The logical form of psychological reports

David Lindeman

Georgetown University

I present a unified view of the form and interpretation of indirect, direct, mixed, and attitude reports which results from combining and subtly modifying Davidson's accounts of action sentences, indirect reports, and quotations. The view provides a solution to substitution problems from Frege and Prior while maintaining semantic innocence and respecting Davidson's view that semantic and inferentially relevant structure should correspond. Although arrived at by different lines of consideration, this view has resonances with views recently elaborated by linguists. This result should be of independent interest. It should also be of interest to scholars of Davidson inasmuch as it helps to clarify and vindicate Davidson's thoughts about logical form and the proper methodology for semantic theorizing.

Advocacy and the Function of Folk Psychology

Henry Schiller

Summary

We often use terms like 'wants' to advocate for our interests and for the interests of others. I explain how such advocating speech works, and in which contexts such uses of 'wants' are the default (and in which they are not). I use these observations to argue that there is an *irreducible normative function* to folk psychology.

Abstract

Folk psychology is our ordinary framework for understanding human behavior in terms of mental states, like belief and desire. Folk psychology is often thought of in terms of a proto-scientific *theory of mind* (Gopnik & Wellman 1994); this endows us with a particular way of understanding the behavior of complex physical systems. Correspondingly, it is often claimed that we acquired a folk psychological capacity in order to have dominion over our environment (an environment populated with intentional systems). This motivates the following reductive position:

(1) Folk-psychological reduction:
   a. We acquired a folk psychological capacity in order to predict, explain, and control the behavior of other individuals,
   b. All uses of folk psychology can be reduced to one of these basic functions.

Though the precise details of this claim have been disputed at the margins (cf: Knobe 2006), the view that we can explain all uses of folk psychology *in terms of its status as scientific theory* is widely endorsed.[1] When combined with another plausible claim about the *use* of psychological verbs—namely, that when we use

1 This includes not only proponents of the so-called 'theory theory' of folk psychology, but also those who expound simulation (Goldman 1989) and hybrid (Stich & Nichols 1992) views as well.

terms like ‘thinks’ and ‘wants’ we are tokening concepts from folk psychology—we get the position that all use of psychological language reduces to one of the three basic functions.

What I argue in this paper is that folk psychological verbs are part of a *normative practice*, and that their use cannot be reduced to things like prediction and explanation. Specifically I argue for the following view of *desire ascriptions*:

(2) ‘Wants’ as a verb of advocacy:
In cases of collective decision making or group deliberation, the default use of ‘wants’ is *to advocate* when the subject of the ascription is a member of the group (or is proxied by one).

I start with an observation about a typical use of verbs of desire—we often use such verbs to advocate our interests and the interests of others. Here is a characteristic example:

> *Sam and Mikhail are trying to decide where to get some food in New York. Their friend Jake will also join them. “I think I want pizza” says Sam; “I think Jake wants something that reminds him of the midwest” says Mikhail. “Let’s compromise and go to Speedy Romeo’s—they have St. Louis style pizza”.*

To advocate for something is to draw attention to reasons in favor of endorsing or doing that thing. To advocate someone’s interests is to put their interests on the table as potential sources of such reasons. Advocacy is connected to a notion of ‘doing for’—to do something for someone is to take their interests into account as one of your reasons.

After establishing that these uses of ‘wants’ occur, as well as their intended effect, I go into more detail about their *discursive function*. What I argue is that such uses establish a commitment to taking someone’s interest into account—the psychological effect of this is to draw attention to that thing as an object of (potential) value in our joint decision making. This is borne out by the observation that uses of desire ascriptions *in collaborative contexts* license the subsequent use of *deontic* modal verbs:

(3) *Context: We’re planning a camping trip with Dan;*
   a. A: Dan wants to make omelettes.
   b. B: So there should be eggs in the cooler.

Here, the natural reading of 'should' is performative and deontic (Kaufmann 2019). By contrast, however, when the subject of an ascription is not a collaborator, the same modal verbs tend to get an epistemic reading:

(4) *Context: We're planning on sabotaging Dan's camping trip;*
   a. A: Dan wants to make omelettes.
   b. B: So there should be eggs in the cooler.

When it comes to 'in-groups' and collaborators, the function of folk psychology is a kind of collaborative value management (cf: Tomasello 2022). It is only when we consider 'out-groups'—opponents and neutral parties—that we apply a scientific stance to psychologizing.

The fact that advocating uses of 'wants' cannot be reduced to prediction / explanation is apparent from a number of observations. For instance, we often act to satisfy the desires of the deceased (e.g. 'it's what he would have wanted') where no prediction or explanation is possible. I also consider the relationship between desire satisfaction and gratitude in collaborative contexts.

References

Dennett, Daniel Clement. 1981. *The intentional stance*. MIT Press.

Fodor, Jerry A. 1987. *Psychosemantics: The problem of meaning in the philosophy of mind*. MIT press.

Goldman, Alvin. 1989. Interpretation psychologized. *Mind & Language* 4. 161–185.

Gopnik, Alison & Henry M. Wellman. 1994. The theory theory. In Lawrence A. Hirschfeld & Susan A. Gelman (eds.), *Mapping the Mind: Domain Specificity in Cognition and Culture*, 257–293. Cambridge: Cambridge University Press.

Kaufmann, Magdalena. 2019. Fine-tuning natural language imperatives. *Journal of Logic and Computation* 29(3). 321–348.

Knobe, Joshua. 2006. The concept of intentional action: A case study in the uses of folk psychology. *Philosophical studies* 130. 203–231.

Stich, Stephen & Shaun Nichols. 1992. Folk psychology: Simulation or tacit theory? *Mind & Language* 7(1-2). 35–71.

Tomasello, Michael. 2022. *The evolution of agency*. Cambridge, MA: MIT Press.

How Do Indexical Beliefs Motivate Group Action?

Juliana Faccio Lima

January 7, 2025

Perry [1979] has famously argued that we need indexical beliefs to explain and rationalize intentional actions of individuals. I climb up a tree upon seeing a bear running towards me because I believe that "I am about to be attacked by a bear". The beliefs that "The inattentive hiker is about to be attacked by a bear" or "John is about to be attacked by a bear" are about me. But, if I don't believe that "I am the inattentive hike" or "I am John", indexical beliefs, I won't be motivated to act in a way to save myself.

Interestingly, we can generate similar cases with groups and conclude that a group indexical belief, a *we*-belief, is necessary to explain a group's behavior. Suppose I am hiking with a group called the 'Bear Food hiking team'. Upon seeing a bear coming towards us, we all climb up the nearest tree together. We are motivated to behave this way because we believe that "we are about to be attacked by a bear". The beliefs that "The inattentive hiking group is about to be attacked by a bear" or "Bear Food Hiking Team is about to be attacked by a bear" are also about us. However, unless we believe that "we are the inattentive hiking group" or that "we are Bear Food Hiking Team", we won't be motivated to act in a way to save ourselves.

A popular account of indexical beliefs, 'Self-Locating View', holds that indexical beliefs play such a fundamental role in explaining and rationalizing action because they reveal an object's location relative to the person who holds the belief. Thus, what is special about my belief that "I am about to be attacked by a bear" is that it locates me, the believer, in relation to the object of the belief. In other words, it is a belief that carries information about where, when, or who the object of the belief is in my immediate environment.

Given the similarities of the individual and group cases, it's reasonable to expect and *prima facie* desirable to have one account of indexical beliefs that explains why they are necessary to motivate individual and group actions.

