AND SPEAKER ABOUTNESS, AND SUPER THICK PROPOSITIONS - Tamara Dobler	30
AUTHORITY, CONVENTION AND GROUP ASSERTION - Daniel Garibay-Garcia	31
THE SIGNIFICANCE OF UNINFORMATIVE CONTENT. THE CASE OF UNINFORMATIVE LIES - Grzegorz Gaszczyk	33
PRE-REFLEXIVE SELF-CONSCIOUSNESS AND ORIGINS OF AGENCY - Michael Hegarty	37
OBJECTIONABLE THICK TERMS AND INFERENTIALISM - Leopold Hess	39
LINGUISTIC MISTAKES AND THE PUBLIC LANGUAGE - Aleksi Honkasalo	40
Conceptual Considerations on Modeling referential ambiguity in Natural Language - A. Jampozik	42
LOCAL CONTEXTS FOR SITUATION SEMANTICS - Lukas Lewerentz	44
MENTAL IMAGERY AND CO-PREDICATION - Michelle Liu	47
COMMUNICATIVE NORMS ANND RATIONAL EXPECTATIONS - Neri Morsili	48

MENTAL IMAGERY AND THE REPRESENTATION OF MATERIAL LOOKS - Stephen Müller	49
LINGUISTIC VAGUE EXISTENCE - Elisa Paganini	52
Is transfer of epistemic warrant the mark of inferential transitions - IIIa Patronnikov	53
INTENTIONALISM OUT OF CONTROL - Carlotta Pavese, Alexandru Radulescu	54
ACTIVE INFERENCE BETWEEN REALISM AND INSTRUMENTALISM - Niccolò Aimone Pisano	55
HOW TO BE FICTIONALIST ABOUT INSTRUMENTAL GROUP MINDS - Krzysztof Posłajko	57
LANGUAGE AS EXPRESSSION OF THOUGHT - Indrek Reiland	59
PROPOSITIONS AND NON-DOXASTIC ATTITUDE REPORTS - W. Rostworowski, K. Kuś, B. Maćkiewicz	69
EXPERIMENTAL ACQUAINTANCE WITH ABSTRACT OBJECTS - Işık Sarıhan	65
COMMON GROUND IN DEVELOPMENT - Antonio Scarafone	66
BREAKING THE FOURTH WALL AND METAFICTIONAL REFERENCE - Merel Semeijn	69
EXISTENTIAL GENERIC AND INFORMATION STRUCTURE - Zhengjie Situ	79
THE AFFECTIVE BASIS OF INFERENTIAL REASONING - Sanja Sreckovic	73
MY PHILOSOPHICAL POSITION IMPLIES 'P BUT I DON'T BELIEVE THAT P' - Maciej Tarnowski	75
HOW GÖDEL BOTH DID AND DID NOT STEAL FROM SCHMIDT - Jeske Toorman	78
THE BURDEN OF COMMUNICATION: IS THERE A REAL 'DEVELOPMENTAL DILEMMA'? - E. Vaccargiu, D. Mazzarella	80
GENERIC VIEW OF GENDERED SLURS AND THE SUBSET ARGUMENT - Pasi Valtonen	83
EVALUATIVE DEFLATION, SOCIAL EXPECTATIONS & THE ZONE OF MORAL INDIFFERENCE - Willemsen, Baumgartner, Cepollaro, Reuter	85
WHAT WE TALK ABOUT WHEN WE TALK ABOUT 'BELIEVE' - Simon Wimmer	87
UNDETERMINISTIC COUNTERFACTUALS - Tom Wysocki	90





# KEYNOTE SPEAKERS





Complex words like 'reactionary', 'transmission', and 'liquidate' have non-compositional (atomic) meanings, not predictable from the meaning of their parts and their structure. This creates an issue for single syntactic engine theories, which subsume word structure into syntactic structure and take roots as their basic elements, because the structures that syntax generates (phrases) typically have compositional meaning. In this paper, I discuss work in syntax that aims to provide a principled account of those structural domains that allow atomic content and those that do not, and put it together with work in pragmatics that aims to explain how these atomic senses arise (as ad hoc concepts) in on-line communication. The result is a pragmatic lexicon, which stores words with their families of related senses (polysemy), and other information relevant to their production and comprehension in communication. Words are the basic units of communication, psychologically salient to language users, although they have no existence as such in the syntax. The pragmatic lexicon lies outside the narrow language faculty but interfaces with it at specific points (the domains that allow assignment of atomic content).



Is the expression of anger a legitimate form of counter speech? When is the expression of anger permissible in deliberative democracy? Amia Srinivasan (2018) and Maxime Lepoutre (2018) argue that the expression of anger has such a place. Their views contrast with, and respond to, Martha Nussbaum's (2016). Nussbaum argues that anger is counterproductive and has harmful effects, even though it can sometimes be instrumentally useful. Lepoutre's and Srinivasan's views rely on the premise that anger has fittingness conditions. Srinivasan argues that anger is apt when it is the fitting response to injustice, but denies any essential connection between anger and intentions to act retributively. Lepoutre addresses the counterproductivity objection, arguing that anger can be epistemically productive, directing listeners to unjust features of a situation that had not been previously registered.

Here, I argue that there is an essential connection between anger and retributive action. I contest the interpretation of anger's fitting conditions that Srinivasan and Lepoutre rely on. Emotions are evolved cognitive mechanisms that trigger thought and action. Anger in particular triggers actions to redress what the subject registers as an offense. It is not constitutive of anger that it is apt only when an offense is unjust or immoral. Also, anger is often epistemically defective: it displays authority (Park et al. 2013), it plays a role in intergroup violence (Claassen 2014), with narcissistic entitlement, it's correlated with aggressiveness (Reidy et al 2007), with male aggrieved entitlement, it plays a role in misogyny (Kimmel 2014, Manne 2017), with contempt and disgust, it correlates with violent action (Matsumoto et al 2015), and acting out of anger leads to conflict escalation (Schuman and Ross 2010). These are reasons to reject the claim that anger is a reliable indicator of injustice. They are also reasons to doubt its reliability as a motor of social justice movements. I suggest that it is necessary to distinguish the appropriateness of the emotion of anger from that of its expression. In particular, I suggest that the appropriateness of the social expression of anger is strongly dependent on the context and on the form of its expression. On a positive note, I suggest that there are some constrained appropriate conditions for anger expression (see e.g. the Truth and Reconciliation Commission in South Africa (Chakravarti 2014)). But permissible anger expression in democracy, I will suggest, requires contexts of common ground and mutual trust, which are hard to ensure.

### Bibliography

Chakravarti, S. (2014): Sing the Rage. University of Chicago Press.

**Claassen**, C. (2014): "Group Entitlement, Anger, and Participation in Intergroup Violence", British Journal of Political Science 46: pp. 127-148

Kimmel, M. (2015): Angry White Men: American Masculinity at the End of an Era. Nation Books.

**Lepoutre**, M. (2018): "Rage Inside the Machine". Politics, Philosophy, and Economics 17(4): pp. 398-426.

**Matsumoto**, D. et al. (2015): "The Role of Intergroup Emotions in Political Violence". Current Directions in Psychological Science, 24(5), 369-373.

Nussbaum, M. (2016): Anger and Forgiveness, OUP.

**Park**, J. et al, (2013): "Social status and anger expression: the cultural moderation hypothesis". Emotion. 213(6): pp. 1122-1131.

**Reidy**, D.E. et al. (2007): "Effects of narcissistic entitlement and exploitativeness on human physical aggression" Personality and Individual Differences 44 (2008) 865–875.

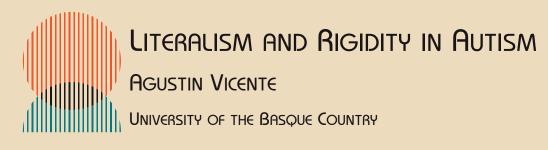
**Schumann**, K. and Ross, M. (2010): "The Benefits, Costs, and Paradox of Revenge". Social and Personality Psychology Compass, 4: 1193-1205.

Srinivasan, A. (2018): "The Aptness of Anger". Journal of Political Philosophy, 26(2): pp. 123–144.

Strawson, P. (1974): Freedom and Resentment and other Essays. Routledge.



This papers offers a theory of spatial indexicals like "here" and "there" on which such expressions are variables associated with presuppositional constraints on their values. I show how this view handles both referential and bound uses of these indexicals, and I propose an account of what counts as the location of utterance on a given occasion.



One of the most characteristic traits in autism is "rigidity" or "inflexibility". Autistic people have a tendency to need routines, to be strict about rule following, to insist on sameness, to experience intolerance to uncertainty, and to have difficulties switching tasks or perspectives. At the same time, autistic people are known to be more literalist than neurotypical people, in that autistic people have a stronger tendency to interpret non-literal uses of language literally. My talk will have two parts. In the first one, I will try to clarify the "rigidity" notion, or family of notions. In the second part, I will argue that current explanations of literalism in autism fail, and that such literalism is probably related to some rigidity facet.



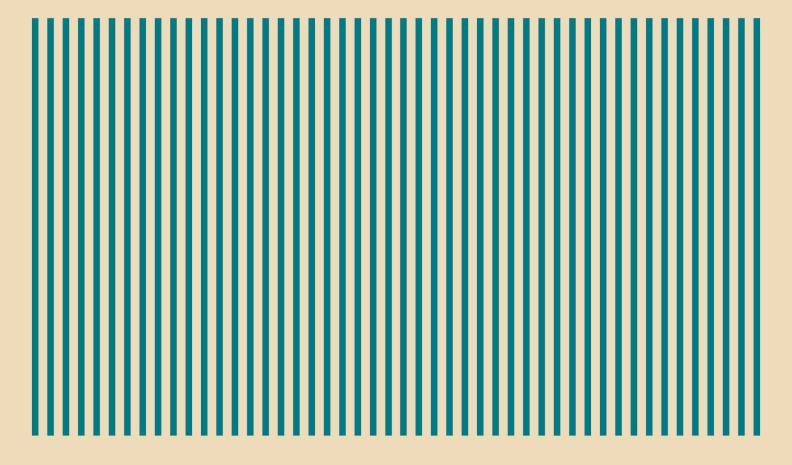
Full title: Focusing on frames: What has been presupposed about presuppositions in central debates within philosophy of language

Kellogg's Raisin Bran has for long been advertised as a "heart healthy" cereal choice --- even though a small bowl contains almost half of the recommended daily limit intake of sugar, given all the added sugar. The information content of this claim about the cereal being "heart healthy" has been challenged, even in court; yet the framing devices at play, including in the very use of the label "Kellogg's Raisin Bran" – accentuating the raisin and bran content – and their influence on associated discussion have received considerably less attention. This phenomenon of the discussion focus shifting to information content and away from the effects of framing devices is widespread, holding implications, cautionary tales for scholarly inquiry as well, as I aim to show in this talk.

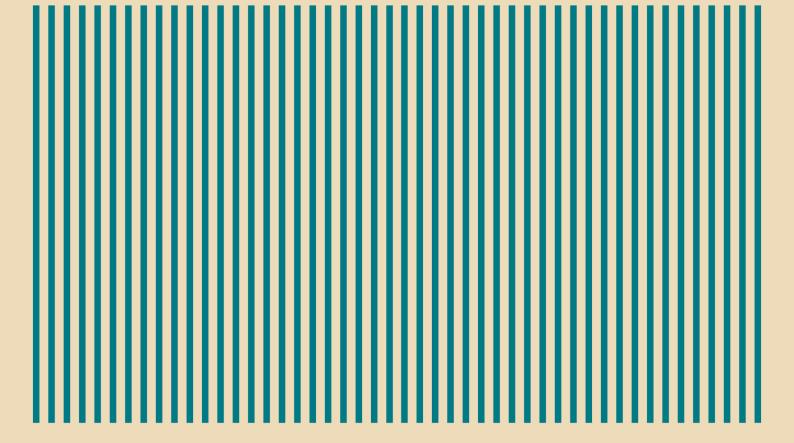
The choice of labels is one of various ways in which a discussion is affected by framing devices, which shape the way we see the world by invoking memories, feelings, emotions, associations, opinions, preferences, images that go well beyond information content. Crucially, a discussion is also affected by frames imposing a certain structure on approaching, understanding an issue, and that structure includes a taxonomy of contrasting alternatives.

Frames affect discussions of all sorts, including scholarly inquiry. My focus in this talk is on frames at play in a series of interconnected case studies in 20th-21st-century debates in philosophy of language about presuppositions in work by and on H. P. Grice and Saul Kripke. I will, first, illustrate how leaving these frames – with their concomitant taxonomical assumptions, for example -- undiscussed has hindered, slowed those debates. Second, I will point out how discussion and better understanding of frames can – by imposing checkpoints about frames at play, as well as checkpoints for needed taxonomical and other frame-related revisions – can make such debates more effectively conducted in future research (within philosophy of language as well as outside of it). Third, I will reflect on some of the ways in which philosophy of language as a subfield can shift, expand its traditional focus – which has been on truth conditions, propositional content, information expressed by linguistic utterances -- and contribute to the theoretical toolbox for understanding frames.

The interconnected case studies on debates about presuppositions within philosophy of language present counterexamples to a proposal of Elisabeth Camp's ("Perspectives and Frames in Pursuit of Ultimate Understanding" 2019) who holds that a chosen frame's role diminishes, disappears in later stages of inquiry; as details are understood, theories transcend initially introduced frames as being oversimplified. In fact, I argue, we see ill-fitting frames' and their taxonomical assumptions' lasting influence – unless we shift the traditional focus on information content and subject those very frames to systematic discussion and theorizing.



# CONTRIBUTED PAPERS





# A PUSHMI-PULLYU PSYCHOLOGY

# Nimra Asif, Alison Springle

University of Connecticut,
University of Tübingen / University of Oklahoma

"Pushmi-pullyu" representations, and similar such posits, are a) more primitive than propositional attitudes and b) have aspects in common with both beliefs and desires. As we explain in §1, representations of this sort are frequently posited to account for nonhuman animal mental states, original intentionality, at least some perceptual representations, in-between belief cases, the representations that underwrite skillful action, and emotions. However, as we'll argue in §2, it is difficult to make sense of pushmi-pullyu representations against the background of the standard propositional attitude model. One way to remedy this is to re-think the propositional attitude model: rather than trying to squeeze pushmi-pullyus out of propositional attitudes, perhaps propositional attitudes can be fitted into a model of mental representations that treats pushmi-pullyu representations as basic. In §3 we propose that understanding pushmi-pullyu representations as "embodied instructive representations" (Springle & Humphreys 2021; \*\*\*\*\*) may provide a way to do this.



Dual Character Concepts (DCCs) have been introduced as a distinct concept class with two independent dimensions for classification: a descriptive and a normative one. For exam- ple, Amy is a descriptive philosopher, for she has tenure at the Department of Philosophy at XYZ. Joe, on the other hand, is the janitor at XYZ, and thus is not a philosopher, descriptively speaking. But Joe also convinced the university to invest into eco-friendly detergents on the whole campus, based on an elaborate moral argument. So, while Joe might not be a philosopher descriptively speaking, he arguably is a philosopher in a normative sense. Inversely, if it becomes public that Amy mostly plagiarized her way into her position, we might say that she technically is a philosopher, but certainly is not a true philosopher (normatively speaking). Hence, the concept Philosopher (C) allows for a classificatory double-dissociation, according to which descriptive Cs can fail to be true Cs, and descriptive non-Cs can qualify as true Cs (Del Pinal, 2015; Del Pinal et al., 2017; Del Pinal and Reuter, 2017; Knobe et al., 2013; Leslie, 2015; Reuter, 2019). Rather than arguing at the concept level, this paper stresses the importance of looking at how terms are used in everyday classifications, e.g. in the statement "Joe the janitor is a true philosopher."

It is widely assumed that DCCs have a default normative sense, which is part of their generally accepted standard meaning (see Leslie, 2015, for a discussion). Most authors suggest tentative semantic explanations related to lexical meaning (e.g., Knobe et al., 2013; Leslie, 2015; Reuter, 2019). Pragmatic explanations (e.g., Baumgartner, ms; Hesni, 2021), on the other hand, are less represented. The aim of this paper is to show that dual character classification patterns can be explained by lexical-pragmatic processes like sense narrowing and broadening, and can involve the figurative use of words (e.g. metaphors). These findings suggest that dual character—as a classification scheme—is generally more adequately described as the result of lexical-pragmatic processes rather than constituting a genuine and distinct concept class. Hence, this proposal goes against the received notion of dual character as a semantic feature of certain terms, and breaks with the traditional position that dual character is a distinct structure of certain lexical concepts. Instead, I suggest that dual character is a feature of speaker meaning, which we arrive at via pragmatic inferencing.

The pragmatic understanding of dual character is probably best illustrated by ad hoc metaphorical classifications, such as:

- (1) These socks are true magicians! Anytime I do the wash, one of them magically disappears.
- (1) is arguably a dual character classification, for socks (i.e. descriptive non-magicians) are said to be true magicians. The socks' propensity to disappear might even put to shame certain professional magicians who always fail their disappearance tricks. At the same time, we would suspect that the speaker in (1) uses figurative language: First, magicians are usually human beings rather than socks. Second, it seems odd to put socks literally in the same category as a true magician, like Houdini. This is partly due to the fact that the property for the normative classification, i.e. to magically disappear, implies some kind of agency on

behalf of the magician, whereas it does not for inanimate objects like socks. Thus, applying said predicate to socks results in a linguistic mismatch, which triggers the pragmatic process of predicate transfer, i.e. the pragmatic sense modulation of the predicate in order to accommodate what it is applied to (Nunberg, 1995; Recanati, 2004, 2005). Third, it seems plausible that the speaker does not literally think of their socks as true magicians, but rather uses the classification hyperbolically (and metaphorically), to express just how unnerving washing socks actually is. In sum, the pragmatic meaning of the term magician in (1) is both narrower and broader than its lexical meaning: it only includes entities that are good at magically disappearing, among others certain non-magicians (such as the socks in (1)), while excluding professional magicians who are bad at disappearing.

This shows that dual character classifications do not require a lexical concept with a distinct structure, i.e. what is traditionally understood as DCC. Rather, these classifications can be explained by lexical-pragmatic processes and can include figurative elements.

### References

Baumgartner, L. (ms). Dual Character Concepts Gone Wild.

**Del Pina**l, G. (2015). Dual Content Semantics, privative adjectives, and dynamic compositionality. Semantics and Pragmatics, 8(0):1–53.

**Del Pinal, G., Madva, A., and Reuter, K**. (2017). Stereotypes, Conceptual Centrality and Gender Bias: An Empirical Investigation. Ratio, 30(4):384–410.

**Del Pinal, G. and Reuter, K.** (2017). Dual Character Concepts in Social Cognition: Commitments and the Normative Dimension of Conceptual Representation. Cognitive Science, 41:477–501.

**Hesni**, S. (2021). Normative generics: Against semantic polysemy. Thought, 10(3):218–225. Knobe, J., **Prasada, S., and Newman, G. E**. (2013). Dual character concepts and the normative dimension of conceptual representation. Cognition, 127(2):242–257.

**Leslie**, S.-J. (2015). "Hillary Clinton is the Only Man in the Obama Administration": Dual Character Concepts, Generics, and Gender. Analytic Philosophy, 56(2):111–141. Nunberg, G. (1995). Transfers of Meaning. Journal of Semantics, 12(2):109–132. Recanati, F. (2004). Literal Meaning. Cambridge University Press.

**Recanati**, F. (2005). Literalism and Contextualism: Some Varieties. In Preyer, G. and Peter, G., editors, Contextualism in Philosophy: Knowledge, Meaning and Truth, pages 171–196. Oxford University Press. **Reuter**, K. (2019). Dual character concepts. Philosophy Compass, 14(1):e12557.



It is hard to see how pain could be a brain state. Physical facts don't look like the kind of fact that could entail nor explain phenomenal facts and painfulness does not seem like a property of soggy grey matter. These hard problem intuitions suggest dualism, but dualists have a hard time integrating consciousness with the physical world. This conundrum has pushed philosophers to abandon their posts and embrace russellian monism (RM), a view that occupies a middle ground between dualism and physicalism and promises to avoid their shortcomings.

Russellians argue that we cannot see that pain is grounded on a brain state due to our impoverished understanding of physical. Our epistemic access to the world only reveals its relational/dispositional properties. But where we have relational/dispositional structure, they say, there must be some underlying intrinsic properties that ground it. RM claims that there is a single class of inscrutable intrinsic properties at the fundamental level which are the common base of everything else, of the chair in which I now sit and of my experience while writing this sentence. Russellians, however, disagree about the nature of such properties. Those who cannot give up the view that consciousness is fundamental embrace a panpsychist version of RM (Goff, 2017) and those who cannot accept that consciousness is fundamental embrace a panprotopsychist or physicalist version (Alter and Nagasawa, 2012; Brown 2017).

This talk will concern russellian physicalism (RP), the view that the intrinsic properties of matter are purely physical. And while it may seem that more physical properties cannot help with the hard problem, according to RP, postulating robustly physical intrinsic properties that constitute consciousness can explain our problem intuitions. Indeed, the physical facts that we are familiar with do not explain phenomenal facts because they hide consciousness relevant facts, they are silent on the intrinsic aspect of the physical. Similarly, we can't accept that phenomenal properties are physical because we underestimate the physical due to our ignorance of its full nature.

In this way, RP delivers what everybody wants, a metaphysics of mind that does not lead to a disjointed ontology and that is compatible with the existence of consciousness. To increase the lure of their physicalism even further, RP insists that standard physicalism is a form of illusionism. But this lumping of all standard forms of physicalism into the consciousness-deniers camp is mistaken. The orthodox physicalist is a realist through and through. Type-B physicalism, the most popular form of physicalism, shares the "dualist" problem intuitions and, like RP, it promises to accommodate them within a physicalist framework. Does this mean that physicalists have two ways of reconciling with realism about consciousness? Unfortunately no, I will argue that there is just one way of defending physicalist realism and it is not the russellian way.

My argument will be familiar to anyone who rejects the standard physicalism by reflecting on their phenomenal concepts (as many russellians do). It has been argued that type-B physicalism cannot accommodate the substantial knowledge afforded by our phenomenal concepts. Goff (2011) provides a

clear exposition of the problem. Our phenomenal concepts, he says, reveal what been in some phenomenal state essentially consists of, and this is incompatible with type-B physicalism. I will focus on this formulation of the revelation challenge and show that although it is a surmountable difficulty for type-Bs, the argument, applied to RP, is deeply problematic.

The problem for the russellian arises due to two commitments of the view: A) that revelation is true and B) that if we had a transparent account of the physical that includes all facts about the intrinsic nature of the physical there would be an a-priori entailment from physical to phenomenal facts. I argue that RP is committed to A because denying revelation is incompatible with realism about consciousness. And I argue that RP is committed to B because denying a-priori entailment would leave no explanatory work for the tenets of the theory.

The reason why A and B lead to trouble is the following: Our phenomenal concepts do not reveal that what it is to be in a phenomenal state is to instantiate any recognisably physical property. When we reflect on how pain feels nothing tells us that what it would take for someone to feel pain is to be in a state of Cfiber firings. The claim is not that phenomenal concepts reveal phenomenal properties as non-physical. But rather, that phenomenal concepts do not reveal phenomenal properties as physical. This is perfectly compatible with such properties being physical insofar as there is a physically acceptable explanation of how this is so. Type-B physicalism, appealing to the cognitive isolation of phenomenal concepts, claims that we can know about the nature of consciousness in two conceptually distinct ways. Thus, they can accommodate revelation by showing that our phenomenal concepts reveal essential features of physical properties but not under a physical characterisation or mode presentation. However, this strategy is incompatible with RP because there can be no a-priori entailment between two radically distinct modes of presentation. If our physical concept "RED+" (where "+" adds facts about intrinsic physical properties) is a-priori connected with our phenomenal concept "RED\*" then both concepts must capture their shared referent (assuming identity) by a recognisably similar characterisation. RED+ must characterise intrinsic physical properties recognisably as the properties revealed by RED\*. But if RED\* reveals intrinsic properties as they would be characterised by an account of the intrinsic nature of physical then the view is incompatible with physicalism because it would mean that the fundamental intrinsic properties of reality are akin to the properties we find in introspection. And the only properties of this sort are either phenomenal or protophenomenal.

A natural response for a defender of RP is to insist that our phenomenal concepts do reveal physical properties, but we fail to conceptualise them as such due to our limited (structural) conception of the physical. But I argue that this runs into a seemingly unavoidable dilemma:

If intrinsic properties are recognisably physical then introspection does not reveal any such property and RP would have to deny that revelation is true.

If intrinsic properties are like the properties we find in introspection, the russellian can accept revelation but then there can be no physicalist version of RM.

As mentioned above, denying revelation is incompatible with RPs phenomenal realism and I argue that horn B is incompatible with RPs commitment to a robustly physicalist version of RM. If successful, my argument shows that, unlike type-B physicalism, RP cannot defeat the charge from revelation without

weakening its physicalist or realist credentials. And, therefore, that russellians fail to deliver on the promise of reconciling realism and physicalism about consciousness.

The plan for the talk is the following: Section one briefly explains russellian and type-B physicalism. Section two presents the revelation challenge against type-B physicalism and shows that it is more problematic for RP. I will expand on the schematic version of the argument presented above and strengthen the aforementioned dilemma. Section 3 considers two ways of resisting horn B of the dilemma and argues that they fail.

# References:

Alter, T, and Nagasawa, Y. (2012). "What Is Russellian Monism?". Journal of Consciousness Studies, 19 (9-10), 67-95.

**Brown**, C. D. (2017). "A properly physical Russellian physicalism". Journal of Consciousness Studies, 24(11–12), 31–50.

**Goff**, Philip. (2011). "A posteriori physicalists get our phenomenal concepts wrong". Australasian Journal of Philosophy, 89(2), 191-209.

———— (2017). Consciousness and Fundamental Reality. Oxford: Oxford University Press.

**Papineau**, D. (2019). "Response to Chalmers' 'The Meta-Problem of Consciousness". Journal of Consciousness Studies, 26 (9-10), 173-181.



# (NOT NECESSARILY CREDIBLE) DENIABILITY BIANCA CEPOLLARO, DAN LOPEZ DE SA, ANDRÉS SORIA RUIZ

Università Vita-Salute San Raffaele Milano, ICREA, IFILNOVA

Why do we often choose to insinuate something, when we could just openly state it?

One key factor is aiming to preserve so-called plausible deniability—that is, that we can deny that we have insinuated, if challenged (Mazzarella et al 2018). "Can" here cannot merely mean that it is possible to deny something (pace Viebahn 2020). People deny all sorts of things, including things they have openly asserted. Our main claim in this paper is that "can" cannot mean that denials are credible either. In this work, we argue that the philosophical debate on deniability has inappropriately narrowed its focus on an epistemic understanding of what it is for deniability to succeed. We argue that so-called 'plausible deniability'—insofar as it is one of the main features that we aim to preserve when insinuating (and using other forms of indirect speech)—covers more than the epistemic incarnation as credible deniability.

There are contexts where reasons for insinuating are indeed epistemic: perhaps we don't know the extent to which our interlocutors are cooperative, and insinuation is, in this sense, part of strategic speech. Suppose that a driver stopped for speeding tries to indirectly offer a bribe to the police officer:

"I'm in a bit of a hurry. Is there any way we can settle this right now?" (Lee & Pinker 2010)

Their attempt can be accepted by corruptible officers, and disregarded by incorruptible ones. A big part of the literature focuses on this kind of cases, understanding the bribe example as pivoting around how credible denials would be, and thus offering an epistemic account of deniability.

But these, we claim, are at best partial accounts of the phenomenon. The reason is that sometimes we choose to insinuate even when we all know—indeed, it is mutually known—that the insinuation takes place. And so, we all know—indeed it is mutually known—that denials would be false. For instance, in a conversation among intimate partners, she may choose to merely insinuate (rather than explicitly say) that she has met her lover again with:

"Sorry I'm late for dinner, I had a very long day"

just for the sake of avoiding yet another conversation about the issue.

Or imagine that the Republican congressman denied that he ever insulted President Biden on the House floor after concluding his remarks with:

"Let's go Brandon"\*

As Camp observes, in this kind of case the recipients of the denial may just "roll their eyes but acquiesce" (Camp 2018, 51), showing that successful denial does not amount to credible denial. That is, deniability can be sought—and kept!—even in cases where denials would be (mutually known to be) just totally incredible. Although this has been observed in the literature before, we aim here to highlight its crucial significance for the debate. The epistemic theories that aim to account for plausible deniability but merely focus on credible deniability prevent a proper understanding of the broader linguistic phenomenon. But what is then deniability in general, and why is it worth preserving, if it doesn't necessarily entail that the speaker's denial would be credible?

We do not aim to provide here a full-fledged theoretical articulation of deniability in general, but contribute some tools which, we hope, provide a fruitful framework for further research onto the nature of the phenomenon. In so doing, we elaborate, amend, and also depart from Elisabeth Camp's (2018) suggestive remarks. As she says,

"Insinuation constitutes a kind of communicative bluff: an attempt to make a conversational move without paying the conversational cost" (Camp 2018, 46).

According to her, the conversational cost is avoided by exploiting the gap between what has happened in the conversation and what participants are prepared to acknowledge.

Camp claims that this gives reasons for contending that there are things that are believed but not accepted (in the sense of Stalnaker 2002 inter alia, see Camp 2018, 56). Instead, we will argue that this motivates distinguishing acknowledging within "mere" accepting. Thus, fully successful cases of insinuation manage to get their message into the "common ground" without being acknowledged. The phenomena of insinuation and deniability thus vindicate, in turn, the need to distinguish between such a broad conception of common ground and a more restricted layer that we label conversational record (here departing again from Camp's own suggestion). This latter notion would encapsulate the consequences of our linguistic actions that Camp was calling "conversational cost."

A full-fledged theoretical elaboration of deniability (credible or not) will have to articulate the nature of acknowledging versus mere accepting, and the conversational record versus the general common ground. In order to do so, it is important to consider cases where the relevant consequences are not merely conversational—like fines and convictions in the courtroom. This may help clarify, for instance, the role of the audience's responses in successful insinuations. Reflection on these cases reveals that the relevant audience may not be the most immediate hearer—in the bribe case, e.g., it could be someone other than the police officer, say, some judge who would have to assess the case. In this sense, relevant deniability may be relative "to a virtual audience," "outsiders or authorities" (Lee & Pinker 796). Note that we're not solely interested in these strict and codified contexts. Rather, our aim is to do justice to the heterogeneous character of plausible deniability, which encompasses a wide variety of cases (cooperative and non-cooperative, private and public, formal and informal, etc.), including those in which the "conversational costs" that a speaker may aim to avoid are much more diffuse and elusive.

Note, in conclusion, that conceiving of plausible deniability as broader than credible deniability has the potential to shed further light on related phenomena such as lying and misleading, and their potential normative upshot (see Saul 2012, Stokke 2018), as well as the recent debate on counterspeech strategies against toxic discourse (see Langton 2018).

### References

**Camp**, Elisabeth (2018). "Insinuation, common ground, and the conversational record". In: New Work on Speech Acts. Ed. by Daniel Fogal, Daniel W. Harris, and Matt Moss. Oxford: Oxford University Press, pp. 44–66.

<sup>\*</sup> https://en.wikipedia.org/wiki/Let%27s\_Go\_Brandon#Politics

**Langton**, R. (2018). Blocking as counter-speech. In: New Work on Speech Acts. Ed. by Daniel Fogal, Daniel W. Harris, and Matt Moss. Oxford: Oxford University Press, 144-156.

**Lee, James and Steven Pinker** (2010): "Rationales for Indirect Speech: The theory of the strategic speaker," Psychological Review 117: 785–807.

**Mazzarella**, Diana et al. (2018). "Saying, presupposing and implicating: How pragmatics modulates commitment". In: Journal of Pragmatics 133.133, pp. 15–27.

**Saul**, J. M. (2012). Lying, misleading, and what is said: An exploration in philosophy of language and in ethics. Oxford University Press.

Stalnaker, Robert (2002). "Common Ground," Linguistics and Philosophy 25: 701–21.

**Stokke**, Andreas (2018). Lying and insincerity. Oxford University Press.

Viebahn, Emanuel (2020). "Lying with presuppositions". In: Noûs 54.3, pp. 731–751.



# A NOTE ON THE MEANING OF 'I'

# Tadeusz Ciecierski, Jakub Rudnicki

# University of Warsaw

According to a widespread belief, the reference of the first-person indexical 'I' is given by a simple rule that links the reference of a particular use of 'I' to the speaker-parameter of the context of use. This belief has been put into question as a result of numerous arguments provided by Maximillian de Gaynesford (de Gaynesford 2006). De Gaynesford points out three kinds of problems connected with the idea of the simple rule regulating the uses of 'I': (i) the diversity of formulations of the simple rule, (ii) problems connected with the exact logical form of the rule, (iii) problems connected with the user-producer ambiguity that can be identified in the case of most formulations of the simple rule.

Ad. (i) We argue that the diversity of formulations cannot by itself be taken as an argument either against the existence of the simple rule or against the correctness of one of the existing formulations of the simple rule.

Ad. (ii) We argue that one should consider as the most reasonable candidate for the logical form of the simple rule the following scheme:

(R)  $\forall x \forall y [x \text{ is the use/production of 'I'} => (\exists z \text{ x refers to } z => (\forall z' \text{ x refers to } z' <=> z' = up(x))]$  where up is the function that assigns an agent to the act of using or producing 'I'.

Ad. (iii). We discuss various counterexamples to user and producer oriented versions of (R). Consider, for instance:

Sentry. "When a king sends his sentry to another king, who delivers his message 'I surrender, let's discuss terms', it may be the sentry who is speaking, but the king is the agent of the context and the one who is really communicating with the receiving king. That's why the receiving king doesn't think it's the sentry who is surrendering, but the enemy king." (Dodd and Sweeney 2010: 344)

Office hours. "Joe is not in his office one day and Ben notices that a number of students keep approaching his door and knocking. They then stand around and look bemused for a while before leaving. Taking pity on these poor souls wasting their time, Ben decides to attach his "I am not here today" note to Joe's door". (Corazza, Fish, Gorvett 2002, 5)

Executor. "Suppose I am the executor of your will and thus sole user of the signed requests you have produced legitimizing various necessary funerary arrangements—'I wish that my body be given to science.'" (Gaynesford 2006, 41)

We critically discuss attempts to deal with such counterexamples offered respectively by intentionalism ((Predelli, 1998) and conventionalism (Corazza, Fish, Gorvett 2002) and argue against both accounts. At the very same time, we argue that potential counterexamples do not constitute a homogeneous class:

some cases (Office hours) should be explained away as cases of demonstrative uses of indexicals (cf. Kaplan (1989)).

Next, we propose a formulation of the rule that successfully unifies the analysis for the non-demonstrative uses of "I". The proposed version of the rule is:

(RR – short version) The referent of the use of "I" is the object (person) responsible for the speech act of which that use of "I" is a part of.

or more specifically:

(RR – long version) The referent of the use of "I" is the object (person) responsible for the satisfaction of the felicity conditions of the speech act of which that use of "I" is a part of as being the speech act of the particular illocutionary type.

The idea behind the formulation of the simple rule is that the referent of a particular token of 'I' is the person responsible for the speech act containing the token in question. We analyze the notion of responsibility in terms of commitments and obligations connected with particular kinds of speech acts. From the viewpoint of (RR), the problematic scenarios like Sentry and Executor receive an intuitive analysis: both the king and the author of the will are persons responsible for the satisfaction of the felicity conditions involved in cases of particular speech acts that contain tokens of 'I'.

### References

**Corazza**, E.; Fish, W.; & Gorvett, J. (2002). Who Is I? Philosophical Studies 107 (1):1-21. de Gaynesford, M. (2006). I: The Meaning of the First Person Term. Clarendon Press.

**Dodd, D. & Sweeney, P.** (2010). Indexicals and utterance production. Philosophical Studies 150 (3):331-348.

**Kaplan**, D. (1989). Demonstratives: An Essay on the Semantics, Logic, Metaphysics and Epistemology of Demonstratives and other Indexicals. In Joseph Almog, John Perry & Howard Wettstein (eds.), Themes From Kaplan. Oxford University Press. pp. 481-563.

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

Predelli, S. (1998): 'I am not Here Now'. Analysis 58(2), 107–115.



# A Puzzle about Mental Lexicons and Semantic Relatedness: Why Irregular Polysemy is a Type of Homonymy

# ALICE DAMIRJIAN

STOCKHOLM UNIVERSITY

This talk is on the topic of lexical ambiguity, and aims to tackle the question Why do irregular polysemes enjoy a processing advantage compared to homonyms? It will do so in two steps: First, I will criticize the popular view that the reason why irregular polysemes enjoy a processing advantage is because the meanings associated with irregular polysemes are represented and stored in single lexical entries in our mental lexicons, whereas the meanings associated with homonyms are represented in separate lexical entries. Views that accept this explanation will be called one-entry theories of irregular polysemy (an umbrella term intended to cover both so-called over-specificationist and under-specificationist accounts of irregular polysemy). What I will argue is that one-entry theories fail to predict and account for the fact that processing is facilitated in varying degrees. Secondly, I will defend a separate entries account of irregular polysemy according to which irregular polysemes are represented and stored just as homonyms are in our mental lexicons. I will then survey some alternative explanations for why facilitation effects in processing occur, but ultimately argue that such effects are best explained with reference to our metalinguistic beliefs about meaning relations.

It is becoming increasingly common within the study of ambiguity resolution, processing, and representation to distinguish between different types of lexical ambiguity, two of which are irregular polysemy and homonymy. It is usually assumed that our judgments of semantic relatedness are what distinguishes between the two, such that homonymy is a lexical ambiguity between semantically unrelated (or only weakly related) meanings, whereas irregular polysemy is a lexical ambiguity between highly to moderately related meanings (standardly referred to as senses). Now, what has been observed by psycholinguists is that semantic relatedness between associated meanings correlates with observable differences in processing, such that irregular polysemes enjoy a processing advantage compared to homonyms. Having multiple related meanings, it seems, facilitates processing and word recognition (see Eddington & Tokowicz, 2015), and the more related two meanings are judged to be, the more prominent the facilitation effects are. So, the degree to which the meanings associated with an irregular polyseme are judged as related varies, and ambiguous expressions can be placed along a continuum stretching from those associated with highly related meanings to those associated with completely unrelated meanings (Klepousniotou & Baum, 2007).

The fact that irregular polysemes enjoy a processing advantage is usually taken as evidence for a psychologically real distinction between homonymy and irregular polysemy. It is then hypothesized, and often accepted, that this must reflect a difference in mental representation. The hypothesis is that our judgments of relatedness track a difference in mental representation, and it is these differences in mental representation which explain why facilitation effects occur. Or so the one-entry theorists claim (see Falkum & Vicente, 2015; Vicente, 2018).

But my criticism is this: One-entry theories of irregular polysemy fail to predict and explain why a continuum is observed. Both semantic relatedness and processing facilitation is a matter of degree; meanings can be more or less related and processing can correspondingly be more or less facilitated. As observed by Klepousniotou and Baum (2007), the line between an expression's being polysemous or homonymous is fine-grained. However, the distinction as construed by one-entry theorists is not fine-grained – it is sharp – either an expression is associated with one (polysemous) lexical entry or several (homonymous) ones. On their picture, neither irregular polysemy nor homonymy comes in degrees.

So, the one-entry theorists face a puzzle, what I call the Continuum Puzzle: If we choose to explain why the facilitation effects occur by stipulating that it is because the meanings associated with irregular polysemes are stored under one and the same lexical entry, then we have problems explaining the fact that facilitation effects occur in varying degrees. We will also make the implausible prediction that minimal, single case differences in degrees of semantic similarity will settle whether an ambiguity is represented in a subject's mental lexicons as a single polysemous lexical entry or multiple homonymous ones.

In the second half of the talk, I will argue that one-entry theorists are mistaken in assuming that we must look to the structure of our mental lexicons in order to explain why people judge meanings as related and why such judgments affect their processing. Could there not be alternative explanations out there? One-entry theorists do not seem to think so. They often accuse proponents of separate entries accounts of being incapable of explaining how it can be the case that our judgments of semantic relatedness facilitate processing (see Brocher et al., 2018; Carston, 2021). For if a separate entries account were correct, it is argued, irregular polysemes would have to be processed just like homonyms. This is ill-considered, as advocates of separate entries accounts just need to provide a different explanation. In fact, several such alternative explanations are possible, and have been suggested (see Devitt, 2021; Löhr, 2021).

The explanation I will propose is that the relevant difference between irregular polysemes and homonyms is precisely the fact that we judge the meanings associated with irregular polysemes as related. And we judge them to be related because we believe that they are. This view rejects the idea that our judgments of semantic relatedness reflect a fact about the structure of our mental lexicons. Instead, it argues that these judgments reflect our metalinguistic beliefs, and it is precisely the fact that we have these beliefs which explains why processing is affected. As Koskela and Murphy point out: "[I]t may be questioned whether people's metalinguistic reasoning about meaning relations can reflect how lexical meaning is mentally represented" (2006, p. 743). What such metalinguistic reasoning can reflect, however, is what metalinguistic beliefs about meaning relations people have, and that is precisely what I suggest we take them to reflect.

### References

**Brocher, A., Koenig, J.-P., Mauner, G., & Foraker, S.** (2018). About sharing and commitment: The retrieval of biased and balanced irregular polysemes. Language, Cognition and Neuroscience, 33(4), 443–466. https://doi.org/10.1080/23273798.2017.1381748

**Carston**, R. (2021). Polysemy: Pragmatics and sense conventions. Mind & Language, 36(1), 108–133. https://doi.org/10.1111/mila.12329

**Devitt**, M. (2021). Semantic polysemy and psycholinguistics. Mind & Language, 36(1), 134–157. https://doi.org/10.1111/mila.12327

**Eddington, C. M., & Tokowicz, N**. (2015). How meaning similarity influences ambiguous word processing: The current state of the literature. Psychonomic Bulletin & Review, 22(1), 13–37. https://doi.org/10.3758/s13423-014-0665-7

**Falkum, I. L., & Vicente, A**. (2015). Polysemy: Current perspectives and approaches. Lingua, 157, 1–16. https://doi.org/10.1016/j.lingua.2015.02.002

**Klepousniotou, E., & Baum, S. R.** (2007). Disambiguating the ambiguity advantage effect in word recognition: An advantage for polysemous but not homonymous words. Journal of Neurolinguistics, 20(1), 1–24. https://doi.org/10.1016/j.jneuroling.2006.02.001

**Koskela, A., & Murphy, M. L.** (2006). Polysemy and Homonymy. In K. Brown (Ed.), Encyclopedia of Language & Linguistics (2nd ed.) (pp. 742–744). Oxford: Elsevier. https://doi.org/10.1016/B0-08-044854-2/01045-2

**Löhr**, G. (2021). Does polysemy support radical contextualism? On the relation between minimalism, contextualism and polysemy. Inquiry, 1–25. https://doi.org/10.1080/0020174X.2020.1868329 Vicente, A. (2018). Polysemy and word meaning: An account of lexical meaning for different kinds of content words. Philosophical Studies, 175(4), 947–968. https://doi.org/10.1007/s11098-017-0900-y



The notion of aboutness (also: subject matter, topic) has been recently deployed to explain several linguistic phenomena that the standard possible world semantics has been unable to account for (hyperintensionality, logical omniscience, epistemic closure, partial truth, conceptual content, imagination etc.). Yablo (2014) understands aboutness as "the relation that meaningful items bear to whatever it is that they are on or of or that they address or concern, (Yablo 2014: 1). It is an "independent factor in meaning, constrained but not determined by truth-conditions" (Yablo, 2014: 2). Following Yablo, 2C semanticists (Berto 2020, Hawke 2016, Spolaore and Plebani 2021) propose that the thick proposition expressed by p encompasses not only p's truth-conditions (the set of worlds where p is true) but also its subject matter. In this paper I appeal to Travis cases to motivate the idea of super thick propositions. I argue that to account for this data we need to add another component to our semantics, that is, if we want to capture what an utterance of p in context c is intuitively about ("speaker aboutness"). A super thick proposition expressed by p encompasses (i) p's truth-conditions (i.e., the thin proposition), (ii) its sentential topic that is determined by its focus structure (i.e., "sentence aboutness") and (iii) the topic of discussion determined relative to domain goals.



# AUTHORITY, CONVENTION AND GROUP ASSERTION

Daniel Garibay-Garcia

UNIVERSITY OF ST. ANDREWS

Recently there has been a lot of interest in collective phenomena; group assertion and group testimony among those. In everyday life, it is very common that we attribute assertions and lies to groups like different institutions, governments, companies, and research groups. For example, we may say that Coca-Cola asserted that there is no relation between its drinks and obesity; there is an intuitive difference between that and having a room full of people that coincidently have all said that there is no relation between sugar and obesity. What Coca-Cola did was a group assertion, while the latter is a collection of individual assertions. If we are interested in these cases and their differences, there are many questions that we may try to answer: what is the nature of group assertion? Is it the same kind of speech act as individual assertion? What is the relation between group and individual assertion? We may also be interested in its epistemic normativity, are there any epistemic norms or commitments groups have when asserting? Or we might be interested in answering when an utterance is a group assertion. In one sense, group assertions are performed by individuals, they are the ones who do the physical actions required to assert, so, there is a question about which utterances are group assertions and why. In this work, I answer the last question, but hopefully, my account also has some important insights on the relationship between individual assertion and group assertion and the nature of assertion in general. I hold a conventionalist account of assertion, according to which an utterance is a group assertion if it follows the accepted procedure. These procedures (in cases of speech acts) can be understood in general as uttering some words in the right context by the right individuals. The key component of my account is having the authority to assert on behalf of the group.