In this paper, I will argue that the Self-Locating View can't accommodate group indexical beliefs for basically two reasons. First, in some cases, for me to have a *we*-belief, I do not need to know *where* the group the belief is about is in my immediate environment, as the Self-Locating View requires. For instance, if I'm leading the hike with the Bear Food team and I'm in front of the group, then I'm not seeing every member. As a result, I won't know where they all are in relation to me, but I can still have a *we-belief*. Second, I do not need to know *which group* the belief is about to have a *we*-belief – this is the equivalent of knowing *who* the belief is about in the individual case. Take the last case and further suppose that I don't know who some of the members of the Bear Food Hiking team are. This means that I do not know which group it is, in what I will argue is the relevant sense, but I can still have the *we*-belief.

The problems for the Self-Locating View become more interesting when we move to *demonstrative beliefs*, another type of indexical belief. Suppose I see that Nora is about to be shot and push her to the floor to save her. For reasons explained before, the motivating belief has to be indexical, in this case, a belief like "she is about to be shot". According to the Self-Locating View, this belief is indexical because it reveals Nora's location in relation to the person who holds the belief.

Now suppose that Andrew, a security guard, sees that Nora is about to be shot through security

cameras. He also sees that Jasmin, his co-worker, is near Nora but can't see her. He calls Jasmin and gives her precise directions to find Nora, which she successfully does. Together, Andrew and Jasmin saved Nora. The motivation was that they (as a group) believed that "she$_{[\text{Nora}]}$ is about to be shot". However, to be motivated to save Nora, Andrew didn't need to know where Nora was in relation to him or to both of them; just in relation to Jasmin. But Andrew's knowledge of where Nora is relative to Jasmin is not self-locating.

I will present an alternative approach to indexical beliefs where these are practical ways of thinking of objects. In my proposed view, indexical beliefs require a *know how* to interact with the object, instead of a knowledge of where, who, or when the object of the belief is in relation to the believer. This way of understanding indexical beliefs can uniformly explain both individual and group actions. In the *we*-belief case, I do not need to know where the group is relative to me, but just how to act as a group to perform a group action. In Andrew and Jasmin's case, Andrew needs to know how to get the group to get to Nora, which he does by guiding Jasmin to Nora. So, Perry is right that there's a special kind of belief to motivate actions. But, if I'm right, group action and indexical beliefs suggest that he is wrong about what an indexical belief is.

References

Anscombe, G. E. M. (1975). *The first person*. Clarendon Press.

Cappelen, H. and Dever, J. (2013). *The Inessential Indexical: On the Philosophical Insignificance of Perspective and the First Person*. Oxford University Press.

Castañeda, H.-N. (1968). On the logic of attributions of self-knowledge to others. *The Journal of Philosophy*, 65(15):439–456.

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

García-Carpintero, M. and Torre, S. (2016). *About Oneself: De Se Thought and Communication*. Oxford University Press.

Kaplan, D. (1989a). Afterthoughts. pages 54–57.

Kaplan, D. (1989b). Demonstratives. In Almog, J., Perry, J., and Wettstein, H., editors, *Demonstratives. Themes from Kaplan*, pages 481–563. Oxford: Oxford University Press.

Magidor, O. (2015). The myth of the de se. *Philosophical Perspectives*, 29(1):249–283.

Perry, J. (1979). The problem of the essential indexical. *Noûs*, pages 3–21.

Perry, J. (1990 [1977]). Frege on demonstratives. In Yourgrau, P., editor, *Demonstratives*, pages 50–70. Oxford University Press.

Perry, J. (1998a). Indexicals, contexts and unarticulated constituents. In *Proceedings of the 1995 CSLI-Amsterdam Logic, Language and Computation Conference*, pages 1–16. Stanford University: CSLI Publications.

Perry, J. (1998b). Myself and i. *Philosophie in synthetischer Absicht, Stuttgart*, pages 83–103.

Prosser, S. and Recanati, F. (2012). *Immunity to error through misidentification: New essays*. Cambridge University Press.

Shoemaker, S. (1968). Self-reference and self-awareness. *The Journal of Philosophy*, 65(19):555–567.

Shoemaker, S. (1970). Persons and their pasts. *American Philosophical Quarterly*, 7(4):269–285.

Boban Arsenijević (University of Graz) & Magdalena Kaufmann (University of Connecticut)

Trying our best

Conative predicates like English *try* (Try-verbs) tend to take infinitival or subjunctive complements for which they display a mix of extensional and intensional effects (Sharvit 2003). Syntactically, Try-verbs are argued to have small complements ('events', Ramchand and Svenonius 2014; Wurmbrand 2019), which reflect dependence in temporal (simultaneous) and in participant structure (obligatory control, i.e. identity of matrix and embedded subject). Syntactic and semantic deviations from this pattern are subject to cross-linguistic and dialectal variation and are typically attributed to coercion (Grano 2017). Building on Grano (2017), we argue that the investigation of a richer inventory of possible complements reveals further complexities in the semantic profile of Try-events and how they relate to the realization of different argument types in the syntax.

Known ingredients Grano (2017) provides a semantics for Try-predicates that captures its selection of control-clause arguments:

(1) a. TRY(P)(x)(e)(w) is defined only if $\forall y \forall e' \forall w' [P(y)(e')(w') \rightarrow \text{Ag}(e', y) \text{ in } w]$;
b. where defined, TRY(P)(x)(e)(w) = 1 iff $\text{Ag}(e, x) \wedge \forall w' \in INT_{x,w} : \exists e' [e' <_{init} P(x)(e')(w)]$

try-ascriptions are thus true of an individual if they are the agent of an event that is carried out with the intention that it develop into an event as described by complement predicate. The requirement that an actual event (physical or mental) is instantiated in the evaluation world is the ***action component*** (Sharvit 2003; Nadathur 2019), that it is intended as an initial stage of a *P*-event is the ***intention component***. Temporal modifications that specify *P* as having no initial stages at the time of evaluation are predicted to be infelicitous. Complements that do not semantically specify that the individual argument (subject) of *try* is the agent give rise to undefinedness. Both restrictions are avoided if a coercion operator is placed on the complement. The matrix subject is thereby guaranteed to be the agent of this enlarged complement; if CAUSE does not require temporal overlap, *P* can be specified to follow the causing event.

(2) $[\![\text{OP'}_C]\!] = \lambda P_{\langle e,\langle \epsilon,st\rangle\rangle} \lambda x \lambda e \lambda w . \exists e' [\text{CAUSE}(e, e') \wedge \text{Ag}(e, x) \wedge P(x)(e')(w)]$

Whether this operator is available and how CAUSE is specified captures patterns of crosslinguistic/dialectal variation.

Trying more data Upon closer inspection, even in the absence of a coercion operator, agentivity of the matrix subject in the complement is neither sufficient nor necessary for felicity. (3a) meets the agentivity presupposition but is infelicitous in dialects of English that do not allow coercion; non-agentive predicates as in (3b,c) are acceptable even in non-coercion dialects:

(3) a. Bill tried (*for himself) to read a book. b. Bill tried not to cry. c. Try to forget it.

(1) also predicts too strong a connection with intentions. Holguin and Lederman (2023) argue convincingly that (i) unlike *intend*, *try* does not entail that the agent believes that they can bring about *P*, (4), nor (ii) does it entail that the agent wants *P*, specifically, agents can try something to demonstrate that they cannot bring it about (cf. (4b)).

(4) a. I {try / #intend} to win the lottery.
b. (To demonstrate that the Cybertruck's windows cannot be broken,) Musk's assistant Franz { tried /#intended } to break the window.

The German Try-predicate *versuchen* can take finite complement clauses that describe an eventuality that the Try-event aims to bring about (***goal***); in this case, intention is entailed.

(5) Wir versuchen, dass sich die Gäste wie zu Hause fühlen.
we try that REFL the guests like at home feel
'We try to make the guests feel at home.' (#But we don't want this.) from context.reverso.net

Try can also surface with nominal arguments or as a noun:

(6) a. Try the book! b. Try reading a book! c. Give { books /reading books} a try!