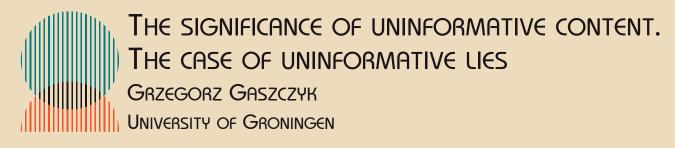
In the traditional debate about individual assertions, intentions or conventions are the elements that explain why a particular utterance is the speech act it is (see Kölbel, 2010; Stalnkaer, 1999). Intentionalist proposals hold that assertion is characterized by a particular communicative intention. Conventionalist proposals hold that it is characterized by a convention. In the debate about group assertion and group testimony there are different interesting proposals (see Lackey, 2018; Tollefsen, 2009; Fricker, 2012). Some of them seem to be intentionalist proposals, but in general, the debate hasn't been framed in this fashion. My objective in this work is two-folded. One aim is methodological. I try to bring some tools from previous debates about assertion to analyse group assertion. My other aim is to defend a conventionalist proposal of group assertion. Under this kind of proposal, an utterance is a group assertion if it follows the convention of this speech act. This kind of proposal does not claim that intentions aren't relevant for speech acts, but they are not essential, and they are not what defines the kind of act. Intentions are interesting and problematic in the case of group assertions and collective phenomena. Some believe that a particular kind of communicative intention is what distinguishes assertion (see Stalnaker 1999). And others believe that intentions are what distinguishes individual actions from group actions (see Searle1990, Habgood-Coote 2020). I aim to defend a proposal that doesn't require these kinds of intentions. But it is compatible with groups having them.

There are several important cases and phenomena that are important to account for when explaining group assertion. There are cases of institutions or corporations making assertions. These are commonly made by spokespersons with some guidelines in official settings. We also have assertions made by research groups. These are commonly made in papers or presentations and they involve different levels of coordination and deliberation to produce the assertions. There are also cases of less organized groups coordinating to give testimony or assert. We need to explain what these cases have in common, when are these genuine cases of group assertion, and why. There is also the important phenomenon of silencing, which gains complexity in the case of groups: groups can be silenced, but also just their members or the spokespersons speaking or trying to speak on behalf of the group.

My Proposal, the conventional account of assertion, can be summarized in the following conditions:

- I. There is a set S of speakers such that the utterance that means that P is performed by the members of S and at least one member of the set has the intention to express that P.
- II. S licitly has the authority to assert on behalf of the group about subject D.
- III. S's authority is recognised by the relevant community.
- IV. P belongs to subject D and doesn't conflict with the discourse of the group.
- V. S performs the utterance in virtue of S's authority.
- VI. The utterance takes place in the appropriate scenario or channel.

Against Lackey (2018), Fricker (2012), and Tollefsen (2009), I hold that my proposal has a better explanation of the different cases of group assertion and also a better explanation of silencing in group assertion. There are cases of collaborative work where the members of the group decide that each member will produce a portion of the work on their own and all the parts together will form a single piece of work. These assertions are made by one individual as part of collaborative work with no collaboration from the other participants to produce the assertion. I hold that these are genuine group assertions and previous accounts of assertion fail to explain them. I also hold that they fail to accurately account for silencing in group assertion. Previous accounts fail to accurately distinguish between different cases of silencing: Spokesperson silencing, group silencing and group members silencing. Finally, I hold that there is an important distinction between group locutions (utterances made by groups) and group illocutions (speech acts made by groups) that these accounts fail to make. I argue that group locutions and group illocutions are independent. Groups can assert through collections of individual utterances, these are cases where we have a group illocution but no group locution. And not all cases of utterances produced by groups (utterances where the group coordinates to produce the utterance) are cases of group assertion, these are cases where we have a group locution and no group illocution.



Assertions are primarily a source of information. It is unsurprising then that one way to theorise about assertion is to put its informativeness in its centre, i.e., think about them as essentially informative speech acts. However, such formulation can receive two readings:

STRONG INFO Assertions essentially deliver new information. WEAK INFO Assertions essentially deliver information.

The straightforward way of being informative is to introduce to the common ground an unknown piece of information, like asserting to my friend p: "There is a box of cookies on the table." STRONG INFO counts this as a proper assertion. WEAK INFO maintains that asserting p can be proper even if p is already part of the common ground. Thus, it makes space for uninformative assertions.

Assertions perform certain unique functions. The one I will focus on concerns lying, i.e., only assertions are lie-prone (Dummett 1981; Stainton 2016). Moreover, so far lies have been analysed as informative along STRONG INFO. However, I will argue that we can lie with uninformative content.

## 1. Informativeness of primary and secondary content

STRONG INFO follows from many theories of assertion (Searle 1969; Stalnaker 1978; García-Carpintero 2004, 2020; Pagin 2011; Kölbel 2011; cf. Farkas 2020). This idea is expressed in many ways. On the one hand, some propose a similar condition for a proper assertion, i.e., an assertion that p is proper if p is not already common knowledge; for Searle (1969, p. 66), this is one of the "preparatory conditions" for asserting; for Stalnaker (1978, pp. 88-89 in 1999), this is the first "principle" about assertion; for Farkas (2020, p. 11), this is one of the "default pragmatic assumptions" accompanying assertions. On the other hand, some argue that STRONG INFO follows from our conversational patterns. García-Carpintero (2004, 2020), for instance, claims that one can be criticised for asserting something that is commonly known; this follows from an observation that "when it is correct to presuppose p, it is incorrect to assert it" (2020, p. 23).

Notice that not only assertions are informative. This feature is ascribed to declarative speech acts in general (Searle 1969; Alston 2000; Kissine 2013), i.e., to primary or at-issue content (Simons et al. 2010). Just as it is improper to reassert, it is equally so for conjecturing, swearing, or admitting. Furthermore, informativeness can be naturally extended to some secondary or not-at-issue contents, i.e., to conventional implicatures (Potts 2005, 2015) that can be treated as secondary assertions, i.e., as assertions that are carried by other speech acts, and to informative presuppositions (Potts 2015; Beaver et al. 2021).

### 2. Uninformativeness of primary and secondary content

Interestingly, uninformative content, just as informative, can be mapped along two dimensions, i.e., primary and secondary content. Firstly, not all declaratives are informative. One, there are unique speech

acts, like reminding, that do not provide any new information into common ground. Reminding is not just reasserting; it is inappropriate to assert when one is in a position to remind, for instance, if I would reassert "There is a box of cookies on the table," you can rightly complain and respond "I know, you already told me that!" However, it is appropriate to remind it (often in a performative way, like "Just to remind you, p"). Two, in certain circumstances, some speech acts can be performed in an uninformative way. Consider guaranteeing. I can firstly assert that p, and later on guarantee that p. However, guaranteeing in this sense does not deliver any new information, rather it reinforces the content that was already stated. Repeating the same content may not be informative in a standard way, however, it is not redundant since it can perform other functions, like reassuring the

audience to act on its basis and commit the speaker even stronger to what is communicated (Crone 2019). Secondly, the paradigmatic case of uninformative content on the not-at-issue level are standard cases of presuppositions that are treated as commonly known information. Table 1 summarises the paper to this point.

	Primary (at-issue) content	Secondary (not-at-issue) content		
Informative content	assertions, conjectures, retractions, predictions	conventional implicatures, informative presuppositions		
Uninformative content	remindings, guaranteeing	presuppositions		
Table 1. Examples of informative/uninformative content along two dimensions, i.e., primary/secondary content.				

## 3. Lying with informative content

The widely accepted definition of lying states that lies are insincere assertions (Carson 2006; Sorensen 2007; Fallis 2009; Stokke 2017; 2018; Viebahn 2020, 2021; Marsili 2020). Consider Stokke's says-based and Viebahn's commitment-based definitions of lying. Both assume STRONG INFO, however, what they count as assertions is very much different. Stokke (2018) restricts the scope of lies to what is said, while Viebahn (2020; 2021) is much more permissive and allows for lying with what is implicated or presupposed, as long as one commits to the expressed content. Crucially, both views restrict lies to informative content. Stokke because his definition is based on a Stalnakerian (1978) theory of assertion that maintains that asserting something that already belongs to the common ground is infelicitous. Viebahn (2021, p. 301) explicitly argues that we lack commitment in uninformative cases, like when we presuppose p while taking p to be common ground.

### 4. Lying with uninformative content

I propose to extend the possibility of lying to the second tier of Table 1. Such lies can be dubbed follow-up lies: one can make a follow-up lie about p in a context c only if someone already lied about p in c. To perform a follow-up lie one must lie by means of a speech act that can transfer information that was already stated before. Following STRONG INFO, this cannot be done by assertions or conventional implicatures. However, there are alternatives. Firstly, consider the primary content. Lying is possible with speech acts at least as committal as assertions (Marsili 2020; Viebahn 2021). If there are no cookies, I am lying by asserting p: "There is a box of cookies on the table." Now, imagine that you inquire further and I

say "I guarantee you that p;" I am lying again even though p is uninformative since it was already asserted and accepted into common ground. The same concerns other uninformative declaratives. Secondly, let us look at not-at-issue content, specifically presuppositions. I can presuppose the content that was already asserted, by saying, for instance, "The cookies are tasty!" Again, if there are no cookies, I am lying by presupposing it just as I firstly lied by asserting. Uninformative presuppositions can be just as committal as informative ones (Potts 2015; Peters 2016). I propose a properly modified commitment-based definition of lying that includes uninformative lies.

### References

Alston, W. P. (2000). Illocutionary acts and sentence meaning. Cornell University Press.

https://doi.org/10.7591/9781501700422

**Beaver, D. I., Geurts, B., & Denlinger, K.** (2021). Presupposition. In E. N. Zalta (Ed.), The Stanford Encyclopedia of Philosophy (Spring 2021). Metaphysics Research Lab, Stanford University. https://plato.stanford.edu/archives/spr2021/entries/presupposition/

**Carson**, T. L. (2006). The Definition of Lying. Noûs, 40(2), 284–306. https://doi.org/10.1111/j.0029-4624.2006.00610.x

**Crone**, P. (2019). Assertions of Clarity & Raising Awareness. Journal of Semantics, 36(1), 53–97. https://doi.org/10.1093/jos/ffy006

Dummett, M. (1981). Frege: Philosophy of Language. Harvard University Press.

Fallis, D. (2009). What Is Lying? The Journal of Philosophy, 106(1), 29-56.

https://doi.org/10.5840/jphil200910612

Farkas, D. F. (2020). Canonical and non canonical questions. Manuscript.

García-Carpintero, M. (2004). Assertion and the Semantics of Force-Makers. In C. Bianchi (Ed.),

The Semantics/Pragmatics Distinction (pp. 133–166). Stanford: CSLI Lecture Notes, 155.

**García-Carpintero**, M. (2020). On the Nature of Presupposition: A Normative Speech Act Account. Erkenntnis, 85(2), 269–293. https://doi.org/10.1007/s10670-018-0027-3

Kissine, M. (2013). From Utterances to Speech Acts. Cambridge University Press.

**Kölbel**, M. (2011). Conversational Score, Assertion, and Testimony. In J. Brown & H. Cappelen (Eds.), Assertion: New Philosophical Essays (pp. 49–77). Oxford University Press.

Marsili, N. (2020). Lying, speech acts, and commitment. Synthese.

https://doi.org/10.1007/s11229-020-02933-4

**Pagin**, P. (2011). Information and Assertoric Force. In J. Brown & H. Cappelen (Eds.), Assertion: New Philosophical Essays. Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199573004.003.0005

**Peters**, S. (2016). Speaker commitments: Presupposition. Semantics and Linguistic Theory, 26, 1083. https://doi.org/10.3765/salt.v26i0.3951

**Potts**, C. (2005). The logic of conventional implicatures. Oxford University Press. 10.1093/acprof:oso/9780199273829.001.0001

**Potts**, C. (2015). Presupposition and Implicature. In S. Lappin & C. Fox (Eds.), The Handbook of Contemporary Semantic Theory (pp. 168–202). John Wiley & Sons, Ltd. https://doi.org/10.1002/9781118882139.ch6

**Searle**, J. R. (1969). Speech Acts: An Essay in the Philosophy of Language. Cambridge University Press. https://doi.org/10.1017/CBO9781139173438

Simons, M., Tonhauser, J., Beaver, D., & Roberts, C. (2010). What projects and why. Proceedings of SALT, 20, 309–327. https://doi.org/10.3765/salt.v20i0.2584

**Sorensen**, R. (2007). Bald-Faced Lies! Lying Without the Intent to Deceive. Pacific Philosophical Quarterly, 88(2), 251–264. https://doi.org/10.1111/j.1468-0114.2007.00290.x

**Stainton**, R. J. (2016). Full-On Stating. Mind & Language, 31(4), 395–413. https://doi.org/10.1111/mila.12112

**Stalnaker**, R. (1978). Assertion. In P. Cole (Ed.), Syntax and semantics (Vol. 9). Academic Press; (References are to the reprint in R. Stalnaker, Context and content. Oxford University Press. 1999. pp. 78–95.).

**Stokke**, A. (2017). II—Conventional Implicature, Presupposition, and Lying. Aristotelian Society Supplementary Volume, 91(1), 127–147. https://doi.org/10.1093/arisup/akx004

**Stokke**, A. (2018). Lying and Insincerity (Vol. 1). Oxford University Press. https://doi.org/10.1093/oso/9780198825968.001.0001

**Viebahn**, E. (2020). Lying with Presuppositions. Noûs, 54(3), 731–751. https://doi.org/10.1111/nous.12282 **Viebahn**, E. (2021). The Lying-Misleading Distinction: A Commitment-Based Approach. The Journal of Philosophy, 118(6), 289–319. https://doi.org/10.5840/jphil2021118621



Are non-human animals (henceforth 'animals') agents? Following philosophical precedent in (for example) Immanuel Kant (1996; 1781/1787/1965) many answer 'no'. True action, these philosophers say, is dependent on rationality in a demanding sense. It is 'acting for reasons', and, they think, it requires the conceptual capacity to self-ascribe thoughts in critical reflection: that is, self-consciousness (McDowell, 1996; 2007a; 2007b). Call such philosophers 'Intellectualists'. Others, following philosophical precedent in (for example) David Hume (1740/1975), answer 'yes'. Catching a frisbee, or deliberately turning one's head towards the source of a noise are actions — and there is no reason to deny agency to animals like dogs, or to think that sophisticated conceptual and metarepresentational abilities are required. Call such philosophers 'Anti-Intellectualists'.

One thing at stake between Intellectualists and Anti-Intellectualists is our understanding of the place of human mentality, broadly construed, in the natural world. Intellectualists find little or no common ground between human action and what animals can do (McDowell, 1996; 2007a; 2007b; Korsgaard, 2009; 2014). Anti-Intellectualists, in contrast, see human action as a more sophisticated version of the same basic capacity, grounded in biological function, shared with other animals (Glock 2019; Burge, 2009; Fulda, 2017; Jones, 2016). Some argue that Intellectualists deny that human mentality's emergence from simpler forms of mentality could be explained in a philosophically cogent way. This has been dubbed 'continuity scepticism' (e.g. Bar-On, 2013). Intellectualists themselves explicitly hold that human mentality is irreducibly normative, and so 'causal' explanations of agency drawing on the concepts and frameworks of natural science are misplaced (Brandom, 1994: 156; McDowell, 1996: 8fn7, 71fn2 & passim; Larmore, 2002: 201; Korsgaard, 2014). Some philosophers sympathetic to this view think that human mentality, and hence agency, is 'permeated' by rationality (e.g. McDowell, 2007b: 368). This poses a problem for the Anti-Intellectualist ambition to understand human agency as an elaboration of a capacity, grounded in biology, and shared with animals. No non-normative explanation is appropriate, they argue. And no animal species could share the essentially rational capacity for agency in humans. Hence, continuity scepticism appears to follow. But surely there is something more to be said about how human agency emerges from simpler nonconceptual precursors, as everyone agrees that it does. Human agency is, after all, one phenomenon among many in the natural world. And, as scientific inquiry has proceeded, more and more such phenomena have been subject to natural scientific explanation. So there had better be a very convincing reason to believe that human agency will not follow suit.

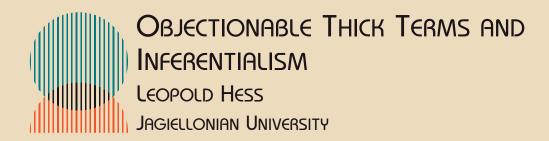
In this paper I propose an account that attempts to satisfy both (i) the Intellectualist's insistence on the essential rationality of human agency, and (ii) the Anti-Intellectualist's demand to explain the emergence of human agency from non-rational animal forms of agency.

I argue that the proposed account can satisfy simpler analogues of two constraints that Intellectualists place on human agency. Intellectualists require that an account of human agency capture both the

aforementioned normativity, and the 'holism' or 'unity' of the mental (as emphasised by Donald Davidson (1980)). The normativity of thought requires that a thinker be 'responsive' to rational reasons. McDowell writes that "the idea of rational response requires subjects who are in charge of their thinking, standing ready to reassess what is a reason for what, and to change their responsive propensities accordingly" (1996: 114). 'Unity', roughly speaking, is the unification of all a subject's conscious mental states into a single point of view. I focus on Intellectualists who follow the broadly Kantian idea that a capacity for self-consciousness is necessary for the normativity and unity of human mentality, and hence of agency. The Kantian Intellectualists, in short, conceive of human agency as a capacity to act for reasons that one can critically reflect on. And self-consciousness is necessary to explain the possibility of critical reflection that is essential to agency so conceived. But the very appeal to self-consciousness appears to rule out any explanation of how this demanding form of agency could have emerged from the pre-reflective, pre-conceptual capacities of animals. Self-consciousness is a demanding conceptual capacity, arguably out of reach of all but the most sophisticated creatures. It therefore imperils the Anti-Intellectualists' explanatory ambition to locate the origins of agency in the natural world.

I propose an account of pre-rational, pre-conceptual mentality ('Pre-Reflective Mentality') centred around a nonconceptual precursor to self-consciousness that shares a structural-functional organisation with human mentality (following work by, e.g.: Bermúdez, 1998; Musholt, 2013; 2015; 2019). In my account, nonconceptual self-consciousness plays a role analogous to that which conceptual self-consciousness plays in human mentality. This offers a potential way to accommodate (i) and (ii) when applied to the capacity for agency. I draw on recent philosophical and empirical work exploring the possibility of 'proto-rationality' in animals (Hurley 1998, 2003; Hurley & Nudds, 2006; Marcus, 2012; Proust, 2018; Camp, 2009). Precursors to the sort of 'middle ground' I propose between Intellectualism and Anti-Intellectualism are found in Susan Hurley (1998, 2003) and Eric Marcus (2012). Hurley (2003) argues that we can use the idea of nonconceptual content to construct a view of 'reasons' suitable for animal minds, at least in the domain of practical reasoning. Combining this work with the central notion of nonconceptual self-consciousness, I claim, can make it plausible how a unified pre-reflective subject can act for non-propositional and preconceptual 'reasons'. Therefore, my account promises to bring together recent philosophical and empirical work on the origins of rationality to show that Pre-Reflective Mentality can support a capacity for agency that can meet simpler analogues of normativity and unity.

Finally, I sketch how my account could provide the foundation for a possible emergence explanation. Both Intellectualism and Anti-Intellectualism, I suggest, share the assumption that to explain the emergence of human mentality it is necessary to posit shared capacities between humans and precursor mentalities. But there is no good reason to accept this assumption. There are the resources in evolutionary biology to make a non-capacity-sharing emergence explanation a viable theoretical option (Gould, 1991; Fitch, 2010). By rejecting capacity-sharing explanations, my strategy aims to add substance to McDowell's (1996) suggestive claim — elaborated by Matt Boyle (2016) — that human and animal mentality — and, as is my main focus, agency — stand to one another as two species of the same genus, though they strictly share no discrete capacities.



Thick terms, such as generous or rude combine description and evaluation. An important problem for accounts of their meaning are so-called objectionable thick terms (OTTs) - ones that some speakers may deem to express an inappropriate evaluation - e.g. lewd or blasphemous. In this paper I will propose to treat OTTs on the model of inferentialist semantics for slurs as proposed by Dummett and Brandom. This inferentialist model, involving objectionable inferences, is inadequate for paradigmatic slurs which are usually assumed to have extensionally equivalent neutral counterparts (NC-slurs), but I will argue it works well for OTTs, as well as so-called gendered slurs, and possibly for a broader range of ideologically distorted expressions. The crucial issue is that of determining reference. The inferentialist model cannot support the reference of NC-slurs. However, OTTs and gendered slurs can be interpreted (by the objectors) as lacking a determinate reference, which can be accounted for in the inferentialist model. This model also supports some important intuitions concerning thick terms and solves some puzzles that have been debated in the literature.



## LINGUISTIC MISTAKES AND THE PUBLIC LANGUAGE

## Aleksi Honkasalo

TAMPERE UNIVERSITY

In the debate concerning the normativity of meaning, there appears to be a common assumption that norms of meaning depend on what speakers mean by their expressions (Eg. Wikforss 2001, Whiting 2016). However, this assumption has some unintuitive consequences regarding linguistic mistakes. If Jane means ancient by 'arcane', it seems she does not misuse the term simply by applying it not non-arcane things. Indrek Reiland (2021) has recently argued that this assumption is tied to the Davidsonian individualist view of language. He on the other hand advocates a public view of language and argues that from this perspective it does make sense to call Jane's use of the term 'arthritis' a linguistic mistake.

However, while the appeal to the public language perspective may help to make room for Jane to be mistaken about the public meaning of 'arcane', it may not be enough to show that she made a linguistic mistake. After all, from the public language perspective, rather than making a linguistic mistake, she successfully expressed the false proposition that a non-arcane thing is arcane. If, on the other hand, linguistic mistakes are understood as meaning something other than the public meaning, we have to concede that Jane can mean something nonstandard by 'arcane'. This is a significant concession to the individualist picture and might even lead to the trivialization of the public language view.

How can this unintuitive conclusion be avoided? Reiland appeals to the idea that expressions have rules for their use. This way it might remain possible to maintain both the public view of language and the possibility of linguistic mistakes by arguing that Jane used 'arcane' with its standard meaning, but since she broke the rules of its use, she also made a linguistic mistake. But what are these rules of use? Reiland builds on David Kaplan's ideas and proposes that expressions with truth-functional content like 'arthritis' could have similar rules as non-truth-functional expressions such as, 'goodbye', 'ouch', and 'oops', which only have conditions of correct use.

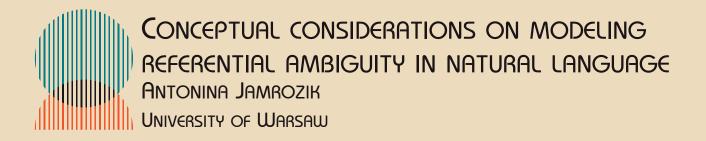
While the conditions of correct use and misuse of expressions like 'goodbye' might in theory be explicated along the lines Kaplan suggested, the conditions for expressions like 'arcane' remain vague despite Reiland's attempts to clarify them (2021, 2022). If rules of use of 'ar-cane' are something to the effect of: 'use 'arcane' to mean arcane', then the dilemma reap-pears. Either the public language determines what 'arcane' means, and Jane cannot help but mean arcane by 'arcane', or Jane can mean something non-standard by 'arcane', which begs the question of what semantic significance, there is to the fact that her meanings do not align with public ones? On the other hand, if the conditions of correct use depend on whether the object in fact is arcane or not, we run the risk of misclassifying some false uses as linguistic mistakes.

## References

Reiland, Indrek (2021). "Linguistic Mistakes". Erkenntnis:1–16.

Reiland, Indrek (2022) "Rules of Use". Mind & Language, Early View

**Whiting**, Daniel (2016) "What is the normativity of meaning?" Inquiry, 59 **Wikforss**, Åsa (2001) "Semantic normativity". Philosophical Studies, 102



Do questions, disjunctions and ambiguity have anything in common with each other? A positive answer to this question is put forward in this talk, and I hope to present the case for it successfully. My main goal here is to establish a conceptual framework that will allow representing different kinds of ambiguity, but especially referential ambiguity, as generating inquisitive content. The latter term is used in the sense that it is used in Inquisitive Semantics (Ciardelli et al. 2018).

First, I will give a brief overview of the phenomenon of ambiguity in natural language. I go through different kinds of ambiguity, moving towards the one that is of most interest to me -- referential ambiguity. The latter term is coined to describe the special status of anaphoric pronouns. As the claim that anaphors are ambiguous is not an uncontroversial one, I put forward an argument for it. The argument is based on two premises. First, there are considerations regarding the undoubtedly ambiguous cases of anaphora, such as the following:

- (1) The city council denied the demonstrators a permit because they feared violence.
- (2) The city council denied the demonstrators a permit because they advocated violence. (example form Winograd 1972, adapted by Kehler 2002)

In (1) the anaphoric pronoun "they" refers to the city council, and in (2) "they" refers to the demonstrators. However, the syntactic structure of the sentences is identical, so there is nothing in the syntax itself that secures those coreferences. I argue that cases like this motivate the claim that actually all anaphora is ambiguous, and pragmatic mechanisms of disambiguation always play (more or less) prominent role in anaphora processing. Second, I point out that anaphoric constructions pass some of the linguistic tests designed to detect ambiguity (Zwicky and Sadock 1975). I finish this section of my talk with further justification as to why a new kind of ambiguity is needed in order to capture the nature of the sort of ambiguity that is generated by the use of anaphoric constructions. This is achieved by pointing out how other kinds of ambiguity are insufficient to fully capture the features of such constructions.

In the second part of the talk, I will dive into how the notion of "raising an issue" is conceptualised in the framework of inquisitive semantics. I will show that this conceptualisation can be used to represent at least some kinds of semantic ambiguity, including referential ambiguity. I propose a way to represent ambiguity as generating inquisitive content. It is first done for ambiguity in general, on a very coarse-grained level, and then for referential ambiguity of anaphoric expressions. The latter is explicated using the framework of Dynamic Inquisitive Semantics (Dotacil & Roelofsen 2015), especially the feature of this system that allows for raising an issue as to the identity of a discourse referent. In the last part of the talk I give a general argument for the claim that the treatment of anaphora that relies on coindexing the anaphor and its antecedent on the syntactic level does not capture many important features of anaphora. It is still a mainstream treatment of this phenomenon, with just a few notable exceptions (Bittner 2001, Dekker 2012).

I give an overview of how the account proposed in this talk compares to abovementioned theories and of the possibilities of further research that come with it.

## References

**Bittner**, M. (2001). Topical referents for individuals and possibilities. In Semantics and Linguistic Theory (Vol. 11, pp. 36-55).

Ciardelli, I., Groenendijk, J., & Roelofsen, F. (2018). Inquisitive semantics. Oxford University Press.

**Dekker**, P. J. (2012). Predicate logic with anaphora. In Dynamic Semantics (pp. 7-47). Springer, Dordrecht.

**Dotlačil, J., & Roelofsen, F.** (2019). Dynamic inquisitive semantics: Anaphora and questions.

In Proceedings of sinn und bedeutung (Vol. 23, No. 1, pp. 365-382).

Kehler, A. (2002). Coherence, reference, and the theory of grammar.

**Zwicky, A. M., & Sadock, J. M**. (1975). Ambiguity tests and how to fail them. In Syntax and Semantics volume 4 (pp. 1-36). Brill.

# LOCAL CONTEXTS FOR SITUATION SEMANTICS LUKAS LEWERENTZ UNIVERSITY OF GIESSEN

Situation semanticists have argued that definite descriptions and other definite terms such as pronouns are evaluated at situations, parts of possible worlds that can be very small. On this view, a definite description the F, when evaluated at some situation s, presupposes that there is exactly one individual that is F at s. If the presupposition is satisfied and there is exactly one F at s, the F denotes that individual; otherwise, the F fails to denote.¹ Since this view claims that definite terms trigger presuppositions, it needs to be supplemented with a theory of how presuppositions work. Whether the view is compelling depends crucially on whether it can be supplemented with a plausible (and ideally independently motivated) theory of presuppositions in such a way that the combined view makes plausible predictions. This is a challenge that situation semanticists so far have not done enough to address.² In this talk, I explain the challenge, I lay the groundwork for one way of addressing it, and I explore some of the problems that arise. One of the main tasks of theories of presuppositions is to explain how presuppositions project, that is, to explain how the presuppositions of a complex sentence depend on the pre- suppositions of its parts. Consider the following examples:

(1)

- a. Sarah stopped smoking.
- b. Sarah used to smoke cigars and Sarah stopped smoking.
- c. Either Sarah never smoked cigars or Sarah stopped smoking.

As is well-known, (1a) presupposes that Sarah used to smoke, and this presupposition disap- pears (doesn't project) when the sentence is embedded as in (1b) or (1c). (1b) and (1c) don't presuppose that Sarah used to smoke. Theories of presupposition projection try to account for such facts. Similarly, consider the following cases:

(2)

- a. The book is heavy.
- b. There is a book and the book is heavy.
- c. Either there isn't any book or the book is heavy.

Whereas (2a) presupposes that there is exactly one book (in some contextually relevant situation), this presupposition doesn't project onto (2b) or (2c). (2b) and (2c) don't presuppose that there is exactly one book. Theories of presupposition projection should also be able to account for those facts. However, it's not clear how a unified explanation of these projection facts could be achieved. It seems that (1b) doesn't presuppose that Sarah used to smoke because the first conjunct entails that she used to; and it seems that (1c) doesn't presuppose that Sarah used to smoke because the negation of the first disjunct entails that she used to. But note that there is a book doesn't entail that there is exactly one book, and neither does the negation of there isn't any book. As Mandelkern and Rothschild (2020) point out, this makes it difficult

to see how situation semanticists could provide a unified account of presupposition projection. If the relevant entailments explain why the presupposition of (1a) doesn't project onto (1b) and (1c), the same explanation can't be given for why the uniqueness presupposition of (2a) doesn't project onto (2b) and (2c), since the relevant entailments do not hold in latter cases.

In this talk, I develop one idea for how a unified account of such projection data could look like. One prominent approach to projection assumes that the presuppositions of an embedded sentence project unless they are entailed by the local context of that sentence. Roughly, a context of a sentence is a body of information that provides the background for interpreting the sentence. Whereas global contexts are contexts for unembedded sentences, local contexts are contexts for parts of sentences. In a conjunction A and B, the local context of B depends on what the global context is and on A; it incorporates information provided by A. How this works depends on one's conception of contexts. Following Stalnaker (1978), a context is often taken to be a set of possible worlds, namely, of those possible worlds that are compatible with the background information.<sup>3</sup> Given such a view, existing theories of local contexts predict that the local context of B in A and B is the set of those worlds of the global context where A is true.<sup>4</sup> But this wouldn't work for the situation semanticist. After all, the set of all worlds in which there is a book is true comprises worlds in which there is more than one book, and so the uniqueness presupposition of the second conjunct of (2b) would not be guaranteed to be entailed by its local context. Thus, we would fail to predict that the presupposition doesn't project. To solve this problem, I suggest that the situation semanticist adopt a different conception of contexts: I propose that a context is a set of situations rather than a set of worlds. Moreover, I assume that a context doesn't contain all the situations that are compatible with the background information; rather, it contains only the smallest situations compatible with it. The smallest situations in which there is one book contain only one book, and so the local context of the second conjunct of (2b) contains only one book (unless other books have been mentioned before), in which case the local context does entail the uniqueness presupposition triggered by the book. In the talk, I show how such minimal contexts can be determined pragmatically, by adopting and modifying the pragmatic algorithm suggested by Schlenker (2009). I then discuss whether the predictions this algorithm makes are plausible, by looking at disjunctions and at donkey sentences such as Every farmer who owns a donkey feeds the donkey. I argue that in both cases, the predictions raise serious questions about the underlying situation semantics, since the predictions are plausible only given non-standard assumptions about the semantics.

## References

Barwise, Jon and John Perry (1983). Situations and Attitudes. Cambridge: MIT Press.

Büring, Daniel (2004). "Crossover Situations". Natural Language Semantics 12.1, pp. 23-62.

**Coppock**, Elizabeth (2014). "Review of Paul Elbourne, Definite Descriptions". Nordic Journal of Linguistics 37.1, pp. 112–120.

Elbourne, Paul (2005). Situations and Individuals. Cambridge, MA: MIT Press.

— (2013). Definite Descriptions. Oxford: Oxford University Press.

**Heim**, Irene (1982). The Semantics of Definite and Indefinite Noun Phrases. (2011 Edition). — (1990). "E-Type Pronouns and Donkey Anaphora". Linguistics and Philosophy 13.2, pp. 137–177.

— (2002). "On the Projection Problem for Presuppositions". Formal Semantics. The Essential Readings. Ed. by Paul Portner and Barbara H. Partee. Oxford: Blackwell, pp. 249–260.

**Kratzer**, Angelika (1989). "An Investigation of the Lumps of Thought". Linguistics and Philosophy 12.5, pp. 607–653.

Mandelkern, Matthew and Daniel Rothschild (2020). "Definiteness Projection". Natural Language Semantics 28, pp. 77–109.

**Schlenker**, Philippe (2009). "Local Contexts". Semantics and Pragmatics 2.3, pp. 1–78. **Stalnaker**, Robert (1978). "Assertion". Syntax and Semantics 9, pp. 315–332.

<sup>&</sup>lt;sup>1</sup> The view that definite descriptions are evaluated at situations goes back to Barwise and Perry (1983); for a more recent, influential approach to situation semantics, which I follow here, see Kratzer (1989); for more recent situation semantics accounts of definite terms, see e.g. Heim (1990), Büring (2004), Elbourne (2005), Elbourne (2013). Pronouns may be treated similarly, as discussed, for example, by Elbourne.

<sup>&</sup>lt;sup>2</sup> Heim (1990), discussing the prospects for a situation-based account of definites, acknowledges that the predictions about how the presuppositions of definites project "ought to fall out from a general theory of presupposition projection" (p. 141), but then doesn't provide such a theory. Elbourne (2013, pp. 52–80) discusses presupposition projection, but I don't think his view amounts to a systematic account of projection; see Coppock (2014) for some critical discussion. See also Mandelkern and Rothschild (2020), who argue that situation semanticists haven't provided a systematic account of presupposition projection that covers the relevant data about the behaviour of these presuppositions.

<sup>&</sup>lt;sup>3</sup> Or, in dynamic semantics, sets of assignment functions or sets of pairs of an assignment function and a world; see e.g. Heim (1982).

<sup>&</sup>lt;sup>4</sup> See both dynamic semantic approaches such as Heim (2002) and pragmatic accounts such as Schlenker (2009).



## MENTAL IMAGERY AND CO-PREDICATION

MICHELLE LIU

University of Hertfordshire/Leverhulme Trust

Polysemy is a linguistic phenomenon where a word has multiple distinct but related meanings or senses. In this paper, I focus on one particular issue with respect to polysemy, i.e., co-predication. Co-predication is associated with a type of polysemy known as 'logical polysemy'. In the case of co-predication, a nominal logical polyseme allows multiple predications selecting different senses of the polyseme. Consider the sentence: 'The book is covered with dust but still insightful'. The word 'book', which has two predications, is used in two distinct senses – it denotes 'book-as-tome' and 'book-as-content'. Drawing on recent research in psycholinguistics which suggests that language processing frequently involves mental imagery, I show how mental imagery can illuminate the phenomenon of co-predication. I address five objections against my account.



According to a received paradigm in speech act theory, speech acts are regulated by their illocutionary rules by conceptual necessity. Necessarily, any given illocutionary type (e.g. asserting, promising) is defined by the unique set of rules that governs its performance. (e.g. Pollock 1982; Williamson 2000; Alston 2000, chap. 8; García-Carpintero 2021). Following Levinson (Levinson 1983), I shall refer to this family of views as the "checklist theory" of illocutionary normativity, since it presupposes a 'checklist theory' (Fillmore 1975) of illocutionary concepts and their rules (specifically, one that stipulates a strict relation of ontological dependency between illocutionary types and their rules).

In the first part of the talk, I will present some objections to this received view. I will show that checklist theories have trouble accommodating (A) widespread disagreement between scholars about which cooperative rules govern each illocution, (B) competent speakers' inability to articulate illocutionary rules, and (C) the looseness of the normative constraints imposed by illocutionary rules.

The second part of the talk will attempt to sketch an alternative paradigm for understanding illocutionary normativity. Crucial to this alternative proposal is the idea that expecting others to cooperate in a joint project (like exchanging information) involves expecting others to (i) try to help achieve the established goals of the joint project, and (ii) act as rational agents. Focusing on assertive speech acts, I will show how these assumptions can be used to derive illocutionary norms from the distinctive features of different illocutionary acts.

The third part of the talk will review some of the implications of the alternative proposal. On checklist theories, illocutionary norms impose precise obligations on illocutionary performances. By contrast, the alternative paradigm understands illocutionary norms as loose, complex, and highly sensitive to features of the context. I will argue that this is a strength, rather than a weakness, of the alternative account ¬– an unsurprising claim, since rejecting an overly stiff conception of normativity is the only way to avoid objections (A-C). On top of this, the alternative paradigm is able to offer a plausible genealogical story about how illocutionary norms develop and come to be adopted within a linguistic community, while avoiding some controversial metaphysical commitments of the orthodox view.

## References

Alston, William P. 2000. Illocutionary Acts and Sentence Meaning. Cornell University Press.

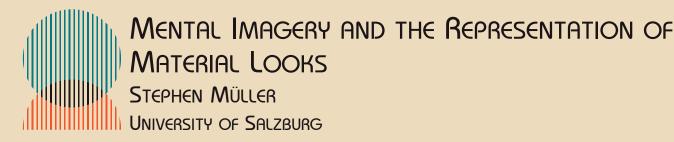
**Fillmore**, Charles J. 1975. 'An Alternative to Checklist Theories of Meaning'. Annual Meeting of the Berkeley Linguistics Society 1 (0): 123–31. https://doi.org/10.3765/bls.v1i0.2315.

**García-Carpintero**, Manuel. 2021. 'How to Understand Rule-Constituted Kinds'. Review of Philosophy and Psychology, July. https://doi.org/10.1007/s13164-021-00576-z.

Levinson, Stephen. 1983. Pragmatics. New York: Cambridge University Press.

Pollock, John L. 1982. Language and Thought. Princeton University Press.

**Williamson**, Timothy. 2000. Knowledge and Its Limits. Oxford University Press. <a href="https://doi.org/10.5040/9781474284974.00620">https://doi.org/10.5040/9781474284974.00620</a>



Some philosophers claim that mental images represent *how things look like* (e.g., Gregory, 2013, 2018). The aim of this talk is to present a novel and precise analysis of this proposal. In a first step, I will define the concept of material looks (short: looks<sub>M</sub>). While philosophers traditionally construe looks as experiential properties, material looks are properties that physical entities objectively have (roughly, they are a visual scene's surface reflectance properties set in relation to a potentially empty perspective). So, a scene has a look<sub>M</sub> irrespective of its being consciously perceived. Though rarely pointed out as such, related conceptions of looks feature in contemporary philosophy of language and mind. For example, they can (arguably) be identified in debates on perceptual justification (e.g., McGrath, 2018) and in the philosophical semantics of ordinary looks talk (e.g., Martin, 2010).

Then, having clarified what material looks are, I will argue that mental images represent them. To a first approximation, the presented theory claims that a mental image x represents a look<sub>M</sub> L by inheriting this content from a corresponding perception x\* of a scene that has L. To be more precise, we must first distinguish between types and tokens of representation. A token of representation is a neurobiological particular, while a type of representation is a set of functionally equivalent tokens. If a token perception x represents a particular scene y (as any mainstream externalist theory would have it), then according to the presented account this token's type T represents the look<sub>M</sub> L shared by the particular scenes that its members represent (amongst them y)1. Now, for every type of perception T, it is possible to individuate a corresponding type of mental imagery T\*. The pertinent notion of correspondence derives from the extensive body of empirical evidence showcasing that mental imagery and perception exhibit substantial overlap with respect to their neurobiological realization (e.g., Winlove et al., 2018; Dijkstra, Bosch & van Gerven, 2019). Tellingly, this similarity makes itself felt at the level of content (e.g., Pearson, Naselaris, Holmes & Kosslyn, 2015)<sup>2</sup>. Hence, it is expedient to say that a type of perception T corresponds to a type of mental imagery T\* iff their members are (approximately) neurobiologically equivalent. Finally, in order to determine the referent of any given mental image x, we must first identify its type T\*. If T\* corresponds to the type of perception T that represents L (in virtue of containing members that represent scenes exhibiting L), then T\* inherits this content L. Thus, we can say that the members of T\* (such as x) represent (non-actual) scenes that look L. It should be noted that in practice this procedure usually happens in reverse. That is, we speak of a mental image of L in order to individuate a mental image as belonging to a certain type (i.e., as structurally corresponding to perceptions of scenes looking L).

Generally speaking, the present theory is premised on the observation that the content of mental imagery tracks its cognitive function, and that (approximate) structural equivalence (between perception and mental imagery) implies (approximate) functional equivalence (see Pearson, Naselaris, Holmes & Kosslyn, 2015). By analyzing the content of mental imagery as described, we appreciate these connections. But the theory has further upshots, of which I will here describe only two. First, the present account helps us articulate in which sense the contents of perception and mental imagery are similar (similar content view),

but also in which sense they are different. Usually, a similar content view is proposed to explain the salient phenomenological likeness that relates perception and mental imagery (e.g., Kind, 2001; Nanay, 2015). In the theory at hand, similar content results naturally from only functional and neurobiological considerations. Thus, it provides independent evidence for similar content views while sidestepping any need for introspection. Second, the theory helps us make sense of the unconscious representation of looks. If looks are construed as experiential properties, it is not clear how they could be unconsciously represented (though Gregory 2013, 2018 makes such a proposal). Yet, in the service of theoretical unification, unconscious mental imagery is enjoying increasing acceptance within cognitive science (e.g., Brogaard & Gatzia, 2017; Nanay, 2021). So, by shifting focus to material looks, the present theory neatly accommodates for this development. For this and the aforementioned reasons, the representation of material looks should be especially interesting as applied to cognitive theories of mental imagery.

To conclude, I will discuss apparent counterexamples to the given account and argue that they pose no threat. There are diverse intuitions as to what mental images "ought to" represent, some of which seem in conflict with looks<sub>M</sub> content<sup>3</sup>. For instance, some may feel that mental images can represent particular objects (such as the Eiffel Tower) or even non-existent particular objects (such as the tooth fairy). Then, some may claim that mental images can represent generic content (such as a non-particular apple), or things that are unperceivable even in principle (such as electrons or even impossible geometrical configurations). Instead, so my claim, mental images can merely represent the looks<sub>M</sub> of scenes that prominently feature such objects. As for unperceivable objects, we cannot visualize them whatsoever, but merely related objects that do prominently feature in visual scenes (such as a sketch of the Bohr model or the Penrose triangle). So, given its methodological expediency, we should forfeit these intuitions in favor of the presented account.

## Rerefernces

**Brogaard, B., Gatzia, D.E.** (2017). Unconscious Imagination and the Mental Imagery Debate. Front. Psychol., 8, 799. doi: 10.3389./fpsyg.2017.00799

Dijkstra, N., Bosch, S.E., van Gerven, M.A.J. (2019). Shared neural mechanisms of visual perception and imagery. Trends in Cognitive Science, 23, 423-434.

**Gregory**, D. (2013). Showing, sensing and seeming: Distinctively sensory representations and their contents. Oxford: Oxford University Press.

**Gregory**, D. (2018). Visual Expectations and Visual Imagination. Philosophical Perspectives, 31(1), 187-206.

**Haugeland**, J. (1991). Representational Genera. In: W. Ramsey, S. Stich & D. Rumelhart (Eds.), Philosophy and Connectionist Theory. Hillsdale: Lawrence Erlbaum.

Horikawa, T., Tamaki, M., Miyawaki, Y., Kamitani, Y. (2013). Neural Decoding of Visual Imagery During Sleep. Science, 340, 639-642.

Kind, A. (2001). Putting the image back in imagination. Philos. Phenomenol. Res., 62, 85-109.

**Kulvicki**, J. (2003). Image Structure. The Journal of Aesthetics and Art Criticism, 61(4), 323-340. Kulvicki, J. (2006). Pictorial Representation. Philosophy Compass, 1(6), 535-546.

**Martin**, M.G.F. (2010). What's in a Look? In: B. Nanay (Ed.), Perceiving the World, pp. 160-225. Oxford: Oxford University Press.