In the constructions in (6), the argument does not name a course of events the agent is instructed to realize with the intention of finding out if they can. Instead, *reading* (or some other action involving the book, (6a)) is presented as something the agent should realize to find out whether it pleases them or helps them achieve some other contextually salient goal (***means***).

A new try We assume that TRY-events are characterized by having an AGENT, a THEME (an event predicate often specified by the *to*-infinitival), and a GOAL that the agent wants to bring about by carrying out the TRY-event. We assume that the TRY-event itself (***extensional trace***) is either an initial event in a continuation path towards the realization of the THEME-event (Sharvit (2003); Nadathur (2019)) or an instantiation of the THEME-predicate. Table 1 shows the alignments for a few of our examples (boldfaced: the role(s) realized overtly):

	Extensional trace	THEME	GOAL	MEANS
(3a)	initial subevent of a book-reading	to read a book	Bill has read a book	
(4b)	initial subevent of a window-breaking	**to break the window**	it is demonstrated that the window cannot be broken	initial subevent of a window-breaking
(5)	treating guests nicely	for the guests to feel at home	**the guests feel at home**	treating guests nicely
(6b)	book reading event	to read a book	you gain knowledge	**reading a book**

To provide a semantics that allows for flexibility and impact of complement type, we assume that the verb is lexically associated with an event predicate that can instantiate either the THEME (denoted by the control infinitival or resolved anaphorically; v the type of events) or the MEANS. We follow Sharvit (2003) and Grano (2017) for the action component and assume that TRY-events themselves are events in the world of evaluation (***extensional trace***) that have an intensional connection to the THEME. To allow for (4), we avoid an entailment to intention (pace Grano) or desire (pace Sharvit) for realization of the THEME. Instead, we assume that the Agent believes that, if the THEME can be realized (circumstantial possibility $\Diamond^{circ}$), then it is possible that an event of the same kind as the extensional trace brings it about ($e \approx_k e'$).

(7) $[\![\text{TRY}]\!] = \lambda P_{\langle s,\langle vt\rangle\rangle}\lambda x\lambda e.\exists P'_{\langle s,\langle vt\rangle\rangle}.[\text{AGENT}(e) = x \;\&\; \text{THEME}(e) = P' \;\&$
$[P(e) \vee [\; P = P' \;\&\; \Box^{dox(x)}[\Diamond^{circ}\exists e'[P'(e')] \rightarrow \Diamond^{circ}\exists e', e''[e' \approx_k e \;\&\; e' <_{init} e'' \& P'(e'')]]]]]$

A finite complement clause (German *dass*-clause, Serbian non-control *da* complement clauses, Todorović and Wurmbrand 2020; Kaufmann et al. 2023) does not fill the event predicate argument slot, but can specify the desired result state (GOAL) of the TRY-event. We assume that these complements contain modal operators that are anchored to the matrix verb event (Kratzer 2006); TRY-events can serve as the anchor for a prioritizing modal and the complement clause specifies something the agent is trying to bring about. In the presence of an overt theme argument (pronominal or infinitival), the goal can only surface as a rational clause (for syntactic reasons).

Conclusion. The investigation of different syntactic complement types uncovers a more complex event structure of TRY-predicates as has been assumed previously. The resulting unified semantics allows also for flexible integration of the complements as advocated for in the recent literature. Further research should systematically test the alignments we suggest as well as effects of volition and intention and their correlations with different complement types.

Sel. Ref. Grano 2017 Control, temporal orientation, and the cross-linguistic grammar of *trying*. • Holguin & Lederman 2023 Trying without fail. • Kratzer 2006 Decomposing attitude verbs. • Nadathur 2019 Causality, aspect, and modality in actuality inferences. Sharvit 2003 Trying to be progressive: the extensionality of *try*.

Belief-Intention alternation with Italian *convincere*

Achille Fusco[1], Tommaso Sgrizzi[1]
[1]IUSS Pavia

Introduction Crosslinguistically, it has been observed that certain clause embedding verbs exhibit a belief/intention semantic alternation wrt their complement clause (Grano 2019, 2024; Giannakidou & Mari 2022). Such alternation correlates with specific morpho-syntactic features of the complement clause, for example finiteness for English *persuade* (Grano 2019, 2024) and Russian *dumat* (Kasenov, 2023), indicative vs. subjunctive mood for Greek *pitho* 'persuade' (Giannakidou & Mari 2022). Interestingly, Italian *convincere* 'convince', shows the same alternation within the non-finite domain. The verb can select two infinitival prepositional complementizers:

(1) a. Marco ha convinto Gianni **di** avere un figlio.
Marco$_1$ has convinced Gianni$_2$ *di* PRO$_{1/2}$ have.INf a child.
Marco has convinced Gianni that he has a child.
b. Marco ha convinto Gianni **ad** avere un figlio.
Marco$_1$ has convinced Gianni$_2$ *a* PRO$_{*1/2}$ have.INf a child.
Marco has convinced Gianni to have a child.

In (1a), Mario causes Gianni to have a belief about some state of affairs *s*, whereas in (1b) the result is that Gianni has an intention to bring *s* about. We will argue that the two infinitival clauses differ in structural size (Wurmbrand 2001), and that this difference, paired with a suitable lexical semantics for the two complementizers, explains the semantic alternation observed.

Previous accounts of belief/intention predicates: Belief/intention predicates show an indicative/subjunctive alternation in the complement clause (Grano 2024). Giannakidou & Mari (2021) propose for Greek that indicative selection presupposes subjective veridicality of the embedded clause wrt the modal background provided by the embedding verb. Grano (2024) on Romance languages suggests that subjunctive complements involve eventuality abstraction, i.e., that their eventuality argument is not existentially bound. For English, Grano proposes that the same distinction is conveyed by finite vs. non-finite complements. This, however, does not readily apply to cases like Italian, where the two senses map to two infinitivals. More importantly, this proposal implies that complement clauses are of different semantic types, committing us either to assume lexical ambiguity of the embedding predicates, or to implement some type adjusting mechanism.

Empirical evidence We observe the following contrasts. First, (1a) allows the embedded PRO to be bound by either the matrix subject or object, whereas (1b) involves obligatory object control. This contrast aligns with Landau's (2021, 2024) distinction between predicative and logophoric control: in (1a), PRO may refer to either the AUTHOR or ADDRESSEE of the matrix clause (logophoric), whereas in (1b), this is not possible (predicative). According to Landau (2015, 2024), this distinction pertains to different layers of the infinitival clause. Second, modals are strictly prohibited with *a*-infinitive but allowed with *di* -infinitive (2), suggesting that *di-infinitive* includes a full TP, whereas *a*-infinitive only contains a lower functional layer capable of hosting low aspectual verbs (Cinque 2006; Grano 2015).

(2) Marco ha convinto Gianni **di**/***a** voler/poter/saper/dover studiare.
Marco has convinced Gianni *di*/**a* want/can/being-able/have-to study.INf

Third, only the content of the infinitive below *di* can be assessed for truth, unlike the infinitive selected by *a*. This implies that the former, but not the latter, is fully propositional.

(3) Marco ha convinto Gianni **di**/***ad*** avere un figlio, ma non è vero.
Marco has convinced Gianni that he has/*to have a child, *but it is not true.*

Our Proposal: Based on the above evidence, we propose that there is a structural difference between the two constituents selected by *di* and *a*. *a* selects minimally a vP infinitive, which denotes an event description with an abstracted individual argument, thus forcing *de se* readings (unless *x* is existentially bound) (Chierchia 1989; Grano 2024).