**McGrath**, M. (2018). Looks and Perceptual Justification. Philosophy and Phenomenological Research, XCVI(1), 110-133.

Michelon, P., Koenig, O. (2002). On the relationship between visual imagery and visual perception:

Evidence from priming studies. European Journal of Cognitive Psychology, 14(2), 161-184. Nanay, B. (2015). Perceptual content and the content of mental imagery. Philosophical Studies, 172, 1723 – 1736.

**Nanay**, B. (2021). Unconscious mental imagery. Phil. Trans. R. Soc. B, 376: 20190689. doi: https://doi.org/10.1098/rstb.2019.0689

**Peacocke**, C. (1992). Scenarios, Concepts and Perception. In: T. Crane (Ed.), The Contents of Experience, pp. 105-135. Cambridge: MIT Press.

**Pearson**, J., Clifford, C.W.G., Tong, F. (2008). The functional impact of mental imagery on conscious perception. Current Biology, 18(13), 982-986.

**Pearson**, J., Naselaris, T., Holmes, E.A., Kosslyn, S.M. (2015). Mental Imagery: Functional Mechanisms and Clinical Applications. Trends in Cognitive Sciences, 19, 590-602.

**VanRullen**, R., Reddy, L. (2019). Reconstructing faces from fMRI patterns using deep generative neural networks. Communications Biology, 2(1), 1-10.

**Winlove**, C.I.P., Milton, F., Ranson, J., Fulford, J., MacKisack, M., Macpherson, F., Zeman, A. (2018). The neural correlates of visual imagery: A co-ordinate-based meta-analysis. Cortex, 105, 4-25.

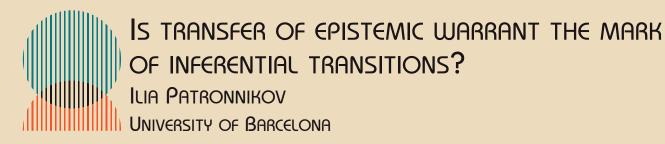
- $^{1}$  A perception's looks<sub>M</sub> content conceptually relates to a picture's barebones content (as analyzed in Haugeland 1991 and Kulvicki 2003, 2006), which consists of variations of incident light with respect to direction. Likewise, a scene's looks<sub>M</sub> determine the variations of incident light at a particular perspective.
- <sup>2</sup> Put differently, the content of mental imagery supervenes on its neurobiological structure. In particular, studies on perceptual priming (e.g., Michelon & Koenig, 2002), binocular rivalry paradigms (e.g., Pearson, Clifford & Tong, 2008) and especially neural decoding studies (e.g., Horikawa, Tamaki, Miyawaki & Kamitani, 2013; VanRullen & Reddy, 2019) illustrate this fact.
- $^{3}$  Note that many non-conceptual notions of content, such as Peacocke's (1992) scenario content, cohere very well with looks<sub>M</sub> content.



According to deflationist metaphysicians, the question "how many objects are there?" has an indeterminate answer because of different and equally available linguistic conventions, there not being a single convention better than others. My claim is that such philosophers are entitled to support vague existence for linguistic reasons.

My claim is challenged by a famous argument, first presented by Lewis, then revived by Sider and others, according to which, if the question assumes an unrestricted quantifier ranging over precise concepts, the answer has a definite answer corresponding to the wider domain the quantifier is allowed to range over and the quantifier is therefore not vague.

I argue instead that the extension of a quantifier is not its domain of quantification, but that the existential quantifier is a second-order function applied to first-order functions (using a Fregean assumption), and that any quantifier's extension is therefore constituted by concepts. Under this assumption, I will claim that quantifiers' vagueness depends on the equal availability of different concepts as extensions of the existential quantifier.



The paper focuses on an attempt to capture the notion of inference by an appeal to the idea that, unlike other kinds of transitions, inferential transitions transfer epistemic warrant. An account along these lines is defended by Susanna Siegel. In a nutshell, I argue that her account fails to adequately distinguish inference from other kinds of transition between mental states. I point out that the distinction between inferential and non-inferential transitions cuts across the distinction between transitions that transfer warrant and transitions that fail to do so. It is because, assuming a plausible picture of justification, transfer of epistemic warrant doesn't supervene only on the properties of the states featuring in a transition – instead, it supervenes on the properties of a larger system of states, of which the former states are part. If so, it is to be expected that some non-inferential transitions might transfer epistemic warrant, whereas some inferential transitions might not do it. For example, associative transitions can in principle transfer warrant under some circumstances, so on Siegel's account, such associative transitions would count as inferential. This conclusion is problematic because one of the things that the account aims to do is to capture the difference between inferential and associative transitions. Therefore, the account needs to be modified.



## Intentionalism out of control Carlotta Pavese, Alexandru Radulescu

CORNELL UNIVERSITY, UNIVERSITY OF MISSOURI

Suppose I say "That is my dog," and manage to refer to my dog, Fido. Intentionalism claims that my intention to refer to Fido is part of the explanation of the way that the demonstrative gets Fido as its referent. Extreme intentionalism claims that my intention completely determines the referent. Moderate intentionalism claims that my intentions do not suffice: for instance, that the audience needs to be able to figure out that I intended to refer to Fido. But if these conditions are met, the speaker's intention determines the referent; the speaker is in control of this semantic fact. In this paper, we argue that intentionalism must give up the claim that the speaker is always in control: we will provide new examples where reference happens but the speaker either lacks the intention to refer, or does not know how to refer, or yet again does not have the requisite knowledge for intentionally referring; in all of theses cases, we argue, the speaker is not in control of reference. We will argue that these cases provide better objections against intentionalism than the ones offered by proponents of conventionalism (such as Stojnic et al 2013; Stojnić 2021), and we go on to offer an improved formulation of intentionalism, which gives up claims of control, but retains the claims about the importance of the speaker's mental states in securing a referent for demonstratives.



In recent years, the theory known as the Free Energy Principle has gained increasing popularity in the philosophy of the cognitive sciences (see Friston (2009, 2010, 2012) for some earlier presentations of the theory, and Parr, Pezzulo, Friston (2022) for one of the most recent treatments). There is a number of declinations of this theory (often under different names, for instance "Active Inference", which I will use as an equivalent label), as it is not unexpected from cutting-edge scientific theories, but the core concepts of this theory remain stable: self-organising adaptive systems, and cognitive systems in particular, are to be conceived of as approximating optimal Bayesian inference, and in doing so they manage to remain attuned to their environment, thus staying alive. In particular, and this is the characteristic feature of the Free Energy Principle (FEP) that distinguishes it from earlier, more general Bayesian approaches to the brain and cognition while making it also a theory of living systems, adaptive systems achieve these accomplishments by minimising an information-theoretic quantity known as (variational/expected) free energy, which is used as a tractable proxy for another quantity, namely surprise. Surprise is in turn an information-theoretic notion closely related to entropy, and by keeping it to a minimum, living organisms, and cognitive systems in particular, manage to remain far from thermodynamical equilibrium while exchanging matter and energy with their environment. This, according to FEP, is done by performing Bayesian inferences so as to anticipate the sensory inputs coming from the environment, on the one hand, and to actively influence the environment itself through action, thus selectively sampling it.

In my paper, I will discuss a philosophically relevant question that can be, and has been raised (e.g. Colombo and Seriès (2012); Colombo, Elkin and Hartmann (2021); Van Es (2020)) concerning how to interpret the status of the statistical models employed within FEP, whether in a realist or an instrumentalist key. Specifically, I will elaborate a novel argument in favour of an instrumentalist interpretation of the use of statistical models in the context of this theory, and I will examine the implication that this interpretation has if one intends to elaborate an account of cognition based on FEP (as Kiverstein and Sims (2021) do). I will proceed as follows. First, I will introduce FEP, explaining as informally as possible the core ideas behind it, as well as the most important theoretical notions that it employs. Of particular relevance will be the notions of generative model (i.e. the statistical model involved in active inference) and of Markov blanket (i.e. the mathematical tool, comprising sensory and active states, which separates statistically the agent engaging in active inference from its environment).

Secondly, I will explain what adopting a realist approach to the statistical models involved in FEP would consist in, and I will present my case against this approach. The idea is that, if the models elaborated by scientists are meant to capture the statistical models that are "really" embodied by agents (a neutral term that I use to refer to systems meeting FEP or, equivalently, engaging in active inference), then the scientists' models should plausibly be interpreted as being isomorphic to the systems' embodied models. But isomorphism-based modelling (famously advocated by Van Fraassen (1980), Bueno (1997), French (2003), Bueno and French (2011), be it in terms of isomorphisms or partial isomorphisms) has been shown

to face important difficulties in avoiding that misrepresentation, i.e. inaccurate or incomplete representation of the target system, collapses into failure to represent the target system altogether (see Suárez (2003), Pero and Suárez (2016)). Hence, given that in setting up their statistical models scientists unavoidably have to make use of approximations, idealisations, and, more generally, design choices driven by their explanatory goals, FEP-scientists' models would systematically fail to represent the real models "embodied" or "instantiated" by their target systems, in virtue of the failure of the intended isomorphism to obtain.

This would obviously be quite problematic, and the solution that I recommend, in the third part of the paper, is adopting an instrumentalist, rather than realist, interpretation of the claim made by proponents of FEP that agents "embody" statistical systems. In this way, the distinction between the scientists' models and the systems' models would be, consciously and harmlessly, eliminated. The statistical models FEP makes use of would be straightforwardly conceived of as scientific constructs, which are meant not to capture the "real" models instantiated by agents engaging in active inference, but to simply describe the real statistical relations occurring within the target systems (in this sense, I concur with Van Es (2020) that the proposed instrumentalism is compatible with certain forms of realism, not concerning the statistical models but the entities and properties that those models are meant to capture). As a consequence, there would be room for misrepresentation in elaborating models of agents engaging in active inference (i.e. meeting FEP), and different models could be assessed in terms of their empirical adequacy, without losing their legitimacy as a consequence of their variable degree of accuracy or completeness.

However, instrumentalism is a double-edged sword. If, on the one hand, it protects the legitimacy of design choices in models of active inference scientists make use of, it has consequences over the usefulness of providing a "mark of the cognitive" (an account of what cognition is) in terms of FEP. To be sure, trying to elaborate one, as Kiverstein and Sims (2021) do, would still be a worthwhile enterprise, as it would further our understanding of the nature of cognition. But, unfortunately, it would make it impossible, in practice, to settle some of the debates from which the need for a mark of the cognitive stemmed (see Adams (2010, 2019), Adams and Aizawa (2001), Adams and Garrison (2013), Rowlands (2009). For instance, there would not be a fact of the matter as to whether there actually exist extended cognitive systems (Clark and Chalmers (1998) is the locus classicus of the literature on extended cognition; see also Kirchhoff and Kiverstein (2019)): whether a cognitive system is extended or not would depend on the initial design choices made by modellers, guided by their explanatory interests (similar points have been raised by Clark (2017)), so that one would need to assume from the start what one may be interested in finding out.



The aim of this paper is to provide a coherent and metaphysically plausible interpretation of fictionalist intuitions regarding the status of institutional groups (such as parliaments and corporations) and their alleged mentality. Such fictionalist intuitions are often expressed by non-philosophers, including legal theorists and theorists of the firm. In the debate in analytic social ontology, however, the fictionalist position is currently not very popular: philosophers tend to opt either for non-reductive realism or some form of reductionism.

The general question is how to interpret the statements, used both in ordinary talk and in legal language which suggest that we treat parliaments, corporation, courts etc, as possessing actual beliefs, desires and intentions. According to the strong realist view such statements can be, and often are, literally true; thus, corporations as as real believers as individual persons. Reductionist think that such statements might be reduced to statements about individual mental states and actions. Fictionalism opposes both reductionism and realism, and claims (in most general terms) that statements about institutions can be true only relative to a certain fiction.

In my view, the rejection of fictionalism about institutional minds by analytic social ontologists is premature. What is needed is a charitable reading of fictionalist intuitions. The way fictionalist intuitions about institutional mentality are presented by both legal theorists and theorists of the firm might seem inconsistent as these formulations seem to employ two mutually exclusive fictionalist approaches: the abstract artifact view and pretense theory. On the abstract artifact view, fictional entities are abstract creations of the relevant authors; in the specific case of parliaments and corporations, these entities are taken to be intentional creations within specific legal/institutional frameworks. On the pretense theory, on the other hand, our talk about fictional entities is a mere pretense which does not track any features of reality.

These approaches are, on the face of it, incompatible: abstract artifact theories entail that parliaments and corporations exist and statements about them might literally be true. Pretense theory, on the other hand, shies away from ontological commitments and does not treat fictional statements as true.

In my view, however, it is possible to provide a coherent interpretation of these pronouncements if we adopt a "two-stage" fictional approach to institutions and their mentality. On this two-stage approach, it is postulated that parliaments and corporations are primarily abstract artifacts which are created within legal/institutional frameworks. These abstract artifacts are subsequently used as props in the relevant games of make-believe, in which they are fictionally treated as agents that are capable of having intentions and performing actions.

On this view, corporations and parliaments exist, and some statements about them are literally true (e.g., "such-and-such corporation was incorporated on such-and-such date"). Many other statements about

them, however, are only made in pretense mode, especially those which ascribe actions and intentions to institutional groups. In such locutions, we only pretend to treat such abstract and artifactual entities as "persons" or "agents" that have "intentions", "beliefs" and such.

This two-stage approach might, in my view, capture certain important theoretical commitments of certain strands in both legal theory and the theory of the firm, namely that firms and parliaments are indeed real, but they lack important features of concrete entities and genuine intentional agents.



In his famous paper "Language as Thought and as Communication", Sellars argued (perhaps implicitly against Grice) that language should primarily be viewed as a form of expression of conceptual thought and not as a vehicle for communication (Sellars 1969). However, he also suggested understanding expression of thought and thought itself in a way which makes language prior to thought.

This view presents a package of answers to two separate questions. First, whether we can understand language and linguistic meaning in non-communicative terms (Sellars) or whether they're somehow inherently communicative or audience-directed (Grice, Davidson, Dummett)? Second, whether language is prior to conceptual thought and can explain it (Sellars), whether conceptual thought is prior to language and can explain it (Grice, Lewis), or whether they're somehow on a par (Dummett, Davidson)?

Even though there are outliers, the most standard package answers are Sellars's own non- communicative + language-priority view and Grice's communicative + thought-priority view. The latter is taken for granted by many working in contemporary foundational semantics (Keiser 2020, Schiffer 2017). In this talk, I will explore the prospects of a Lockean or Lewisian view which steers a middle course by taking public language to be expression of thought and thus not inherently communicative, but on which at least some thought is taken to be prior to language so that it can get the latter off the ground. In a slogan form: Language has a life of its own, independently of communication. Thought has a life of its own, independently of language.

In the first part of the talk I will defend the idea that public language, in its current form, is autonomous and independent of communication.2 Even though language plausibly got started by our making noises in order to provide evidence to others of our rudimentary thoughts, over time the sound- thought pairs got conventionalized and language became independent of our particular communicative aims. This can be supported by at least three sorts of arguments. First, there are ways of using language which are very hard to press into the communicative mold despite the best Gricean attempts. One of the best examples is writing down your thoughts to organize them. Second, there can be mismatches between what we say or otherwise linguistically convey (Austinian locutionary acts) and what we communicate which have led many to insist on a strict semantics/pragmatics distinction (Austin 1962). Finally, there are ways of using language deceptively which are very hard to make sense of on communicative views (e. g. see the parallel with cheating in games in Reiland 2020, and discussion of deceptive uses of names and demonstratives in Michaelson 2021).

In the second part of the talk I will discuss how to best understand the view that some thought is independent of language. Since Wittgenstein, this has been vilified as the Augustinian Picture (Wittgenstein 1953). Some traditions in philosophy of language, especially those influenced by Wittgenstein or Sellars, seem to take it for granted that this just can't be the case. Other, more mainstream

traditions seem to take it as obvious that it is the case. And there is very little engagement between them. I will argue that the resistance to the Augustinian Picture by Wittgenstein, Sellars and others stems from their assumption that the capacity to represent things as falling under general categories, to predicate, to operate with concepts, must be acquired (Hicks 2021). However, many contemporary representationalists think instead that basic representational abilities, including some propositional-logical abilities, are primitive and innate, part of the cognitive endowment of higher animals (Burge 2010). From this point of view a limited form of Augustinian picture is entirely defensible when it comes to logical forms (predication, disjunction) and basic concepts (e. g. BODY, CAUSE, ANIMACY).

However, this only gets us as far as some basic forms of thought. Still, one might object that most of our sophisticated concepts are dependent on language and social enculturation, something made vivid by Putnam with his notion of linguistic division of labor and Burge with his notion of anti- individualism (Burge 1979, Putnam 1975). I will argue that this incontrovertible point is not in tension with the priority of thought over language in the sense which is needed to get language off the ground. Some thought is prior to language and can get it off the ground, but mature conceptual thought is also dependent on language, though not on a particular natural language like English or German.

The end result is a view which respects the autonomy of public language together with a strict semantics/ pragmatics distinction while being able to explain how language gets off the ground in terms of expression of independent thought.

## References

Appiah, A. 1985. Assertion and Conditionals. Cambridge: Cambridge University Press

Austin, J. L. 1962. How To Do Things With Words. Cambridge: Harvard University Press

**Burge**, T. 1979/2003. "Individualism and the Mental". In Foundations of Mind. Oxford: Oxford University Press, pp. 100-150

**Burge**, T. 2010. "Steps Towards the Origins of Propositional Thought". Disputatio, 4, pp. 39-67 Davidson, D. 1984. "Communication and Convention". Synthese, 59, pp. 3-17

**Hicks**, M. 2020. "Sellars, Price, and the Myth of the Given". Journal for the History of Analytic Philosophy, 8, pp. 1-16

**Keiser**, J. 2020. "Language Without Information Exchange". Mind and Language. Online First. https://doi.org/10.1111/mila.12303

**Lewis**, D. 1975. "Languages and Language". Minnesota Studies in the Philosophy of Science. Ed. K. Gunderson. Minnesota: University of Minnesota Press, pp. 3-35

**Loar**, B. 1976/2017. "Two Theories of Meaning". Consciousness and Meaning: Selected Essays. Ed. K. Balog, S. Beardman. Oxford: Oxford University Press, pp. 29-47

**Michaelson**, E. 2021. "Speaker's Reference, Semantic Reference, Sneaky Reference". Mind and Language, Online First. https://doi.org/10.1111/mila.12349

**Putnam**, H. 1975/1996. "The Meaning of Meaning". Twin-Earth Chronicles. Ed. A. Pessin, S. Goldberg. Armonk: Sharpe, pp. 3-52

**Reiland**, I. 2020. "Constitutive Rules: Games, Language, and Assertion". Philosophy and Phenomenological Research, 100, pp. 136-159

**Schiffer**, S. 2017. "Intention and Convention in the Theory of Meaning". A Companion to Philosophy of Language, 2nd Ed. Ed. B. Hale, C. Wright, A. Miller. John Wiley & Sons, pp. 49-72

**Sellars**, W. 1969. "Language as Thought and as Communication". Philosophy and Phenomenological Research, 29, pp. 506-527

- <sup>1</sup> That Lewis held such a view, at least at some stage, is clear from the fact that he took linguistic conventions pertaining to uses of declarative sentences to be ones of truthfulness, understood in terms of expression of belief (see also Appiah 1985). This is frequently hidden from view by the fact that many Griceans appealing to Lewis's views of convention state such conventions not in terms of belief, but in terms of Gricean communicative intentions (Loar 1976, Keiser 2020, Schiffer 2017).
- <sup>2</sup> Davidson, despite subscribing to the thesis of 'autonomy of meaning', seems to reject this. What he insists on with his thesis is simply that literal meaning is always separated from ulterior purposes, hence that first meaning is separate from speaker meaning (Davidson 1984: 11-12). But on his view first meaning is essentially communicative.



## PROPOSITIONS AND NON-DOXASTIC ATTITUDE REPORTS

Wojciech Rostworowski, Katarzyna Kuś, Bartosz Mackiewicz University of Warsaw

The aim of the talk is to: (1) present novel evidence for the claim that propositions, or objects of attitudes, are structured entities, based on analysis of non-doxastic attitude reports (and including some experimental data), (2) propose a new semantic analysis of 'structured propositions' that serves as arguments for attitude-verbs operators, in order account for the indicated data related to non-doxastic attitudes.

It has been long recognized that ascriptions of attitudes pose a problem for a semantic theory, in particular, they prevent substitutions of embedded terms with the same denotations but different cognitive values, and fail to preserve logical entailments (e.g., Frege 1892, Cresswell & Bäuerle 1989, Heck 2012). Despite these difficulties, many accounts in philosophy and contemporary linguistics represent the contents on which attitude verbs operate – or at least define the truth conditions of attitude ascriptions – in terms of sets of possible worlds (e.g., Hintikka 1962, Stalnaker 1984, Lewis1986, Geurts 1999 sec. 5, Maier 2015, 2016). Of course, various sophisticated versions of this semantics are proposed in order to deal with specific problems (e.g., the presupposition-projection problem, see Heim 1992, Maier 2015), which delimit some undesirable consequences of the possible-worlds approach, including a move from 'possible worlds' to 'situations' (Barwise and Perry 1983, Elbourne 2005).

However, one challenge for all these approaches is a sensitivity of non-doxastic attitude verbs to the way how the content of an embedded clause is structured. [footnote 1: 1 Non-doxastic attitudes include the ones which do not imply believing in the proposition (like 'wanting something to be true' or 'wondering whether something is true') or the ones which, additionally to a belief, involve another component like emotional or evaluative (e.g., 'be glad that something is true').] This can be best illustrated by embedded conjunctions. In possible-worlds semantics, a conjunctional propositional 'p and q' can represented as the intersection of [p] and [q] (where [p], [q] are sets of worlds where, respectively, p is true, and q is true). Consequently, whenever q entails p, the conjunction 'p and q' is effectively semantically equivalent to its second conjunct (i.e., if [q] is a subset of [p], then the intersection of [p] and [q] equals [q]). So, for example, 'the dictator has been assassinated' is equivalent to 'the dictator is dead and has been assassinated'. However, such conjunctions generate different readings from the ones yielded by the 'sole' second conjuncts when embedded in non-doxastic attitude ascriptions, like the ones below:

(1)

- a. Anne wonders whether the dictator has been assassinated.  $\oplus$
- b. Anne wonders whether the dictator is dead and has been assassinated.

(2)

- a. John wants to lose the duel with honor. #
- b. John wants to lose the duel and to do it with honor.

(1b) attributes a more complex attitude to Anne than (1a); (1a) is true when Anne, let us say, believes that the dictator is dead (so does not wonder about this) but wishes to know whether the reason of death was an assassination; in such a case, (1b) is intuitively false. Based on this data, we argue that attitude verbs do not operate on conjunctional propositions understood as intersections of the sets of worlds representing the meanings of conjuncts. (For related discussions, see Blumberg 2017, Rostworowski 2018). Our argument is further supported by the results of experimental investigations, which have confirmed that (a)and (b)-ascriptions have different readings and that these readings are related to a genuine truth conditional difference, rather than a pragmatic effect. In particular, one of our studies tested the prediction that different attitude ascriptions with embedded conjunctions entail having the attitude towards both conjuncts. In the study, the participants were presented with two-sentence discourses and were asked to assess their coherence. The first sentence attributed an attitude to a subject, and it had the form like (1b) or (2b); the second sentence denied the attitude towards one of the conjuncts (e.g., 'In fact, he doesn't want to lose the duel.') The finding of the study was that such discourses were judged as much incoherent - similarly to discourses with 'canceled' semantic entailments (the ratings significantly below the midpoint, p < 0.001) - and much different from those involving a cancellation of a typical implicature (which were deemed to be coherent, although weakly; the ratings significantly above midpoint, p < 0.01). In order to account for the relevant data, we develop a structured-propositions analysis. Of course, such accounts have been already proposed (e.g., King 1995, von Stechow 1982). The novel idea in our approach is that the informational structure of a sentence is reflected in the way how an utterance of the sentence contributes to a discourse, and this contribution is shaped by various constraints of a pragmatic nature. One such constraint is 'informativeness' or Grice's Quantity, which implies that an utterance should not trigger a piece of information (globally or at 'local' levels; see Singh 2007) that is already established in the discourse. In light of this constraint, conjunctions contribute to a discourse differently from simplestructured sentences, even though the two may be truth-conditionally equivalent on an occasion. Consider:

(3)

- a. The dictator is dead. He [=the dictator] has been assassinated.
- b. The dictator is dead. #He [=the dictator] is dead and has been assassinated.

As (3b) repeats a piece of information that has been just established in discourse, it suffers from redundancy, in contrast to (3a), and so the conjunction makes an infelicitous contribution. Elaborating on our ideas in a more formal way, we will specify the informational structure expressed by a sentence in terms of the sentence's conversational contribution, and the conjunction will be specifically defined in terms of a sequence of updates on a potential context. We will consider how this analysis can be applied to other sentential forms and, finally, argue that once we represent the contents on which attitude verbs operate in terms of the specified informational structures, we get appropriately different readings for ascriptions like (1a) and (1b), or (2a) and (2b).

## References:

**Barwise, J. and Perry, J.,** 1983. Situations and Attitudes, Cambridge: MIT Press. **Blumberg**, K. 2017. 'Ignorance implicatures and non-doxastic attitude verbs', in A. Cremers, T. van Gessel, F.

Roelofsen (eds.) Proceedings of the 21st Amsterdam Colloquium.

Cresswell, M., Bäuerle, R., 1989. 'Propositional attitudes', in Handbook of Philosophical Logic (ed. F.

Guenthner and D.M. Gabbay) Vol IV, Dordrecht, Reidel, 491–512

Elbourne, P., 2005. Situations and individuals, Cambridge, MA: MIT Press.

**Frege**, G., 1892. 'Über Sinn und Bedeutung', Zeitschrift für Philosophie und philosophische Kritik, 100: 25–50.

Geurts, B., 1999. Presuppositions and Pronouns, Oxford: Elsevier.

Heck, R., 2012. 'Solving Frege's puzzle', Journal of Philosophy, 109: 132-174.

**Heim**, I. 1992. 'Presupposition projection and semantics of attitude verbs', Journal of Semantics, 9(3), 183–221.

**Hintikka**, J., 1969. 'Semantics for propositional attitudes', in J. Davis et al. (eds.), Philosophical Logic, Dordrecht: D. Reidel.

King, J. C., 1995. 'Structured Propositions and Complex Predicates', Noûs, 29(4): 516-535

Lewis, D., 1986. On the Plurality of Worlds, Oxford: Basil Blackwell.

Maier, E., 2015. 'Parasitic attitudes', Linguistics and Philosophy, 38(3), 205–236.

Maier, E., 2016. 'Attitudes and Mental Files in Discourse Representation Theory', Review of Philosophy and Psychology 7, 473–490.

**Rostworowski**, W., 2018. 'Descriptions and non-doxastic attitude ascriptions', Philosophical Studies 175(6), 1311-1331.

Singh, Raj. 2007. 'Assertability constraints and absurd assertions', unpublished manuscript, MIT.

Stalnaker, R., 1984. Inquiry, Cambridge, Mass.: MIT Press.

von Stechow, A., 1982. 'Structured propositions', Technical Report 59, Sonderforschungsbereich 99, Universität Konstanz.



# EXPERIENTIAL ACQUAINTANCE WITH ABSTRACT OBJECTS: TOWARDS A RECONCILIATION OF MAJOR VIEWS IN PHILOSOPHY OF PERCEPTION

IŞIK SARIHAN (INDEPENDENT POSTDOCTORAL RESEARCHER)

This paper presents a view concerning the epistemology of phenomenal experience and explores its attractive consequences. It is the view that all experiences involve an acquaintance relation with abstract entities, namely, possible states of affairs (i.e. the ways the world could be) and the property types that figure in these possible states of affairs. For instance, a visual experience of a green leaf (that is, veridically perceiving a green leaf, hallucinating a green leaf, visually imagining a green leaf, etc.) acquaints you with a possible state of affairs (the possibility of there being a green leaf in front of you), and it also acquaints you with a specific shade of the color green. This acquaintance relation is of a different type of epistemic relation to other epistemic relations found in an experience, such as perceptual awareness and introspective awareness, and therefore it does not interfere with our awareness of our environment — it does not constitute 'a veil of perception'.

Theorizing about abstract entities to solve various issues related to perceptual experience is not entirely new (see e.g. Bealer 1982; Butchvarov 1998; Pautz 2007) but the idea has not been developed to its full potential. I will argue that the view can help us to reconcile the main intuitions behind the intentionalist theory with the intuitions that motivate the other prominent theories in philosophy of perception. More particularly, I will claim that:

(Claim 1) The view allows us to satisfy the sense-data intuition that in hallucination we are in direct epistemic contact with various entities, but without postulating concretely existing objects (sense-data) or properties (mental paint, qualia) to explain this contact; and:

(Claim 2) The view also allows us to satisfy the naïve realist intuition that veridical experience puts us in direct epistemic contact with concrete objects in our environment (construed here as actualized possibilities) while not giving up the idea that veridical and non-veridical states have a shared content and phenomenology.

I will begin with clarifying some of the key terms and theoretical positions, and then present the reasons to believe that we are in a direct epistemic contact with abstract entities in perception, illusion, hallucination, imagination and dream. After countering some common misunderstandings regarding abstract entities and our epistemic relation to them, I will present a defense of the above two claims. I will conclude with a brief discussion of what to make of the Knowledge Argument (Jackson 1982) in light of the current view.

## References:

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

Butchvarov, P. (1998) Skepticism about the External World. Oxford University Press.

Pautz, A. (2007) Intentionalism and perceptual presence. Philosophical Perspectives 21(1): 495-541.

Jackson, F. (1982) Epiphenomenal qualia. Philosophical Quarterly 32: 127–136



# COMMON GROUND IN DEVELOPMENT ANTONIO SCARAFONE

CENTRAL EUROPEAN UNIVERSITY

The notion of common ground plays a key explanatory role in theorising about pragmatic reasoning. How can I trust that you will get what I mean when I say that John has already left the meeting? Many individuals go by the name 'John', and each of us might even know more than one, so what entitles me to use the name 'John' to refer to that John on this occasion? For us to communicate successfully, certain assumptions regarding a meeting being held and a person called John having attended it must be common ground. What holds for linguistic communication holds for non-linguistic communication too. How can I trust that, if I point to the door and I address my pointing gesture to you, you will understand that I want you to leave, rather than wanting you to appraise the colour of the newly painted door?

Over the last few decades, a vast number of experiments in developmental psychology has demonstrated that, early in their second year of life, and so before they master anything like a natural language, infants typically are able to communicate flexibly and effectively by means of pointing. Reliance on common ground is thus not a feature unique to the communicative exchanges of competent language users. The problem is that, despite the vast amount of evidence, it remains unclear how the very notion of common ground ought to be conceptualised, especially in accounting for pragmatic reasoning in infancy. This lacuna in the background theorising makes interpretations of the key experimental findings opaque.

In an important series of experiments, Henrike Moll and colleagues found that infants keep track of the common ground they have with others only in interaction, but not when they merely observe someone doing something in conditions which should equally generate mutual knowledge.<sup>1</sup> For instance, if they play with an object together with an adult, and later on the same adult points 'ambiguously' to a raw of three objects, infants preferentially hand the target object to the adult. If another adult does the ambiguous pointing, they are at chance. If they are invited to observe an adult playing with an object, in conditions which are supposed to equally generate mutual knowledge, they are again at chance in responding to the ambiguous pointing gesture. The available interpretations of these findings are, I will argue, either inconsistent or vague, appealing as they do to 'shared experiences', 'assumptions of rationality' or 'personal mutual knowledge'.<sup>2</sup> Furthermore, I will argue, there are independent reasons not to conceptualise common ground in terms of mutual knowledge for the case of infants, as the notion essentially relies on assumptions of normality which are arguably not available to the developing mind.

My proposed recipe to get out of trouble is in four-steps. First, I propose to conceptualise the notion of joint activity in the most general terms, namely, as the participation of multiple agents to the same event (Ludwig 2007). Second, I propose to conceptualise the very notion of common ground as a normative condition, namely, a set of shared commitments (Geurts 2019). The importance of conceiving of common ground as a normative condition becomes apparent when one thinks, first, that it must be possible to be wrong regarding what is in the common ground and, second, that there is no plausible account of common ground as a psychological construct. As a normative condition, the common ground that we have regulates the unfolding of our interactions, whether or not we have propositional knowledge of the commitments which we share. Key to understanding this possibility is, I will argue, an adequate conceptualisation of normative behaviours.

Third, I propose to conceptualise pointing gestures as devices to share commitments to attend to objects and events (Scarafone & Michael 2022). So conceived, communicating by means of pointing is a miniature joint activity which fosters the coordination of action in wider joint activities. This conceptualisation has several advantages. First, it provides a framework to understand how successful communication by means of pointing relies upon, and updates, the common ground. Second, it is a uniform account of the communicative function of pointing which is independent from its (varying) cognitive and motivational underpinnings.

Finally, I formulate a hypothesis regarding some of the representational proxies by means of which infants keep track of common ground. In brief, the hypothesis is that infants keep track of the adult's role in interaction. The notion of role is a normative one. I will argue that, coupled with the previous assumptions, this hypothesis provides a neat systematisation of the experimental findings. To explain how, later in development, behaviours observed outside direct interaction enter the common ground, I will argue that, in the eyes of children and with respect to unfamiliar adults, what changes is the scope of certain commitments or, in other words, the conditions which determine what counts as participating to the same event.

Common ground is key to pragmatic reasoning even in the prelinguistic case. Infants reliance on common ground can be adequately explained only if the very notions of common ground and pointing are conceptualised in normative, rather than psychological, terms, and only if this account is complemented by an account of infants' recognition of adults' role in interaction. This shift in perspective promises an important advancement in our understanding of the nature of prelinguistic communication, and, I will tentatively suggest, may open up new possibilities for understanding how pragmatic reasoning displayed by prelinguistic communicators may foster the acquisition of a language.

## References

**BOHN, M. & KÖYMEN, B**.2018 Common ground and development, «Child Development Perspectives», 12 (2): 104- 106.

**GEURTS**, B. 2019 Communication as commitment sharing: Speech act, implicatures, common ground, «Theoretical Linguistics», 45 (1-2): 1-30.

**LUDWIG**, K. 2007 Collective intentional behavior from the standpoint of semantics, «Nous», 41 (3): 355-393.

MOLL, H., CARPENTER, M., TOMASELLO, M. 2007 Fourteen-months-old know what others experience only in joint engagement, «Developmental Science», 10 (6): 826-35.

MOLL, H., KORING M., CARPENTER, M., TOMASELLO, M. 2006 Infants determine others' focus of attention by pragmatics and exclusion, «Journal of Cognition and Development», 7(3): 411-30.

MOLL, H., RICHTER, N., CARPENTER, M., TOMASELLO, M. 2008 Fourteen-month-olds know what "we" have shared in a special way, «Infancy», 13 (1): 90-101.

**MOLL, H., TOMASELLO, M.** 2007 How 14- and 18-month-olds know what others have experienced, «Developmental Psychology», 43(2): 309-317.

**SCARAFONE, A., MICHAEL, J.** 2022 Getting Ready to Share Commitments, «Philosophical Topics», 50 (1): forthcoming.

**TOMASELLO, M., HABERL, K.** 2003 Understanding attention: 12- and 18-month-olds know what's new for other persons, «Developmental Psychology», 39: 906-12.

<sup>1</sup> See, e.g., Moll et al. 2006, 2007, 2008, Moll and Tomasello 2007, Tomasello and Haberl 2003.

<sup>&</sup>lt;sup>2</sup> See, e.g., Bohn & Köymen 2018.



I argue that the phenomenon of breaking the fourth wall poses a novel puzzle for semanticists of fiction concerning (meta)fictional reference. I suggest that a synthesis of two analyses of fic- tional names can adequately deal with this issue.

A new puzzle of metafictional reference. In the movie On Her Majesty's Secret Service, James Bond is left alone on a beach holding a woman's slippers as she runs away from him. Bond says:

## (1) This never happened to the other fellow

and smiles into the camera. What "other fellow" does Bond refer to? Does he mean Sean Con- nery, i.e., the actor that previously played Bond? How can Bond possibly make reference to Connery? Fictional names as they appear in 'fictional statements' (i.e., statements that are part of a fiction) are often analysed as pretend-referring to actual objects. Consider Tolkien's:

## (2) Frodo had a very trying time that afternoon

The name 'Frodo' in (2) involves pretend-reference to a flesh and blood hobbit and (2) is pretend-true. On the other hand, fictional names in 'metafictional statements' (i.e., assertions about fic-tional entities as fictional entities) are often analysed as referring to abstract objects. Consider:

## (3) Frodo was invented by Tolkien

'Frodo' in (3) refers to an actually existing abstract object invented by Tolkien and (3) is true. As recently described by Recanati (2018), so-called 'parafictional statements' (i.e., state- ments about what is true in/ according to some fiction) complicate this picture by raising issues concerning mixed internal/external perspectives. Consider: (4) In The Hobbit, Bilbo lives in the Shire. Here we seem to have an (external) assertion; (4) is really true, not just pretend-true or true in the fiction. However, (4) seems to involve (internal) pretend-reference to a flesh and blood hobbit rather than reference to an existing abstract object invented by Tolkien; Surely, an abstract object cannot "live in the Shire".

I suggest that a mirror case of parafictional discourse is the phenomenon of 'breaking the fourth wall' (BtFW), i.e., when a fictional character somehow acknowledges the fictionality of the world they are in. For instance, prima facie, (1) is a fictional (internal) statement that in- volves metafictional (external) reference (i.e., reference to a fictional entity as a fictional entity). Here "the other fellow" is taken to abbreviate something like "the other actor that portrayed the fictional character Bond". Bond thus seems to break the boundary (or fourth wall) that exists between the fictional and the real world.

Preliminaries. It is important to distinguish BtFW from the broader category of 'metanarra- tion', i.e., self-reflective utterances that refer to the discourse rather than the story (Neumann and Nu"nning, 2014). BtFW only occurs when we cannot, by adding fictional truths (Cf. Matravers, 2014), accommodate fictional reference to fictional entities to make sense of the fictional utter- ance. So (self-)reference to a personified narrator that reflects on their own narration is analysed as 'plain' fictional reference. The same goes for (self-)reference to a film crew that actually exists in the fictional world as in mockumentaries. BtFW only occurs when a character is interpreted as referring to something outside of their world, e.g., an audience that cannot exist in their world.

Although BtFW occurs in novels (e.g., in Binet's The Seventh Function of Language a char- acter says: "I think I'm trapped in a fucking novel"), comic books (e.g., Deadpool) and video games (see van de Mosselaer, 2022), the phenomenon is mostly known from theatre and film. In order to do justice to all varieties of BtFW, we need an account of reference that can for in- stance analyse a meaningful look into a camera as a kind of pointing or demonstrative reference.

Analysis. Looking at BtFW as a mirror case of parafictional discourse reveals to what extent existing theories of fictional names can deal with this puzzle. On one analysis of parafictional statements (e.g., Evans, 1982) (4) involves pretend-reference to a flesh and blood hobbit. A well- known challenge for this analysis is the fact that statements such as (4) seem to be really true; 'Bilbo' (when taken to refer to an actual hobbit) is an empty name in the actual world and hence, on this analysis, (4) becomes a truth valueless assertion, or (4) becomes merely pretend-true.

A 'face value' approach would analyse (1) as a fictional statement that involves metafictional reference. This seems consistent with our common understanding of BtFW as characters acknowledging their fictionality. However, just as Bilbo does not exist in the actual world, there exists no actor who portrayed the fictional character Bond in the worlds of the Bond movies. In other words, "the other actor that portrayed the fictional character Bond" is an empty definite description in the Bond-worlds. Hence, in the fiction, (1) just becomes a truth valueless (non- sense) statement made by Bond. This does not account for our intuition that what is expressed (i.e., that Connery never portrayed a Bond that women ran away from) is (also) really true.

According to an alternative analysis of parafictional discourse, statements such as (4) involve real reference to abstract objects (e.g., Zalta (1988); Semeijn and Zalta (2021); van Inwagen (1977)). Such an analysis fits a 'strict' approach where fictional pretend-reference is reserved for fictional statements and metafictional reference is reserved for non-fictional statements.

A 'strict' approach would deem (1) a non-fictional statement (assuming it involves metafictional reference). In other words, BtFW would be analysed as a special case of what Gendler (2000) calls 'pop-out', i.e., an assertion by the author that 'interrupts' the fictional discourse. Who is 'the author' in the case of (1)? Maybe George Lazenby, the actor who portrayed Bond in On Her Majesty's Secret Service. Thus Lazenby makes the (true) metafictional assertion that Connery never portrayed a Bond that women ran away from. However, this analysis also doesn't do complete justice to our interpretative intuitions. Isn't it true in the fiction that Bond says something funny/weird, i.e., that Bond says (1)?

I suggest that a synthesis of these two approaches captures the phenomenon best. BtFW statements such as (1) instantiate two speech acts simultaneously: (1) is a (in the fiction nonsensical and truth valueless) statement made by Bond (i.e., a fictional statement made by Lazenby) and (1) is a true assertion made by Lazenby. Although this analysis makes sense of the disruptive and confusing aesthetic experience often associated with BtFW, it is only a tentative answer to the novel and intriguing puzzle of finding a semantic analysis of BtFW.

## References:

Evans, G. (1982) The Varieties of Reference, OUP

Gendler, T. S. (2000) 'The Puzzle of Imaginative Resistance', Journal of Philosophy

Matravers, D. (2014) Fiction and Narrative, OUP

**Neumann, B and Nu'nning,** A. (2014) 'Metanarration and metafiction', The Living Handbook of Narratology

Recanati, F. (2018) 'Fictional, metafictional, parafictional', Proceedings of the Aristotelian Society

Semeijn, M. and Zalta, E.N. (2021) 'Revisiting the 'wrong kind of object' problem', Organon F

Van de Mosselaer, N. (2022) 'Breaking the fourth wall in videogames', Being and Value in Technology

Van Inwagen, P. (1977) 'Creatures of Fiction', American Philosophical Quarterly

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



Generic sentences are used to convey exception-tolerant generalizations over a kind or, roughly, relevant members thereof (Carlson & Pelletier 1995; Cohen 2016). In English they often take the form of a bare plural (BP). The conditions under which generic sentences are acceptable is controversial, but common to many theories is an assumption that they involve some lawlike, universal/majority quantification (Dahl 1973; Pelletier & Asher 1997). However, it was acknowledged that generics receive an existential interpretation in certain environments (von Fintel 1997, Cohen 2004, Sterken, 2015): (i) When embedded as a propositional prejacent of a focus-sensitive particles (ii) existential generics usually judged to be unacceptable when presented out of the blue are considerably more felicitous when served as answers to a current Question Under Discussion (QUD) (Roberts 2012).

Such data pose challenges to most existing accounts. I argue that by incorporating certain principles of Information Structure (Rooth 1992; Büring 2016; Vallduví 2016); and discourse pragmatics (Roberts 2012; Lewis 1979; Beaver and Clark 2008) the semantics of BP generics can be reduced to a simple existential component. However, this requires that we abandon the Quasi-Universal Hypothesis. In this proposal, BP generics have a default weak existential semantics. In most instances where a generic is judged to be degraded, they are strictly true. I replace semantic constraints with felicity conditions that make reference to information structure in discourse. Following Cohen & Erteschik-Shir 2002, I propose that within a BP generic (i) the 'informationally marked' element answers to the immediate QUD. In the case of the predicate, it will be focused; in the case of the subject, it will be marked as a contrastive topic; (ii) the 'informationally unmarked' element should not answer the immediate QUD; (iii) Focus marking can be rendered more obvious through phonetic prominence, focus-sensitive particles, or contextual set-up; and (iv) The apparent quasi-universal, law-like reading of generics is a result of the interaction between BP generics' lexical presuppositions and pragmatic reasoning, and the interpretational variability of generics is partly explained by the open-ended nature of pragmatic implications in general.



## THE AFFECTIVE BASIS OF INFERENTIAL REASONING

## SANJA SRECKOVIC

BELGRADE UNIVERSITY

Animal cognition research offers empirical evidence of intelligent behavior of animals. In reaching their goals, animals often act as if they are capable of different forms of reasoning, such as logical (Call 2004, 2006; Marsh & MacDonald 2012; Erdőhegyi et al. 2007), causal (Blaisdell et al., 2006; Völter & Call, 2017; Schloegl & Fischer, 2017), and preferential reasoning (Moon & Zeigler, 1979; Vasconcelos, 2008), etc., in manners similar to human intelligent behavior. Intelligent behavior of animals is difficult to explain by the minimalist accounts of animal cognition. On the other hand, the accounts of human cognition crucially rely on assumptions concerning linguistic and meta-cognitive capabilities, and these assumptions are not readily extended to animals (Davidson, 1982; Stitch, 1979; Bermudez, 2003, 2006). We are thus left, first, with unexplained intelligent behaviors of animals, and, second, with similar behaviors in human and nonhuman animals, where only the former are explained, by accounts that are prevented by their core assumptions from extending to the latter. Furthermore, even if minimalist accounts could somehow rise to explain the intelligent animal behaviors, it seems that similar behaviors of different, but evolutionarily related, species should not be explained by substantially incompatible accounts. It is common to take the top-down approach, and propose accounts that involve sharing some of the more abstract capacities (e.g. concepts, beliefs) with non-human animals, to help explain their capabilities for inference without linguistic or propositional thought (Bermudez, 2003, 2006; Camp 2009, Beck 2012; Peacocke 2001; Schellenberg, 2013, 2018).