(4) $[[[_{vP}\text{PRO avere un figlio}]]] = \lambda x.\lambda w.\lambda e.\text{HAve-A-CHIld}(e, w) \,\&\, \text{AgeNT}(x, e, w)$

di, on the other hand, selects a larger structure, potentially capable of hosting other functional heads (cf. (2)). Given the propositional nature of this kind of infinitival, we also assume that its eventuality argument has been saturated through some cloSURe operator, realized by a high-TP head (see Grano 2024 for discussion). Therefore, the denotation for the infinitival selected by *di* is along the lines of (5)

(5) $[[\text{cloSURe PRO avere un figlio}]]]] = \lambda x.\lambda w.\exists e'.\text{HAve-A-CHIld}(e', w) \,\&\, \text{AgeNT}(x, e', w)$

We also assume, following Grano (2019), that the denotation for *convincere* takes a property of events *P* and two individuals *x* and *y* and returns another event property such that there is a convincing event with agent *y* and patient *x*, which causes a rational attitude eventuality with experiencer *x* and property *P* (see Grano 2019 for rational attitudes as a uniform class including belief and intention).

(6) $[\text{convincere}] = \lambda P.\lambda x.\lambda y.\lambda e.\exists e'.\text{CONVINCe}(e) \,\&\, \text{AgeNT}(y, e) \,\&\, \text{PATIeNT}(x, e)$
$\&\, \text{CAUSe}(e, e') \,\&\, \text{RATIONAl-ATTITUde}(e') \,\&\, \text{exPERIeNCeR}(x, e') \,\&\, P(e')$

Crucially, Grano (2024) suggests that the semantic encoding of intention requires eventuality abstraction of their complements, since the event argument must be causally related to the attitude eventuality. We propose that, in Italian, this relation is contributed specifically by *a*, thereby inducing the 'cause-to-intend' meaning. On the other hand, *di* simply selects for a proposition with an abstracted individual argument, connecting the attitude content with the attitude eventuality of the matrix clause, thus inducing a 'cause-to-believe' reading.

Further issues Interestingly, only (1a) can be rephrased with an indicative complement. In contrast, when the embedded clause is subjunctive, the interpretation aligns with that of (1b). This aligns with Grano's (2024: ft26) suggestion that existential closure occurs at a position below indicative Mood but above subjunctive Mood. Cinque (2006: 110) similarly notes a distinction in transparency effects, contrasting subjunctive and infinitive ("irrealis") complement with indicative forms. Together, these observations from Grano and Cinque independently imply a reduced structural size for subjunctive clauses, consistent with the parallel observed between *convincere-a-infinitive* and *convincere-che-subjunctive* below.

(7) Gianni ha convinto Mario che [abbia un figlio. ⇝ (1b)]/[ha un figlio ⇝(1a)]
Gianni$_1$ convinced Mario$_2$ that [*pro*$_2$ has.SUBJ a son]/[*pro*$_{1/2}$ has.IND. a son]

More aspects will be addressed, such as the irrealis property of *a-infinitive* (Wurmbrand 2001), empirical arguments for the existential closure, other relevant syntactic contexts (e.g., *avere paura di/a* 'to be afraid of'), the syntax of *di/a* (along the lines of Bocci & Rizzi (2017) and Belletti (2017), respectively).

SelecteD References Belletti, A. (2017) 10.1515/9781501504037-002 * Cinque, G. (2006) 10.1093/oso/9 780195179545.001.0001 * Giannakidou, A., Mari, A. (2021) 10.7208/9780226763484 * Grano, T. (2019) tgrano.pages.iu.edu/cls54_grano.pdf * Grano, T. (2024) 10.1007/s10988 -023-09397-y * Kratzer, A. (2006) * Landau, I. (2015) 978026

2527361 * Moulton, K. (2009) Ph.D. Thesis * Olivier, M., Sevdali, C., Folli, R. (2023) 10.16995/glossa.10135 * Portner, P., Rubinstein, A. (2020) 10.1007/s11050-020-09167-7

Explanatory Reasons, Actions, and the Right Fit

By Robert Reimer

In this talk, I am concerned with the nature of explanatory reasons for actions. Explanatory reasons are reasons that play the role of *explanantia* in *action explanations*. Their purpose is to make the agent's action intelligible and to let it appear rational from the agent's own perspective.

Many philosophers including Jonathan Dancy (2003) and Donald Davidson (2001) do not explicitly distinguish between explanatory reasons and motivating reasons – those reasons upon which the agent actually acted and which function as the premises in the agent's practical reasoning. They assume that explanatory reasons are just the agent's motivating reasons. And this makes sense insofar as explanatory reasons are supposed to let the action appear rational *from the agent's own perspective*. When Sarah calls the ambulance because John is bleeding, what motivates her action is also what explains it, namely *that John is bleeding*. Since both what motivates her action and what explains her action seems to be a *fact* – the fact that John is bleeding –, some philosophers including Joseph Raz (2009, 185) and Eric Marcus (2012, 72-73) conclude that motivating and explanatory reasons are *facts* – facts that the agent believed and acted on. This view on the nature of explanatory reasons is called 'factivism'.

Despite its intuitive appeal, factivism, faces a problem: It cannot explain cases in which the agent acted under the influence of an *epistemic error*. Assume that Sarah called the ambulance believing that John was bleeding whilst, in fact, he was just sullied with tomato juice. Cases like that are called '*error cases*'. In an error case, what seems to motivate the agent was not the case; and whatever is not the case, cannot be a fact. If, however, the agent's motivating reason was not the case, what, if anything, can explain her action?

For Dancy, the problem of error cases is not a problem, at all, because, according to him, sometimes "a nothing (something that is not the case) can explain a something (an action that was done)." (Dancy 2003, 427). To support this view, he introduces the notion of states-of-affairs – states that may or may not obtain. He argues that states-of-affairs (obtaining or not obtaining) explain and motivate actions *in any possible case*. In a veridical case (such as Sarah's calling the ambulance whilst John was actually bleeding), what explains the agent's action is a(n obtaining) state of affairs (that John was bleeding). In an error case (such as Sarah's calling the ambulance whilst John was not bleeding), what explains the agent's action is also a (non-obtaining) state-of-affairs (that John was bleeding). This view is called 'non-factivism'.

In contrast to Dancy, Maria Alvarez argues that proper action explanations require the truth of both the explanandum and the explanans (Alvarez 2018, 3300). She thinks that, in the case of an epistemic error, what explains the action cannot be *what* the agent believed and acted on but rather the agent's *belief* itself or a fact about it: "[I]n 'error cases' – cases when an agent acts on the basis of a falsehood that he believes – the explanans of a true explanation must be a psychological fact." (Ibid., 3300) So, what explains why Sarah called an ambulance cannot be that John was bleeding but rather the psychological fact that John was bleeding. This view is called 'psychologism', and it includes the positions of Davidson (2001) and Michael Smith (1994, 96; 2009) who argue that explanatory reasons are the agent's beliefs (and desires) themselves.

Prima facie, non-factivism and psychologism succeed in specifying explanatory reasons both for veridical cases and for error cases. The current stage of the debate on explanatory reasons, at least, suggests that. In my talk, however, I will show that this is not true. There is a group of error cases that neither factivism, non-factivism, nor psychologism can do justice to. The cases that I have in mind are cases in which the epistemic error concerns both the explanatory reason and the action itself. Consider the following situation: Jane discovers a basket full of fruits in the supermarket. She likes to eat peaches, and she thinks that these fruits in the basket are fresh and juicy peaches. However, the fruits in the basket are apples, and they are also old and wilted. Due to the dim light, however, she mistakes the apples for fresh and juicy peaches. Finally, she takes one of the old and wilted apples and puts it into her basket. Coming home, Paul asks her why she bought an old and wilted apple. Which reason could Jane explain her action?