I take instead the bottom-up approach. I too believe that an adequate account should be consistent with the continuity of mechanisms underlying the (also continuous) intelligent behaviors of human and non-human animals. I propose an embodied approach to animal thought, which may be subsequently extended to human thought, preserving the continuity of mechanisms, and allowing for a more comprehensive account of cognition in general. Here I focus only on the first step of the project, which aims to explain the embodied basis of inferential capability in animals.

Traditional accounts of inferential capability admit solely the abstract elements, such as symbolic representations, and computations performed over these representations. I argue for a wider understanding of the elements of reasoning, involving bodily relations, which open the space for richer inference possibilities. I introduce affective constituents of inferential reasoning. Affects in general are characterized by valence – positive vs. negative, i.e., they can feel 'good' or 'bad'. They are also characterized by strength, or degree – they can be felt in a higher or lower degree. I argue that affects are not only the drivers for acquiring inferential capability, but that they also constitute meaning of at least some logical relations. For example, by being able to experience affects with opposite valences, cognitive agents experience contrariness, which is the basis of the relation of negation. They do not need to possess an abstract concept of negation or of contrariness, since they are able to experience it viscerally.

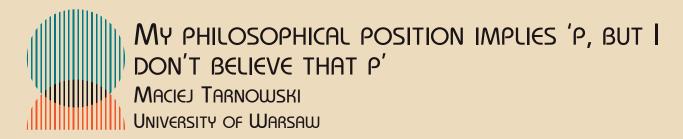
Likewise, when presented with more than one option, experiencing an affect of a particular valence more or less strongly in response to the respective options, enables the agents to experience comparison or

ordering of preferences following the different degrees within the same valence. The whole body is involved in choosing between options, and the bodily experience of choice between differently 'valued' options in turn gives rise to the meaning of exclusive disjunction. Again, the agent need not possess the abstract concept of disjuction, since the relation is experienced by their whole body.

The third example concerns causal reasoning. I argue that animals learn to reason causally not by acquiring an abstract concept of causation, but through their bodily experiences in the context of associated events. When a particular affect accompanies an event, the affective property tends to extend or be projected onto the preceeding event. When presented with the former event, the agent experiences the causal relation between the event as expectation of the consequent event. The expectation is not only contemplated intellectually, rather, the whole body becomes prepared for encountering it. This expectation is what makes correlation 'feel like' causation, and what is in the basis of the common fallacy. This experience, I believe, gives meaning to the relation between the two events, without the need for the agent to possess a concept of causation.

Inference is thus in a way performed by the whole organism, which alleviates the need for the reasoning subjects to possess concepts of logical relations. I will argue that this 'embodied inference' may provide a comprehensive explanation of various animal behaviors, but also shed light on the reasoning of prelinguistic children and aspects of adult human reasoning.

After presenting the main theoretical points of my account, I will show how it applies to particular cases of empirical evidence of intelligent animal behaviors. At the end I will briefly hint at how this account may be extended to humans as linguistic creatures capable of abstract thought.



George Edward Moore (1944) have famously noticed that the following types of assertions (called 'commisive' and 'ommisive' versions of Moore's Paradox [Sorensen 1988]) seem 'absurd' and 'self-contradictory', although the content expressed by them is consistent and may very well be true:

CMP: I went to the pictures last Thursday, but I believe I did not.

OMP: I went to the pictures last Thursday, but I don't believe I did.

In the recent literature, there has been however a growing amount of examples showing that the acceptance of certain philosophical or psychological positions implies that at least some Moorean assertions may be regarded as perfectly acceptable (for a survey see: Hájek 2007). As an example one may consider the following assertions made by a dialetheist (like Graham Priest) or an eliminative materialist (like Paul Churchland) respectively (see: Turri 2010, Williams 2015):

CMP: The Liar sentence is true, but I believe it is not (since the Liar sentence is both true and false).

OMP: I went to the pictures last Thursday, but I don't believe I did (since there are no beliefs).

Other examples may include cases of belief self-ascriptions based on the dispositionalist account of belief in cases of recognizing one's own implicit bias (Gertler 2021), cognitive dissonance (Borgoni 2015, Fileva and Brakel 2019) or acquired from a certain authority, e.g. a therapist (Pruss 2012). These examples are usually used to refute Moore's claim that commisive or ommisive Moorean assertions are always absurd and sometimes to challenge other philosophical positions (such as Evans' account of transparency of self-knowledge) which are used to justify Moore's claim.

In my paper, I will argue for two theses. First will be negative: from the methodological point of view, these observations should lead us to rather reject the philosophical position that allows Moorean assertions than to regard Moorean assertions as non-absurd. I will start by evaluating two arguments, one made by Turri and others:

- P1. Accepting philosophical position P leads to accepting (some) Moorean sentences.
- P2. Accepting P is rational.
- C. Accepting (some) Moorean sentences is rational.

And the other, which turns the structure of the above reasoning from modus ponens to modus tollens:

P1. Accepting philosophical position P leads to accepting (some) Moorean sentences.

P2. Accepting (all) Moorean sentences is absurd.

C. Accepting P leads to accepting absurdities (and therefore should be abandoned or modified).

I will argue that from a methodological standpoint the second premise P2 is better founded in the second argument. Firstly, I will point out that while the premise that accepting Moorean sentences is absurd is grounded in our linguistic and epistemic intuitions, the justification for the philosophical position P being rational ought to be founded in the analysis of the logical consequences of accepting P. By considering other cases of philosophical positions which needed to be modified or abandoned in the face of Moorean challenge (such as direct doxastic voluntarism [Williams 1970] or the claim that every true proposition can be believed), I conclude that if a philosophical position leads to the genuine assertion of CMP or OMP (has an absurd or self-contradictory logical consequence), then it should be abandoned or modified.

I will argue, however, that most of the cited examples – that is dispositionalist and eliminativist examples – are not genuine cases of Moorean assertions and therefore the alleged tension between Moore's observation and a variety of philosophical positions is unfounded. I will start by observing that the usual explanation of the Moorean phenomena involves distinguishing between the mere utterance of a sentence of the form "p, but I don't believe that p" and its assertion and between first-person and third-person ("p, but he/she doesn't believe that p") assertion, which demonstrates that the occurrence of Moore's Paradox is not necessarily intertwined with its syntactic form nor a proposition it expresses. I will argue that in fact, genuine Moore's Paradox arises only in cases where the assertion in question is an expression of one's de se belief, not merely its description. I will label such proper cases expressive, while the improper cases descriptive assertions, building on Dorit Bar-On's (2004) distinction between avowals as (expressive) acts and (descriptive) products.

The thesis that only expressive assertions of the OMP/CMP form are proper cases of Moore's Paradox will be supported by two observations. Firstly, I will notice that the belief ascriptions in improper cases do not enjoy first-person authority, which is usually taken to be a distinctive feature of mental state self-ascriptions (Bar-On 2004) since they can be properly challenged by a third party while in expressive cases they usually cannot. Secondly, I will argue that such belief ascriptions typically differ with respect to the transparency of reasons behind an ascription (which was famously noted by Evans [1982]). While in cases of expressive assertions of belief self-ascriptions the reasons provided by the subject for it may include reasons for the truth of belief's content, while in purely descriptive cases such reasons for self-ascription are ruled out. By noting that belief self-ascriptions in dispositionalist and eliminativist cases are not first-person authoritative and transparent, I will argue that they are much more alike to third-person Moorean belief ascriptions and cannot be regarded as proper examples of CMP or OMP. Such analysis of the absurdity of Moorean assertions is also in line with Evansian and expressivist solutions offered to Moore's Paradox (Moran 2001, Kwon 2019).

I will conclude my talk with a brief taxonomy of cited examples. Philosophical positions allowing for a non-expressive utterance of CMP or OMP, such as dispositionalism or eliminativism, are not threatened by Moorean challenge, but on the other hand, are not providing a basis for genuine Moore-paradoxical assertions. Other positions however, such as dialetheism or direct doxastic voluntarism, which commit subjects to the genuine threat of Moorean absurdity, should be regarded as endangered by this consequence.

### References:

**Fileva, I., & Brakel, L. A.** (2019). Just another article on Moore's paradox, but we don't believe that. Synthese, 196(12), 5153-5167.

**Bar-On**, D. (2004). Speaking my mind: Expression and self-knowledge. Oxford University Press on Demand.

Borgoni, C. (2015). Dissonance and Moorean propositions. Dialectica, 69(1), 107-127.

Evans, G. (1982). The Varieties of Reference, Oxford: Clarendon Press.

Gertler, B. (2021). Smithies on Self-Knowledge of Beliefs. Analysis, 81(4), 782-792.

**Hájek**, A. (2007). My philosophical position says 'p'and I don't believe 'p'. Moore's paradox: New essays on belief, rationality, and the first person, 217-31.

Kwon, H. (2019). Moore's Paradox: An Evansian Account. Erkenntnis, 1-17.

**Moore**, G.E. (1944). Russell's theory of descriptions. In: P. A. Schilpp (ed.), The Philosophy of Bertrand Russell, La Salle, Ill.: Open Court, 177–225.

**Moran**, R. (2001). Authority and Estrangement: An Essay on Self-Knowledge, Princeton NJ: Princeton University Press.

**Pruss**, A. R. (2012). Sincerely asserting what you do not believe. Australasian Journal of Philosophy, 90(3), 541-546. Sorensen, R. A. (1988). Blindspots. Oxford: Oxford University Press.

Turri, J. (2010). Refutation by elimination. Analysis, 70(1), 35-39.

Williams, J. N. (2015). Eliminativism, Dialetheism and Moore's Paradox. Theoria, 81(1), 27-47.



Experimental research in the philosophy of language utilizing setups modelled on Kripke's Gödel case has been taken to show that there exists cross-cultural, intra-cultural and intra-individual variation in reference assignments. At the same time, vignettes and probe questions in these setups have been taken to contain multiple sources of ambiguity by virtue of which the experimental research utilizing these setups might fail to track a semantically interesting phenomenon. Nevertheless, few exceptions aside, follow-up studies attempting to account for these ambiguities have been taken to replicate the original findings, leaving the impression that cross-cultural, intra-cultural and intra-individual variation in reference assignments of a sort that ought to be taken into account by philosophers of language is real.

I will argue for two things. First, I will argue that there is at least one other hitherto unaccounted for way of interpreting the task in setups modelled on Kripke's Gödel case that is such that test subjects' responses fail to be of interest to philosophers of language. On this way of interpreting the task, test subjects interpret the probe question as one that asks for a response based on their own linguistic competence with the term studied. Second, I will argue that because answers to probe questions understood this way fail to track a semantically interesting phenomenon, the only other previously identified way of interpreting the probe question, namely as one that asks for a response that involves test subjects' theorizing about reference, is equally problematic. On the assumption that there are no other ways of interpreting the probe question that delivers responses of interest, the upshot is that setups modelled on Kripke's Gödel case are not well suited for testing what they are supposed to test.

Experimental research utilizing setups modelled on Kripke's Gödel case typically aim to test whether non-philosophers' reference assignments are in accordance with descriptivism or rather with the causal-historical account of reference determination. Descriptivism, at least of the sort that Kripke aims to refute, is an account that predicts specific referents of token expressions as opposed to specific referents of expression types. That is to say, it predicts what the referent of a term is as used by an individual at a time, not what the referent of a term is as used by multiple different individuals who are a member of the same linguistic community. As a result, there can be situations in which descriptivism predicts different referents for different individuals of the same linguistic community and of what is, at least phonetically or orthographically speaking, the same term.

Now in setups modelled on Kripke's Gödel case, test subjects are explicitly told that the descriptions the subject in the vignette associates with the term whose reference is being tested are false of whatever it is that stands at the beginning of the communicative chain ending with the subject in the vignette's use of that term. Consequently, the test subjects and the subject in the vignette associate different descriptions with the term in question, and the term, as used by the test subjects, has, according to descriptivism, a different referent from the referent it has as used by the subject in the vignette.

In setups modelled on Kripke's Gödel case, the probe question is a question about the reference of a term as used by someone else. If test subjects base their response to the probe question on their own linguistic competence with the term studied, descriptivism predicts that test subjects' responses are based on the descriptions they themselves associate with the term whose reference is being tested. In that case, however, descriptivism and the causal-historical account predict the same response, since the test subjects believe that the description the subject in the vignette associates with the term are false of the causal-historically determined referent and therefore the descriptions the test subjects themselves associate with the term are true of the causal-historically determined referent. As a result, there is nothing in test subjects' responses from which conformity with one or the other theory of reference can be deduced.

On another way of interpreting the probe question, test subjects interpret the question as a theoretical one that asks for a response based on the theories of reference test subjects (perhaps implicitly) take to be true. This is how Genoveva Martí thinks the task is interpreted and she thinks this interpretation is likely because the question is about someone else's use of words. Further, this way of interpreting the probe question is problematic, according to her, because the test in question would merely show how test subjects think reference is determined. There is, however, little reason for thinking that test subjects have true theories of reference.

Wikforss and Machery et al. do not accept Martí's criticism to the effect that setups modelled on Kripke's Gödel case are problematic because they merely test what test subjects think the correct theory of reference determination is. Not recognizing that responses based on test subjects' own linguistic competence are equally problematic, their replies to Martí are not such as to be able to counter Martí's objection.

According to Wikforss, Martí is wrong in thinking that setups modelled on Kripke's Gödel case elicit responses based on theorizing about reference because we interpret other individuals' speech based on our own linguistic competence with the terms uttered. But that response just brings one back to the problem with linguistic competence-based replies.

According to Machery et al., setups modelled on Kripke's Gödel case might elicit responses based on test subjects' theorizing about reference, but this is not problematic because the reference assignments that can be deduced in this way correspond to the way in which test subjects use terms and it is use that is taken to be of interest. However, on the assumption that there are no other ways of interpreting the task that might result in responses that are of relevance, the relevant kind of use becomes linguistic competence-based use. But if responses based on theorizing about reference were to correspond to linguistic competence-based use, then the empirical data on which Machery et al. base their claim that use and responses based on theorizing about reference are congruent should show that test subjects' responses all converge. But that is exactly what the data fails to show.



According to the Intentional-inferential model, communication is a matter of expressing and inferring intentions that involve high-order meta-representations: e.g., the intention that the addressee recognizes that the speaker intends him to form a certain belief (Grice, 1957; Sperber & Wilson, 1995). The idea that understanding ordinary acts of communication requires entertaining such an intentional structure - also called the 'recursive mindreading claim' (Scott-Phillips, 2014) – has raised important questions concerning the psychological plausibility of this model.

This work focuses on the question of whether the Intentional-inferential model offers a suitable account of communication in the early stages of development. Several theorists have provided a negative answer to this question by pointing to the existence of a 'developmental dilemma': as young children display sophisticated communicative abilities long before they handle complex meta-representations of beliefs or intentions, the former cannot rely on the latter (e.g., Millikan, 1984; Breheny, 2006).

The first objective of this work is to provide a fine-grained theoretical analysis of the various versions of the 'developmental dilemma' put forth in the philosophical and psychological literature. This will allow us to disentangle distinct – although often conflated – objections to the Intentional-inferential model (following Moore, 2017). The "Belief objection" maintains that young children cannot attribute communicative intentions as this requires a full-fledged concept of belief, which in turn involves understanding that individuals may hold beliefs that misrepresent the actual states of affairs (Dennett, 1978). The scale of this problem goes along with what has long been considered a robust finding in developmental research, namely, that children younger than four lack the ability to attribute false beliefs (see Wellman et al., 2001). The "Complex inference objection" concerns the alleged complexity of the inference about speakers' goal-directed communicative behaviours: communicative intentions are second-order intentions, hence, the ability to attribute them may be difficult to acquire, both ontogenetically and phylogenetically (Bar-On 2013)

The second objective of this work is to assess the extent to which available empirical evidence in developmental psychology can speak in favour, or against, these two objections. In the last two decades, developmental psychologists have extensively investigated the early emergence of false belief understanding (for a review, see Scott & Baillargeon, 2017). While the debate about the interpretation of this data is far from being settled, a growing body of evidence suggests that rudiments of this capacity may emerge early in infancy (see Poulin-Dubois, 2020), thus paving the way for addressing the "Belief objection". In contrast with this flourishing literature, little if no empirical attention has been drawn to the "Complex inference objection". For this reason, the third objective of this work is to spell out the necessary steps to address this gap.

We will provide an in-depth conceptual analysis of the "Complex inference objection" moving from the following questions: Why are communicative intentions complex? Which features make them arguably difficult to be inferred? Geurts (2019) has recently claimed that such complexity depends on the self-referential or reflexive nature of communicative intentions, that is, the fact that the speaker intends the addressee to recognize the very intention with which she produced a communicative act. Here, we suggest that, despite its prima facie complexity, self-referentiality may not be as cognitively taxing as it is claimed to be. In developing our argument, we will integrate two complementary perspectives.

From a philosophical perspective, we will build upon early work from Bach (1987). His analysis of the reflexivity of communicative intentions suggests that their recognition amounts to identifying the communicative attitude expressed by the speaker. In line with this, we will discuss and compare self-referential communicative intentions with other types of second-order intentions (in communicative and non-communicative settings) and suggest that reflexivity can arguably reduce, rather than increase, the inferential complexity of communicative intentions.

From a developmental-cognitive perspective, following Csibra (2010), we will argue that, due to their peculiar character, communicative intentions can be expressed through specific behavioural cues (i.e., "ostensive cues"). This allows infants to decode communicative intentions before they can grasp their embedded content. It follows that processing communicative intentions can be procedurally distinguished from processing ordinary second-order mental states, thus opening up the possibility of observing distinct developmental trajectories in the emergence of these capacities.

Finally, we take stock of the two perspectives by outlining their impact on the "recursive mindreading claim" defended by proponents of the Intentional-inferential model and on the hypothesis of a mental system dedicated to the processing of communicative behaviours (Sperber & Wilson, 2002).

### References

Bach, K. (1987). On communicative intentions: a reply to Recanati. Mind & Language, 2: 141-154.

Bar-On, D. (2013). Origins of Meaning: Must we "Go Gricean?". Mind & Language, 28: 342-375.

Breheny, R. (2006). Communication and folk psychology. Mind & Language, 21: 74-107.

Csibra, G. (2010). Recognizing Communicative Intentions in Infancy. Mind & Language, 25 (2): 141-168.

Dennett. D. (1978). Beliefs about beliefs. Behavioural and Brain Sciences, 49: 568-570.

**Geurts**, B. (2019). What's wrong with Gricean pragmatics. Retrieved online: https://exlingsociety.com/images/Proceedings/ExLing\_2019/10\_0001\_000363.pdf

Grice, P. (1957). Meaning. The Philosophical Review, 66: 377-388.

**Moore**, R. (2017). Gricean communication and cognitive development. The Philosophical Quarterly, 67 (267): 303-326,

Millikan, R. (1984). Language, Thought, and Other Biological Categories. Cambridge, MA, MIT Press.

**Poulin-Dubois**, D. (2020). Theory of mind development: State of the science and future directions. In S. Hunnius & M. Meyer (eds). New Perspectives on Early Social-Cognitive Development. Elsevier: 141-166.

**Scott-Phillips,** T. (2014). Speaking our minds: why human communication is different, and how language evolved to make it special. Basingstoke: Palgrave-Macmillan.

**Scott, R. M., Baillargeon, R.** (2017). Early False-Belief Understanding. Trends in Cognitive Sciences, 21 (4): 237-249.

**Sperber**, D., Wilson, D. (1986/1995). Relevance. Communication and Cognition. Oxford, Blackwell. **Sperber**, D., Wilson, D. (2002). Pragmatics, Modularity and Mind-Reading. Mind & Language, 17: 3-23. **Wellman, H., Cross, D., Watson, J.** (2001). Meta-analysis of Theory of Mind development: the truth about false belief. Child Development, 72, 655-684.



The counterpart assumption is widely accepted in the study of slurs. It provides a simple explanation for the meaning of slurs. A neutral counterpart of a slur is a non-derogatory coextensional expression: The truth conditional contribution of 'Boche' is identical to the truth conditional contribution of 'German'.

The counterpart assumption yields an equally powerful explanation for the derogation associated with slurs. Slurs are especially derogatory because they derogate the target solely on the basis of group membership. This membership is specified with the neutral counterpart. Personal merits or demerits are irrelevant. Slurs may denigrate individuals, for example, on the basis of the skin colour or some other feature that associates them with demographic groups.

Lauren Ashwell questions the counterpart assumption with what I call a subset argument. She argues that gendered slurs do not have neutral counterparts. According to Ashwell, while 'slut' does indeed refer to women, it does not refer to all women. Rather, it refers to women who are promiscuous. 'Slut' then refers to a proper subset of women. As a consequence, Ashwell argues that we should re-evaluate the importance of the counterpart assumption.

In my talk, I defend the counterpart assumption. I develop a generic view of slurs by using an inferentialist framework originally proposed by Michael Dummett so that the modified rules contain a generic element. The proposed rules still capture Dummett's insight that the meaning of slurs is epistemically objectionable. For example, the meaning of 'Boche' allows to infer from the assertion "Hans is a Boche" that the Germans are typically cruel.

Crucially, the generic element accommodates the evidence that supports Ashwell's subset argument. I argue that gendered slurs are no different from ethnic and racial slurs. For example, statements like

Hans is German but he's not one of those cruel Boches.

are very common with ethnic and racial slurs and these should be explained. Indeed, the generic view explains these. The cruelty is attributed generically so it allows counterexamples. The generic view can then be extended to gendered slurs as Ashwell's subset argument relies exactly on this kind of evidence. As a consequence, the generic view yields a unified treatment of slurs across the board.

I conclude by noting that Ashwell's explanation of derogation is problematic. Since, in her view, the derogation does not stem from the counterpart assumption, she argues that the derogation associated with 'slut' stems from normative constraint on how women should behave. She goes on to extend this to ethnic and racial slurs. I am not convinced. Although yelling "You are not act- ing according to what is expected from black folks!" will most likely generate a heated discussion about how exactly black people

should behave, it is certainly not as derogatory as yelling the N- word. This suggests that the counterpart assumption provides a more convincing explanation for the derogation.