To begin with, Jane cannot cite a fact to explain her action of taking an old and wilted apple because what she believed and acted on – that the fruits in the basket are fresh and juicy peaches – was not the case. But she can also not cite her beliefs themselves or the non-obtaining state-of-affairs, that she believed, to explain her action because these things do not *match* her action. How could her (false) belief or the nonobtaining state-of-affairs that the fruits in the basket are fresh and juicy peaches explain her action of taking an old and wilted apple? What is needed instead, it seems, is something that has both a factive and a psychological aspect. I am thinking of something along the following lines: the fact that Jane *mistook* the old and wilted apples for fresh and juicy peaches. This fact, at least, explains why Jane took an old and wilted apple. It makes her action intelligible, and it lets it appear rational from her own perspective. I call these facts – since they have both a factive and a psychological aspect – 'hybrid facts'.

The point of my talk is to show that the currently dominant views on explanatory reasons – factivism, non-factivism, psychologism – cannot do justice to error cases such as Jane's mistaken purchase and that a hybrid view with hybrid facts as explanatory reasons is needed. I first present those currently dominant views. Then, I introduce the case of Jane's mistaken purchase and show why only a hybrid fact can play the role of an explanatory reason. Finally, if time permits, I consider and reject two possible attempts to safe psychologism or non-factivism against my attack.

References

Alvarez, M. 2018. "Reasons for action, acting for reasons, and rationality." *Synthese* 95, 3293-3310.

Dancy, J. 2003. "Précis of Practical Reality." *Philos Phenomenol Res* 67 (2), 423-428.

Davidson, D. 2001. "Actions, Reasons, and Causes." In: *Essays on Actions and Events*. Oxford: Clarendon Press, 3-19.

Everson, S. 2009. "What are Reasons for Action?" In: Sandis, C. (ed.) *New Essays on the Explanation of Action*. London: Palgrave Macmillan, 22-47.

Marcus, E. 2012. *Rational Causation*. Cambridge, Massachusetts: Harvard University Press.

Raz J. 2009. "Reasons: Explanatory and Normative." In: C. Sandis (ed.) *New Essays on the Explanation of Action*. London: Palgrave Macmillan, 184-202.

Smith, M. 1994. *The Moral Problem*. Oxford: Blackwell.

Smith, M. 2009. "The explanatory role of being rational." In: D. Sobel & S. Wall (eds.) *Reasons for Action*. Cambridge: University Press, 58–80.

From artifacts to human lives: Investigating the domain-generality of judgments about purposes

Michael Prinzing & Joshua Knobe

People attribute purposes in both mundane and profound ways—such as when thinking about the purpose of a knife and the purpose of a life. In prior research, these seemingly very different kinds of purpose attributions have been investigated separately. One very plausible hypothesis is that people have fundamentally different ways of attributing purposes in these different domains. Yet, it is also possible that people use the same criteria when attributing purposes in a wide range of domains. If so, then this would reveal an important and surprising connection between programs of research that otherwise seem unrelated. In three studies (total N = 13,720 observations from N = 3,430 participants), we examined purpose attributions in six domains: artifacts, social institutions, animals, body parts, sacred objects, and human lives.

Study 1 manipulated what items in each domain were originally created for (original design) and how people currently use them (present practice). Results revealed that information about the intentions of the original designer and the intentions of the present users each influenced purpose attributions across all six domains. Yet, the sizes of these effects differed across domains. For artifacts and animals, original design and present practice had very similar effect sizes. However, for institutions and human lives, present practice had a larger effect than original design. The reverse pattern held for sacred objects and body parts.

Study 2 manipulated whether items in each domain are good at achieving a goal (effectiveness) and whether the goal itself is good (morality). The results were strikingly similar across domains. People are more inclined to think that a use is an item's purpose if it is highly effective (versus ineffective) for that use. Moreover, people are more inclined to consider morally good uses to be items' purposes than morally bad uses. But, this effect was very small.

In short, Study 1 provided some evidence for domain-specificity in purpose attributions, whereas Study 2 provided evidence for domain-generality. Thus, we ran a third study to test whether the Study 1 results might somehow be compatible with the hypothesis that purposes are attributed in the same ways across domains. We speculated that the impacts of the intentions of the original creators and present users might depend on the kinds of entities that play the roles of creator and user. For example, if the creator of a sacred object is a supernatural being, then perhaps people give greater weight to the creator's intentions because he is a god and not because he is the creator. If so, then, if we reversed which entities play each role (e.g., making a human the creator and a god the user), then we should also reverse the effect sizes for the original design and present practice manipulations.

Following this rationale, Study 3 used a similar design to Study 1 but with an additional experimental factor, "roles," which flipped the kinds of entities that play the roles of original designer and present user. Results revealed that, in most domains, the role-reversing manipulation substantially changed the effects of the original design and present practice

manipulations. Thus, the most striking differences that we observed across domains in Study 1 appear to result from details about which agents usually act as the creators versus users of the items. This suggests that the differences we originally observed don't reflect differences in the importance of original design and present practice across domains. Instead, they reflect inter-domain differences in the kinds of entities that usually play the roles of designer and user.

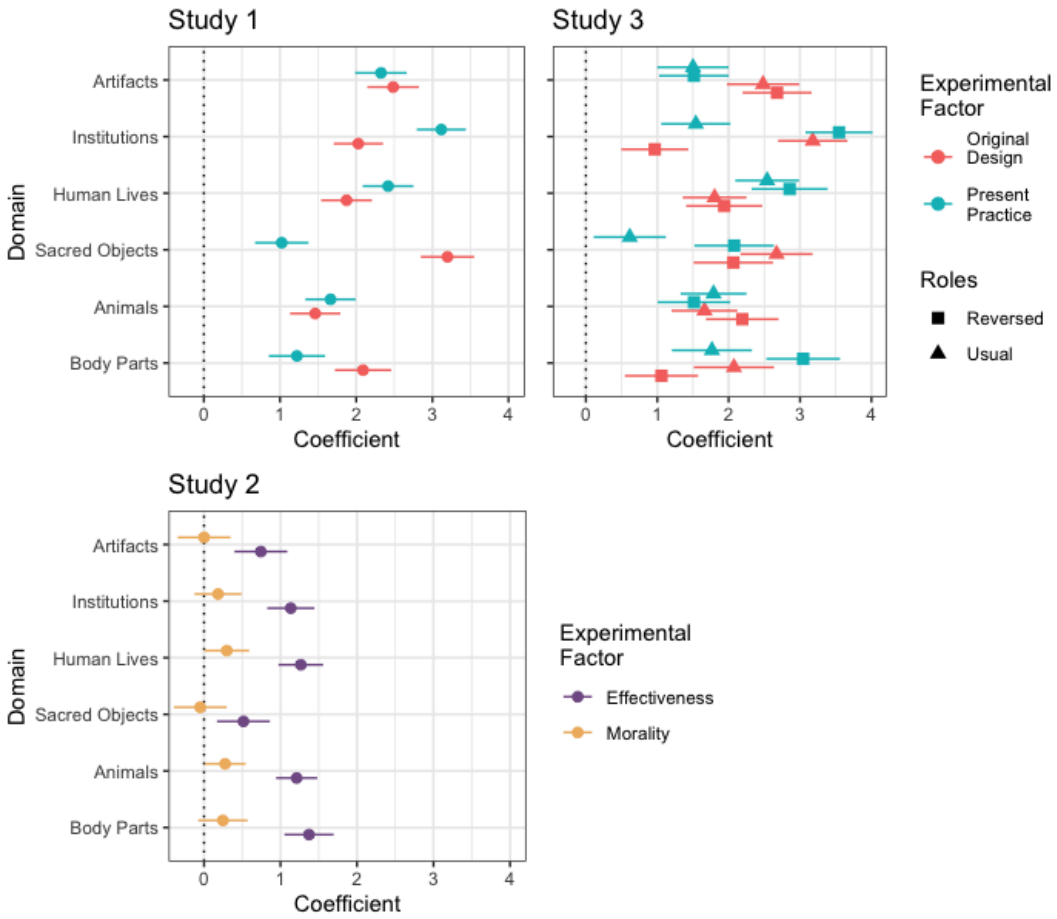

Figure 1: **Effects of experimental factors across domains.** Points and error bars indicate regression coefficients and 95% confidence intervals.

These studies suggest that the factors that influence purpose attributions are strikingly similar across domains. We found that the criteria people use when thinking about the purpose of an ordinary artifact like a chair or a knife are very similar to the criteria they use when thinking about the purpose a body part like a hand or a heart, or even when they are thinking about more deeply meaningful questions such as the purpose of a person's life. If this conclusion is correct, it would have dramatic implications for future research. Whereas reasoning about the purposes of ordinary artifacts seems relatively straightforward, reasoning about religiously or philosophically significant purposes seems complex and impenetrable. Yet, if the criteria for purpose attributions are domain-general, then insight into the more philosophically significant forms of teleological reasoning might be gained by considering findings from research on more mundane forms of teleological reasoning. To understand how people reason about the purpose

of a life, perhaps one should begin by considering how people reason about the purpose of a knife.

What makes an Agent an Agent? Comparing the semantic properties of Instruments and prototypical Agents in subject position

Alice Suozzi, Gianluca E. Lebani
Ca' Foscari University of Venice

This study aims to analyze Italian speakers' perception of the thematic roles of Agent, Patient, and Instrument through the lens of Dowty's (1991) Proto-Role Hypothesis. While previous studies (Kako 2006; Reisinger et al. 2015; White et al. 2016) have demonstrated that the thematic roles of Agent and Patient can be decomposed into basic semantic properties that are inferred by speakers from the syntactic roles of subject and object, as outlined by Dowty (1991), no prior research has specifically focused on Italian or considered the thematic role of Instrument, which constitute the primary novelties of this study.

The Instrument role has traditionally been characterized in causal terms (e.g., Fillmore 1968; Croft 1991; Talmy 2000), as "an entity that serves as an intermediary between agent and patient in a causal chain (Goldberg 2002: 340)". Although non-causal instrumental subroles have also been identified (cf. Marantz 1984; Levin & Rappaport Hovav 1988; Schlesinger 1995; Koenig et al. 2003, 2008), this study focuses on intermediary Instruments for two main reasons: (i) these Instruments are known to share the +cause property with Proto-Agents while being -volitional and -sentient, like Proto-Patients (Dowty 1991); (ii) only intermediary Instruments can undergo the Instrument-Subject Alternation (ISA), i.e., they can appear as subjects when the Agent is left unexpressed (e.g., *John opened the door with the key* / *The key opened the door*).

Intermediary Instruments are thus central to addressing our research questions: (RQ1) Are different properties attributed to Instruments occurring as subjects compared to subjects expressing prototypical Proto-Agents, as semantic properties are inferred from syntactic roles? (RQ2) Does the Instrument role align more closely with the Proto-Agent or the Proto-Patient, depending on its syntactic realization (PP *versus* subject)?

To answer RQ1 and RQ2, 39 pairs of sentences were created, as in (1):

(1) a. *Il pescatore ha catturato il pesce con una rete*
'The fisherman caught the fish with a net'
b. *La rete ha catturato il pesce*
'The net caught the fish'

Five defining properties of Dowty's Proto-roles were selected, starting from previous works (Kako 2006; Reisinger et al. 2015; White et al. 2016), i.e., VOLITIONAL INVOLVEMENT, SENTIENCE, CAUSED A CHANGE, UNDERGOES A CHANGE OF STATE, INDEPENDENT EXISTENCE. Participants were asked to rate, on a continuous scale (1-100), how much a constituent in a sentence displayed each of these five properties. The experiment was presented in two conditions. In the Agent-Instrument condition, 47 Italian native speakers rated how much the Agent-subject (e.g., *il pescatore* 'the fisherman') and the Instrument-subject (*la rete* 'the net') displayed the five properties, for all the sentence pairs. In the Patient-Instrument condition, other 47 Italian native speakers rated how much the instrumental PP (*con la rete* 'with the net') and the Patient-object (*il*

pesce 'the fish') displayed the five properties, for the same sentence pairs. The experiment was conducted online, via the Gorilla Experiment Builder platform (www.gorilla.sc).

The collected data are analyzed with a linear mixed-effects interaction model using *R* (R Core Team, 2023). The dependent variable is score. The fixed effects are the XP/role (i.e., Agent, Patient, Instrument-subject and instrumental PP) and the five properties. Participants and items are considered as random effects. This analysis revealed a significant interaction between the two fixed effects (random intercept $p < .001$). Namely, Agent-subjects are rated significantly higher than Patients, Instrument-subjects and instrumental PPs concerning the VOLITIONAL INVOLVEMENT, SENTIENCE, CAUSED A CHANGE properties ($p < .001$). Patients were ranked significantly higher than the other XP/roles concerning the UNDERGOES A CHANGE OF STATE property ($p < .001$). Finally, no significant difference was found regarding the INDEPENDENT EXISTENCE property ($p > .05$).

The collected data reveal that Agent-subjects are perceived as more volitional and more sentient than Instruments, even when Instruments are realized as subjects (RQ1). Instruments are especially perceived as entities that cause a change, regardless of their syntactic realization (RQ2). They align with Agents in this regard, but with Patients concerning the absence of volition and sentience, thus occupying an intermediate position between Proto-Agents and Proto-Patients, consistently with Dowty's (1991) proposal.

References

Croft, W. (1991). *Syntactic Categories and Grammatical Relations*. Chicago University Press.

Dowty, D. R. (1991). Thematic proto-roles and argument selection. *Language*, 67(3): 547–619.

Fillmore, C. J. (1968). The case for case. In E. Bach & R. T. Harms (eds.), *Universals in Linguistic Theory*, 1-87. Holt, Rinehart, and Winston.

Goldberg, A. E. 2002. Surface generalizations: an alternative to alternations. *Cognitive Linguistics,* 13, 327–56.

Kako, E. (2006). Thematic role properties of subjects and objects. Cognition, 101(1), 1-42.

Koenig, J.-P., Mauner, G., Bienvenue, B. (2003). Arguments for adjuncts. *Cognition*, 89, 67-103.

Koenig, J.-P., Bienvenue, B., Mauner, G., Conklin, K. (2008). What with? The Anatomy of a (Proto)-Role. *Journal of Semantics*, 25(2), 175-220.

Levin, B. & M. Rappaport-Hovav. (1988). Nonevent *-er* Nominals: A Probe into Argument Structure. *Linguistics,* 26, 1067–1083.

Marantz, A. P. (1984). *On the Nature of Grammatical Relations*. MIT Press.

R Core Team. (2023). R: A language and environment for statistical computing (Version 4.3.1). R Foundation for Statistical Computing. https://www.R-project.org/.

Reisinger, D., Rudinger, R., Ferraro, F., Harman, C., Rawlins, K., & Van Durme, B. (2015). Semantic Proto-Roles. *TACL*, 3, 475-488.

Schlesinger, I. M. (1995). *Cognitive space and linguistic case*. Cambridge University Press.

Talmy, L. (2000). *Toward a cognitive semantics Vol. 1. Concept Structuring Systems*. MIT Press.

White, A.S., Reisinger, D., Sakaguchi, K., Vieira, T., Zhang, S., Rudinger, R., Rawlins, K., & Durme, B.V. (2016). Universal Decompositional Semantics on Universal Dependencies. *Proc. of EMNLP 2016*: 1713–1723.