# EVALUATIVE DEFLATION, SOCIAL EXPECTATIONS AND THE ZONE OF MORAL INDIFFERENCE

Pascale Willemsen, Lucien Baumgartner, Bianca Cepollaro, Kevin Reuter

University of Zurich, University of Zurich, Università Vita-Salute San Raffaele, University of Zurich

Philosophers distinguish between two types of terms by which we can communicate moral content, namely thin and thick ethical terms (Eklund, 2013; Roberts, 2013; Väyrynen, 2019; Willemsen & Reuter, 2021. Thin ethical terms express a speaker's positive or negative attitude towards a person or their actions. Typical examples are 'good', 'bad', 'right', and 'wrong' – those terms most frequently used in deontology and utilitarianism and also recent moral psychological research (see Abend, 2013 for an overview). In addition, thick ethical terms also evaluate positively or negative, but they further communicate descriptive content. For instance, the two vices cruelty and dishonesty both communicate negative evaluations. However, being cruel differs from being dishonest in various descriptive ways. Being cruel is about inflicting physical or emotional harm, whereas being dishonest is about lying or cheating.

Philosophers aim to understand the nature of evaluative terms and by what linguistic means they express evaluative content. In doing so, they make two critical assumptions. First, they at least implicitly assume that whatever theory of evaluative language is true for negative terms will also be true for positive ones. They assume some kind of symmetry between positive and negative evaluative terms, such that whatever we learn about terms of one polarity will be the case for the opposite polarity alike. Second, the standard view assumes that positive terms express positive evaluations, viz. either praise itself or an essential preform of praise, and that negative terms express negative evaluations, namely blame or a pre-form of blame.

In a series of papers, both these assumptions have been challenged empirically (X & Y, 2020, 2021; Z, X, & Y, 2022 (names of authors omitted for anonymous review)). In an attempt to decide whether the evaluation of a thick term is semantically and pragmatically conveyed, X and Y (2020, 2021) employ the cancellability test for conversational implicatures (see, e.g. Grice, 1989, but see Blome-Tillmann, 2008; Sullivan, 2017; Zakkou, 2018 for discussions of the test's limitations). Here are two examples of the experimental stimuli they used:

- 1. Negative: Tom is rude [bad], but by that I am not saying something negative about Tom.
- 2. Positive: Tom is friendly [good], but by that I am not saying something positive about Tom.

In both cases, participants were asked whether the speaker contradicts herself. If symmetry holds, whatever contradiction ratings are found for negative statements should also be found for positive statements. However, negative evaluations were significantly harder to cancel compared to positive ones. Statements like (1) were judged to be significantly more contradictory than statements like (2). This asymmetry, called polarity effect, provides a challenge for the symmetry assumption and has not been predicted by metaethicists.

In this talk, we provide an explanation of this effect based on two pre-registered experiments involving over 1000 participants. We argue that the way we use morally evaluative terms to talk about other people's character or behaviour is systematically asymmetrical. Negative evaluative terms are used to assign blame for a person not meeting a moral standard of expectable moral decency. The assertive use of negative evaluative terms almost exclusively serves the purpose of negatively evaluating this moral transgression and assigning blame for doing so. In contrast, positive evaluative terms have two different uses: 1.) a proper evaluative use that is intended to assign praise and to indicate that our moral expectations have been exceeded, and 2.) an evaluatively deflated use that lacks an expression of praise (see Table 1). This evaluatively deflated use indicates that an action lies in the zone of moral indifference in which we remain unmoved and are mostly unwilling to praise or blame the agent.

In Study 1, we demonstrate that positive evaluative language can be used to praise a person but also in an evaluatively deflated way in which the speaker withholds praise. We use an adapted version of the cancellability test to determine whether the evaluation can be neutralised, i.e. whether the evaluation is defeasible. While only very few people consider the evaluation of a negative term defeasible, a significantly larger number of people does so for positive terms. In Study 2, we provide evidence that the defeasibility of positive terms is best explained by the speaker's social expectations as to how people should act. We demonstrate that participants do in fact consider a positive behaviour more expectable compared to negative behaviour. At the same time, however, participants indicate that, in general, they consider positive traits, such as honest or friendliness, something to approve of.



Propositional attitude verbs denote relations to propositions.<sup>1</sup> It's commonly thought that believe is one of these verbs. Based on belief-in ascriptions like John believes in sets, I'll develop a challenge to this view. The dispute. For propositional attituders, believe denotes a relation to a proposition. Their claim can come in several forms. On one view (Hintikka, 1969; Schiffer, 2003, among oth- ers), believe denotes a relation between an individual and a proposition. On another view (e.g. Forbes, 2018), believe denotes a relation between an individual, proposition, and state. Either way, an embedded that-clause (as in Eylem believes that it's sunny, for instance) provides an argument saturating believe's proposition slot.

Contra propositional attituders, I deny that believe denotes a relation to a proposition and propose that it denotes a monadic property of a state instead (compare Pietroski, 2005; Elliott, 2017). Given this, an embedded that-clause doesn't provide an argument saturating believe's proposition slot; its denotation composes with that of believe in a different way. How? As- suming that a that-clause denotes, as Moulton (2015) convincingly argues (see also Moltmann, 2013), a property, namely the property of having a certain propositional content, my proposed denotation composes with that of a that-clause via predicate modification, yielding another property, that of being a belief with a certain propositional content (see Elliott, 2017, pp.178-9, 183 for a detailed analysis).

The challenge. Believe combines with phrases that arguably don't denote propositions, as in John believes in sets. Embedding in sets under believe thus leaves the argument slot that propositional attituders posit unsaturated and so, absent something else that provides the required argument, leaves propositional attituders with a semantically ill-formed output. Yet as the acceptability of John believes in sets illustrates, this prediction is false. On my view, by contrast, no argument slot remains unsaturated, and the denotation of in sets can compose with that of believe via predicate modification.

In response to this challenge, propositional attituders may either 1) say that in XP-phrases indirectly provide propositional arguments via a type-shifter or 2) posit ambiguity/polysemy in believe, one reading being the propositional attituder's denotation, another responsible for the acceptability of believe in XP-constructions.

Response 1. Propositional attituders may posit a type-shifter that takes the denotation of in XP as input and returns a suitable proposition as output; which proposition is returned as output may be constant or vary depending on context. Since the output proposition is available as a propositional argument, my challenge is defused and any belief-in ascription is predicted to be equivalent to some belief-that ascription.

I argue that there's no satisfactory way of specifying the output proposition. This issue arises for both factual (1) and non-factual (2) uses of belief-in reports (for discussion of the distinction see Price, 1969, pp.426-455).<sup>2</sup>

- 1. John believes in sets.
- 2. Sue believes in her friends.

Problem for non-factual uses. 2 seems to require that Sue trusts her friends to do certain (contextually salient) things. But, following the literature on trust (Baier, 1986; Holton, 1994; Marušić, 2017), I argue that no belief-that which Sue could hold is sufficient for Sue to trust her friends to do these things. Thus, for a given non-factual use of a belief-in ascription, there's no equivalent belief-that ascription.

Problem for factual uses. Propositional attituders may suggest that a factual use of 1 is equivalent to John believes that sets exist and define their type-shifter accordingly. But fol- lowing Szabó (2003, p.591), suppose one is ontologically modest: one reasonably holds that one's ontology lacks certain things that exist and so reasonably believes that things one doesn't believe in exist. The equivalence propositional attituders suggest now predicts that one believes in things one doesn't believe in. That, however, isn't a reasonable attitude. Since one is being reasonable, the posited equivalence of factual uses of x believes in XP and x believes that XP exists fails. Developing Szabo's observation further, I argue that neither appeal to context nor loosening the output of the type-shifter successfully avoid this issue.

Response 2. Propositional attituders may posit ambiguity/polysemy in believe, one read- ing being the propositional attituder's denotation, another responsible for the acceptability of believe in XP-constructions. However, I argue that this posit is unmotivated.

Whereas many ambiguities/polysemies of English are lexicalized in other languages (com- pare know and the distinction between wissen and kennen in German), after consulting native speakers of several languages (among others, German, Italian, and Mandarin Chinese) I have yet to find a language that lexicalizes the ambiguity/polysemy posited by the propositional atti- tuder. Moreover, if we control for confounds (for instance, sudden shifts in topic and redundant presuppositions), believe fails the standard test for ambiguity, where one reduces two full state- ments into an elided conjunction and observes whether that conjunction is zeugmatic (Chomsky, 1957, p.36). Examples of acceptable elided conjunctions involving factual and non-factual uses of belief-in ascriptions include the following:

- 3.
- (a) Sue believes in the abundance of human error and that Frege was wrong about sets (in particular).
- (b) Sue believes that Frege was wrong about sets and in the abundance of human error (more generally).
- 4.
- (a) John believes in the teachings of his church and that Christ is his saviour.
- (b) John believes that Christ is his saviour and in the teachings of his church.

Admittedly, these findings leave open the possibility that believe is polysemous, in part because some polysemes allow for co-predication (compare polysemous bottle in David drank and smashed the bottle.).

But, I argue that there is no evidence for preferring polysemy over a claim of lexical generality, according to which both states of belief-in and belief-that fall in believe's extension, but believe doesn't have distinct belief-in and belief-that senses. To support this argument, I draw on an underappreciated diagnostic developed by Textor (2011) and Felappi (2019), which considers admissible corrections by pedantic speakers and highlights a crucial difference between believe and polysemes that allow for co-predication.

#### References

**Baier**, A. (1986). "Trust and Antitrust". In: Ethics 96.2, pp. 231–260. DOI: 10.1086/292745. Chomsky, N. (1957). Syntactic structures. The Hague: Mouton de Gruyter.

**Elliott**, P. D. (2017). "Explaining DPs vs. CPs without syntax". In: Proceedings of the Chicago Linguistic Society. Chicago, pp. 171–185.

**Felappi**, G. (2019). "The Face-Value Theory, Know-that, Know-wh and Know-how". In: Thought: A Journal of Philosophy 0.0. DOI: 10.1002/tht3.406.

**Forbes**, G. (2018). "Content and Theme in Attitude Ascriptions". In: Non-propositional Intentionality. Ed. by M. Montague and A. Grzankowski. Oxford: Oxford University Press.

**Hintikka**, J. (1969). "Semantics for Propositional Attitudes". In: Models for Modalities: Selected Essays. Ed. by J. Hintikka. Synthese Library. Dordrecht: Springer Netherlands, pp. 87–111. DOI: 10.1007/978-94-010-1711-4\_6.

**Holton**, R. (1994). "Deciding to trust, coming to believe". In: Australasian Journal of Philosophy 72.1, pp. 63–76. DOI: 10.1080/00048409412345881.

**Kriegel**, U. (2015). "How to Speak of Existence: A Brentanian Approach to (Linguistic and Mental) Ontological Commitment". In: Grazer Philosophische Studien 91.1, pp. 81–106. DOI: 10.1163/9789004302273\_005.

Marušić, B. (2017). "Trust, Reliance and the Participant Stance". In: Philosopher's Imprint 17.17.

**Moltmann**, F. (2013). Abstract Objects and the Semantics of Natural Language. Oxford: Oxford University Press. DOI: 10.1093/acprof:oso/9780199608744.001.0001.

**Moulton**, K. (2015). "CPs: Copies and Compositionality". In: Linguistic Inquiry 46.2, pp. 305–342. DOI: 10.1162/ LING\_a\_00183.

Pietroski, P. M. (2005). Events and Semantic Architecture. Oxford: Oxford University Press.

Price, H. (1969). Belief. Gifford Lectures. London: George Allen & Unwin Ltd.

Schiffer, S. R. (2003). The things we mean. Oxford: Oxford University Press.

**Szabó**, Z. G. (2003). "Believing in Things". In: Philosophy and Phenomenological Research 66.3, pp. 584–611. DOI: 10.1111/j.1933-1592.2003.tb00280.x.

**Textor**, M. (2011). "Knowing the Facts". In: Dialectica 65.1, pp. 75–86.

— (2013). "Seeing something and believing IN it". In: Perspectives on Perception. Ed. by M. M. McCabe and M. Textor. Berlin: Walter de Gruyter, pp. 65–86.

<sup>1</sup>I use 'proposition' to cover whatever the objects of belief-that are (sets of worlds, structured propositions, properties of individuals (centred/multi-centred worlds), etc.)

<sup>2</sup> See also Szabó (2003, p.585-6), Kriegel (2015) and Textor (2013).



### UNDERDETERMINISTIC COUNTERFACTUALS

TOM WYSOCKI

UNIVERSITY OF PITTSBURGH

One of the classic questions in philosophy of language is how to evaluate counterfactuals (Briggs, 2012; DeRose, 1999; Edgington, 2004; Gillies, 2007; Hiddleston, 2005; Lewis, 1973; Moss, 2012; Stalnaker, 1980). The following theory is concerned with one (though important) class of counterfactuals—causal counterfactuals—and offers a semantics based on underdeterministic causal models, an extension of deterministic causal models (Pearl, 2009; Spirtes, Glymour, & Scheines, 2001). This theory both accounts for standard cases found in the literature and has desirable formal properties.

Consider: Odysseus is planning how to cross the strait. Say, he found himself in a deterministic situation: both Scylla and Charybdis will attack the ship if the weather is clear, but neither will attack if it's misty. He has tools to deal with it: the deterministic interventionist framework (Pearl, 2009; Spirtes, Glymour, & Scheines, 2001). He'll evaluate the outcomes of either decision, and he'll wait for the mist. Say, he found himself in a probabilistic situation. The odds are one in ten that Scylla or Charybdis will attack in the ship if it's misty, and seven in ten that either will attack if the weather is clear. He has tools to deal with it: the probabilistic interventionist framework. He'll calculate the probability of crossing the strait in the mist and in clement weather. He'll wait for the mist. Say now, the monsters will attack if the weather is clear, but either may attack or stay put if it's misty. Obviously, he should wait for the mist, but there's no interventionist framework to help him reach this decision. Now is as good a time as any to formulate one: the underdeterministic framework.

The big picture is this. The strongest causal dependencies are deterministic; they hold if bringing about some events causally necessitate other events. The dependencies underlie the deterministic framework, which has allowed for deterministic theories of would-counterfactuals (Barbero & Sandu, 2021; Briggs, 2012; Galles, & Pearl, 1998; Halpern, 2000; Hiddleston, 2005; Hüber, 2013; Zhang 2013) and of causation (Beckers & Vennekens, 2018; Glymour & Wimberly, 2007; Hall, 2007; Halpern, 2016; Hitchcock, 2001; Weslake, 2015; Woodward, 2003), among others. Weaker are probabilistic dependencies; they hold if bringing about some events affects the probabilistic distribution of other events. The dependencies underlie the probabilistic framework, which has allowed for probabilistic theories of counterfactuals (Pearl, 2009; 2011), of causation (Fenton-Glynn, 2017; Twardy & Korb, 2011), and of causal discovery (Spirtes, Glymour, & Scheines, 2001), among others. Underdeterministic dependencies are the weakest kind of causal dependencies worthy of the name. Based on these dependencies, I introduce the underdeterministic framework and, partially as an illustration of what the framework can do, a theory of counterfactuals.

First, I motivate the framework; you can use it to model situations under uncertainty, which arise in science, ethics, political philosophy, and commonsense reasoning. Next, I introduce underdeterministic causal models; a single model can represent many ways in which the events can unfold—many possible histories. How the histories develop is governed by the model's structural equations, which, unlike in deterministic

models, can return multiple values. Multiple histories, in turn, afford the semantics of modal claims and counterfactuals. An event is causally possible (necessary) iff it happens on some (all) histories. A might-(would-) counterfactuals holds iff the consequent event is possible (necessary) after you bring about the antecedent event with an intervention.

The semantics works only if there's an account of interventions that can bring about any event, and hence I formulate such an account. Interventions come in three types: simple, synchronized, and semaphore. Simple interventions behave a lot like their deterministic counterparts. However, synchronized and semaphore interventions are new species, required to handle disjunctive antecedents. Subsequently, I extend the theory to handle counterfactuals with other counterfactuals in the consequent and any events in the antecedent.

Once the theory is complete, I discuss how it relates to other interventionist theories of counterfactuals. The underdeterministic theory judges the same counterfactuals as true as deterministic interventionist theories of would-counterfactuals with conjunctive antecedents (Galles & Pearl, 1998; Halpern, 2000). It also agrees with Briggs's (2012) theory of would-counterfactuals, which allows for disjunctive antecedents. However, the underdeterministic theory enjoys two advantages over Briggs's: unlike Briggs's, it easily handles variables with infinite ranges, and it uses a single model to evaluate any counterfactual with a disjunctive antecedent or with another counterfactual in the consequent. The main virtue of the underdeterministic theory of counterfactuals, however, is that, in contrast to its deterministic counterparts, it models genuine causal might-counterfactuals, and does it without resorting to probability distributions.

Finally, I apply the theory to the classic debate in philosophy of language on how to interpret might-counterfactuals. Some propose that to say "Odysseus may survive if he attempts to cross the strait" is to say "I don't know that Odysseus won't survive if he attempts to cross the strait," i.e., that all might-counterfactuals are would-counterfactuals that one can't rule out (Stalnaker 1980, DeRose 1999). On this proposal, an epistemic operator ("I don't know") binds an individual counterfactual ("Odysseus will survive if he attempts to cross the strait"). Using the underdeterministic theory of counterfactuals, I develop an alternative: the epistemic operator binds underdeterministic models rather than individual counterfactuals. The proposal explains cases that support the competing theories as well as new cases that the competitors can't handle.

### References

**Barbero, F., Sandu, G.** (2021). Team Semantics for Interventionist Counterfactuals: Observations vs. Interventions. Journal of Philosophical Logic 50(3):471-521.

**Beckers, S., Vennekens, J.** (2018). A principled approach to defining actual causation. Synthese 195:835–862.

Bennett, J. (2003). A philosophical guide to conditionals. Oxford:Oxford University Press.

Biswas, T. (1997). Decision-Making under Uncertainty. Cambridge:MIT Press.

Briggs, R. A. (2012). Interventionist Counterfactuals. Philosophical Studies, 160(1):139–166.

Cohen, M., Jaffray, J. (1983). Rational choice under complete ignorance. Econometrica 48:1281–1299.

**Deœux**, T. (2019). Decision-making with belief functions: A review. International Journal of Approximate Reasoning 109:87–110.

**DeRose**, K. (1999). Can It Be That It Would Have Been Even Though It Might Not Have Been? Philosophical Perspectives 13:385-413.

**Eberhardt, F., Scheines, R**. (2006). Interventions and Causal Inference. Philosophy of Science 74(5):981-995.

**Edgington**, D. (2004). Counterfactuals and the benefit of hindsight. In Dowe & Noordhof (eds.), Cause and Chance: Causation in an Indeterministic World. Routledge.

Ellis, B., Jackson, F., Pargetter, R. (1977). An objection to possible-world semantics for counterfactual logics. Journal of Philosophical Logic 6:355–357.

**Fenton-Glynn**, L. (2017). A Proposed Probabilistic Extension of the Halpern and Pearl Definition of 'Actual Cause.' The British Journal for the Philosophy of Science, 68(4):1061-1124.

Fine K. (2012). Counterfactuals without possible worlds. The Journal of Philosophy 109(3):221-246.

**Galles, D., Pearl J.** (1998). An Axiomatic Characterization of Causal Counterfactuals. Foundations of Science 1:151-182.

Gillies, A. (2007). Counterfactual Scorekeeping. Linguistics and Philosophy 30(3):329-360.

Glymour, C., Wimberly, F. (2007). Actual causes and thought experiments. In Campbell,

O'Rourke, Silverstein (eds.), Causation and Explanation, Cambridge, 43-67.

Hüber, F. (2013). Structural equations and beyond. Review of Symbolic Logic 6(4):709-732.

**Hall**, N. (2004). Two Concepts of Causation. In Collins, Hall (eds.), Counterfactuals and Causation, Cambridge, 225-276.

Hall, N. (2007). Structural equations and causation. Philosophical Studies 132(1):109–136.

Halpern, J. (2016). Actual Causality. Cambridge, MA: MIT Press.

Halpern, J. (2000). Axiomatizing Causal Reasoning. Journal of Artificial Intelligence Research 12:317-337.

Hiddleston, E. (2005). A causal theory of counterfactuals. Noûs, 39:632-657.

**Hitchcock**, C. (2001). The Intransitivity of Causation Revealed in Equations and Graphs. Journal of Philosophy, 98(6):273-299.

Hitchcock, C., Knobe, J. (2009). Cause and norm. Journal of Philosophy 106:587-612.

Knight, F. (1921). Risk, Uncertainty, and Profit. Boston: Houghton Mifflin Company.

Lewis D. (1973). Counterfactuals, Cambridge, MA: Harvard University Press.

Moss, S. (2012). On the Pragmatics of Counterfactuals. Noûs 46(3):561-586.

Norton, J. (2021). Eternal inflation: when probabilities fail. Synthese 198:3853–3875.

Norton, J. (in press). The Material Theory of Induction.

**Pearl** J. (2009). Causality: Models, Reasoning, and Inference, second edition, Cambridge: Cambridge University Press.

**Pearl** J. (2011). The algorithmization of counterfactuals. Annals of Mathematics and Artificial Intelligence 61(1):29–39.

Rawls J. (1971). A Theory of Justice, Cambridge: Harvard University Press.

**Runde** J. (1998). Clarifying Frank Knight's discussion of the meaning of risk and uncertainty. Cambridge Journal of Economics 22(5):539-546.

**Schaffer** J. (2004). Counterfactuals, causal independence and conceptual circularity. Analysis 64(4):299-308.

**Spirtes, P., Glymour, C., Scheines, R.** (2001). Causation, Prediction and Search, 2nd revised edition, Cambridge:M.I.T. Press.

Stalnaker, R. (1980). A defense of conditional excluded middle. In: Harper,

Stalnaker, & Pearce (eds.) IFS, vol. 15. Dordrecht: Springer.

**Szaniawski**, K. (1960). Some remarks concerning the criterion of rational decision making. Studia Logica 9(1):221–239.

**Twardy**, C., Korb, K. (2011). Actual Causation by Probabilistic Active Paths. Philosophy of Science 78:900–913. Vandenburgh, J. (ms). Causal models and the logic of counterfactuals.

**Verma, T., Pearl. J.** (1988). Causal Networks: Semantics and Expressiveness. Proceedings, 4th Workshop on Uncertainty in Artificial Intelligence, Minneapolis, MN 352-359.

**Walters**, L. (2009). Morgenbesser's Coin and Counterfactuals with True Components. Proceedings of the Aristotelian Society 109(1):365–379.

Weslake, B. (2015). A Partial Theory of Actual Causation. British Journal for the Philosophy of Science.

Williams, J. (2010). Defending Conditional Excluded Middle. Noûs, 44(4):650-668.

**Woodward**, J. (2003). Making Things Happen: A Theory of Causal Explanation, Oxford:Oxford University Press.

Zhang J. (2013). A Lewisian Logic of Causal Counterfactuals. Minds & Machines 23:77–93.