Implicativity in Interaction: Effects of conversational context and lexical meaning

Fahima Ayub Khan[1], Eleni Gregoromichelaki, Christine Howes, Robin Cooper
University of Gothenburg

Inferences derived from utterances including presupposition triggers reflect public commitments (de Marneffe et al, 2019; Cornillie, 2018) whose negotiation in dialogue can be taken as an alternative basis for grounding interaction coordination instead of communicating Gricean intentions (see, e.g., Gregoromichelaki et al, 2010; Geurts, 2019). As such, the particular presupposition triggers available in a language can be taken as conventional resources (*affordances*, Gregoromichelaki et al, 2020) for establishing common commitments in interaction. In contrast, Annamalai and Levinson (1992) claim that presupposition triggers and their conversational behaviour are parallel across languages, e.g., the defeasibility properties of all presupposition triggers are similar between English and Tamil, revealing some kind of universal conceptual basis for pragmatic inference. However, their presentation does not cover the full range of presupposition triggers, especially in the category of *implicative verbs*, i.e., verbs which entail the truth of the complement clause while carrying presuppositions constraining the context of utterance (Karttunen, 1971). Nadathur (2023) accounts for such entailments and contextual constraints under a causal model which links the lexical content of the verb to the set of inferences eventually licensed by the sentence. On the other hand, White (2019) points out the contribution of the syntactic and morphological features of particular languages in implicative constructions and argues that the causal framework does not fully account for predicates such as 'remember' and 'forget' since the implicative entailments are rather *actuality entailments* realised by the modal presupposition of these verbs (i.e., 'obligation'). While these studies mostly focus on implicative presuppositions (*not-at-issue*) and entailments (*at-issue*) triggered by lexical items at the individual level, very little is known about how such inferences are co-constructed and processed in interactive contexts where agency and responsibility for common commitments are distributed between interlocutors.

As a preliminary to an investigation of the whole range of factors predicting inferences resulting from the use of implicative constructions, the present study examines how individuals process at-issue implicative inferences (so-called "entailments") in English and Tamil. Specifically, we assume that implicative verbs such as 'remember', 'manage', 'forget', and 'fail' display variable implicative inferences cross-linguistically depending on contextual and lexicogrammatical constraints (cf. Levinson & Annamalai, 1992). For example, we have observed that affirmative assertions involving *remember* in English generally commit the speaker to the truth of the sentential complement (as in (1)), whereas, in Tamil (2), this is not the case, in fact, there is no necessary implication to the truth of complement:

(1) She remembered to lock the door → *She locked the door*

(2) Aval-ukku kadhav-ai poot-a nyabagam-irundha-dhu
she-DAT door-ACC lock-INF memory-have-3.SG.PST.N
She remembered to lock the door →*? She locked the door*

On the other hand, (1) also involves a speaker commitment to the fact that the subject of the main clause was under some obligation or constraint to lock the door, which also holds in Tamil. The speaker's commitment to the truth of the complement of the implicative verb can become the target of a question-answering pair (as in (3) and (4)):

(3) A: Did she remember to lock the door?
B: Yes(, she did lock the door).

(4) A: Avalukku kadhavai poota nyabagam irundhadha?
she-DAT door-ACC lock-INF memory-have-3.SG.PST.Q
Did she remember to lock the door?

[1] fahima.ayub.khan@gu.se

B: Aama, nyabagam irundhadhu
Yes, memory-have-3.SG.PST
Yes, (she) remembered

Unlike the presentation of isolated sentences in a monological context, this use of the implicative construction in dialogical interaction in Tamil has an impact on the potential inferences that can be generated. The inference to the truth of the complement now becomes much more likely.

To investigate such variable inferences, we implemented a judgement task to test the inferential profile of the verbs *remember*, *manage*, *forget* and *fail* in English and Tamil. We recruited 20 native English speakers and 20 native Tamil speakers between the ages of 18 to 75 as participants. The speakers of each language were presented with target items where each implicative verb appeared in an assertion (*monological context*) or in a question-and-answer within a conversation (*conversational context*). This target item is followed by a polar question (see Fig. 1 & 2) where subjects are asked to select an appropriate response with 'yes', 'maybe', and 'no' as the options. Responses were analysed using a mixed-effects logistic regression model.

Our findings reveal that speaker commitment to the truth of the complement was stronger in English in both monological and conversational contexts in comparison to Tamil. However, especially in the case of 'remember' and 'manage', targeting the truth of the complement in questions and through sentential ellipsis (*yes*) in the conversational context results in stronger inferences about its truth in Tamil. This indicates that perceived speaker commitments can vary not only due to the semantic/conceptual content encapsulated in implicative triggers and constructions but also depending on perceived speaker goals in a conversation. Thus, the results provide preliminary evidence that an appropriate account of the processing conditions of such verbs requires not only a fine-grained account of their conceptual structure and the particular syntactic constructions in which they appear but also taking into account the influence of constraints arising from the structure of a fine-grained conversational model. We model cross-linguistic inferences of implicative verbs using the framework of DS-TTR (Gregoromichelaki, 2018) which provides the appropriate contextualisation of lexicogrammatical constraints within a dialogue processing model while taking speaker commitments and goals into consideration.

Supplementary materials

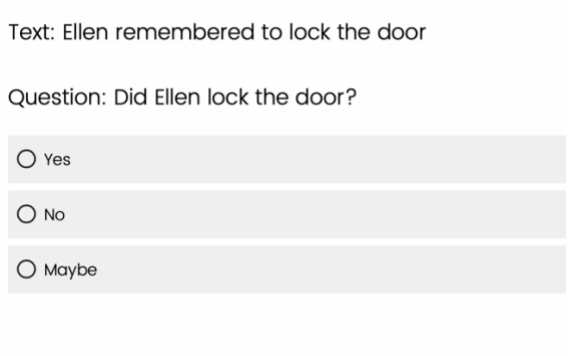

Fig 1. Monological context – English

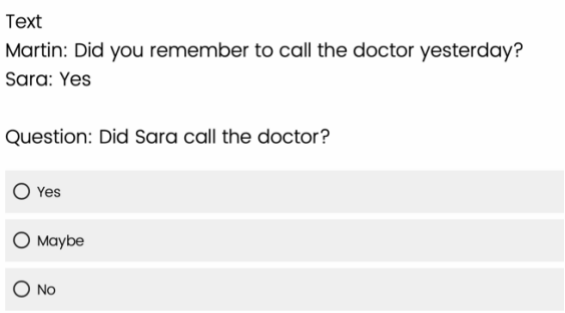

Fig 2. Conversational context – English

References

Cornillie, B. (2018). On speaker commitment and speaker involvement. Evidence from evidentials in Spanish talk-in-interaction. *Journal of Pragmatics*, *128*, 161-170.

De Marneffe, M. C., Simons, M., & Tonhauser, J. (2019, July). The commitmentbank: Investigating projection in naturally occurring discourse. In P*roceedings of Sinn und Bedeutung* (Vol. 23, No. 2, pp. 107-124).

Geurts, B. "Communication as commitment sharing: speech acts, implicatures, common ground" *Theoretical Linguistics*, vol. 45, no. 1-2, 2019, pp. 1-30. https://doi.org/10.1515/tl-2019-0001

Gregoromichelaki, et al. (2010). Incrementality and intention-recognition in utterance processing *Dialogue and Discourse* 2 (1), 199-233

Gregoromichelaki, et al. (2020) Affordance competition in dialogue. ACL 2020
Gregoromichelaki, E., (2018). Quotation in Dialogue. In Saka, P. and Johnson, M. (eds.) The *Pragmatics of Quotation*. Springer.
Karttunen, L. (1971). Implicative verbs. *Language*.
Levinson, S. C., & Annamalai, E. (1992). Why presuppositions aren't conventional. In *Language and text: Studies in honour of Ashok R. Kelkar* (pp. 227-242). Kalinga Publications.
Nadathur, P. (2023). Causal semantics for implicative verbs. *Journal of Semantics*, *40*(2-3), 311-358.
White, A. S. (2019). Lexically triggered veridicality inferences. *Handbook of Pragmatics*, *22*, 115-48.

Lexical Access of Verbs in Parkinson's Disease: Does Agency Matter?

Maura Panozzo Chiomento, Maria Vender, Denis Delfitto

University of Verona

Introduction

Parkinson's disease (PD) is a neurodegenerative condition commonly associated with movement-related symptoms, although non-motor symptoms are also prevalent in nearly all patients (99.7% according to Rodriguez-Blazquez et al., 2020). Linguistic deficits in PD have been observed in many domains including naming tasks, where participants are presented with a picture and are asked to quickly and accurately utter the corresponding word. Research shows that PD patients have more difficulties when naming high-motion verbs compared to low-motion verbs (Herrera et al., 2012; Bocanegra et al., 2017) and find verbs describing fast actions more challenging than those encoding slow actions (Speed et al., 2017). Interestingly, abstract verbs, which lack motion-related semantics, are spared (Fernandino et al., 2013). Other findings attribute verb naming difficulties in PD to syntactic (A-structure complexity) rather than semantic reasons, with PD patients performing poorly compared to controls when naming transitive and unergative verbs but not unaccusative ones (Aiello et al., 2022). To summarise, previous literature has found that naming performance for verbs in PD seems to be influenced by semantic factors like motion and speed, although syntactic-related criteria have also been proposed. However, the role of agency in naming verbs has not been investigated so far. Our study aims precisely to address this gap in the literature, comparing PD patients and healthy adults in a lexical access task eliciting different classes of verbs in which the Agency component has been manipulated.

Methods

The research protocol included three groups: individuals with PD ($N = 31$), age-, sex- and education-matched neurotypical healthy subjects ($N = 31$), and younger controls ($N = 31$). Sample size was determined using a G*Power analysis (effect size of $d = 1$, power of 0.95, $\alpha = 0.05$), which recommended 27 subjects per group. All participants underwent neuropsychological evaluations, encompassing the Frontal Assessment Battery (FAB), the Hospital Anxiety and Depression Scale (HADS), the Montreal Cognitive Assessment (MoCA) to perform a cognitive screening and exclude the presence of mild cognitive impairment in the tested subjects, and an eye-tracking-based Trail Making Test A and B (TMT). Additionally, a neurologist assessed PD patients while OFF medication using the Unified Parkinson's Disease Rating Scale (UPDRS).

The lexical access task that we developed and administered to address our research questions involved naming verbs from five categories determined by the linguistic theory:

(1) unergative verbs with "internal" agentivity (Pinker, 2007), e.g., *to yawn*;
(2) unergative verbs, e.g., *to walk*;
(3) transitive verbs of the "achievement" type (Vendler, 1967), e.g., *to discover*;
(4) transitive verbs of the "accomplishment" type (Vendler, 1967), e.g., *to cook*;
(5) unaccusative verbs, lacking agentivity and constituting our control items, e.g., *to fall*.

The stimuli were AI-generated using DALL-E (see Picture 1) and presented to participants on the computer screen with the software E-Prime 3. Reaction times and accuracy data were collected during the picture naming task.

Expected Results and Implications

Data collection has been completed, and analysis is currently underway, with results available by the time of the conference. The varying degrees of agentivity in the verbal stimuli will enable us to test the hypothesis that this linguistic feature modulates verb naming abilities in individuals with PD. If individuals with PD are sensitive to the Agency component of verbs, the presence of an agent, common to both transitive and unergative verbs, may make these verbs more challenging. This hypothesis is supported by studies indicating altered "awareness of the authorship of action" in PD, which relates to the perceived sense of agentivity (Saito et al., 2017).

The findings of this study will refine the linguistic theory by linking verb naming deficits in PD to broader cognitive and neuropsychological processes while enriching the clinical understanding of PD by identifying linguistic symptoms that might hold diagnostic significance.

Picture 1: Examples of visual stimuli representing the five verbal categories involved

References

Aiello, E. N., Grosso, M., Di Liberto, A., Andriulo, A., Buscone, S., Caracciolo, C., Ottobrini, M., & Luzzatti, C. (2022). Disembodying language: Actionality does not account for verb processing deficits in Parkinson's disease. *Journal of Neurolinguistics*, *61*, 101040. https://doi.org/10.1016/j.jneuroling.2021.101040

Bocanegra, Y., García, A. M., Lopera, F., Pineda, D., Baena, A., Ospina, P., Alzate, D., Buriticá, O., Moreno, L., Ibáñez, A., & Cuetos, F. (2017). Unspeakable motion: Selective action-verb impairments in Parkinson's disease patients without mild cognitive impairment. *Brain and Language*, *168*, 37–46. https://doi.org/10.1016/j.bandl.2017.01.005

Fernandino, L., Conant, L. L., Binder, J. R., Blindauer, K., Hiner, B., Spangler, K., & Desai, R. H. (2013). Where is the action? Action sentence processing in Parkinson's disease. *Neuropsychologia*, *51*(8), 1510–1517. https://doi.org/10.1016/j.neuropsychologia.2013.04.008

Herrera, E., Rodríguez-Ferreiro, J., & Cuetos, F. (2012). The effect of motion content in action naming by Parkinson's disease patients. *Cortex*, *48*(7), 900–904. https://doi.org/10.1016/j.cortex.2010.12.007

Pinker, S. (2007). *The stuff of thought: Language as a window into human nature*. Viking.

Rodriguez-Blazquez C., Schrag A., Rizos A., Chaudhuri K.R., Martinez-Martin P., Weintraub D. (2020). Prevalence of Non-Motor Symptoms and Non-Motor Fluctuations in Parkinson's Disease Using the MDS-NMS. *Movement Disorder Clinical Practice*, *8*(2), 231-239. https://doi.org/10.1002/mdc3.13122

Saito, N., Takahata, K., Yamakado, H., Sawamoto, N., Saito, S., Takahashi, R., Murai, T., & Takahashi, H. (2017). Altered awareness of action in Parkinson's disease: evaluations by explicit and implicit measures. *Scientific Reports*, *7*(1), 8019. https://doi.org/10.1038/s41598-017-08482-0

Speed, L. J., Van Dam, W. O., Hirath, P., Vigliocco, G., & Desai, R. H. (2017). Impaired Comprehension of Speed Verbs in Parkinson's Disease. *Journal of the International Neuropsychological Society*, *23*(5), 412–420. https://doi.org/10.1017/S1355617717000248

Vendler, Z. (1967). Verbs and Times. In *Linguistics in Philosophy* (pp. 97–121). Cornell University Press.

Large Language Models in Large Language Games

Chris Cousens, The University of Glasgow

When I receive an email making me a promise, giving me an order, or offering me a job, have I in fact been promised, ordered, or offered? Until recently, this would have likely been easily answered in the affirmative. We have become quite accustomed to the performance of online speech acts.

The recent proliferation of large language models might seem to undermine this. We can no longer be (as) confident that an email was created by a human 'speaker'. As LLM technology is increasingly embedded into the software and digital platforms used in the workplace (where speech acts have elevated moral significance) we are faced with a dilemma. Either we stop responding to online promises, orders, and offers *as* promises, orders, and offers, recognising illocutionary force only in the utterances whose author we *know* to be human. Or we respond to the 'promises', 'orders', and 'offers' of LLM-generated 'speech' as fully-fledged speech acts. Fisher (2024) argues that the integration of LLM and human speech threatens our ordinary speech practices. I agree—but what to do about it?

If we adopt the restrictive response, we preserve received views in speech act theory (invoking Austin 1962; Strawson 1964). Speech acts are often taken to require *intention* from the speaker, which LLMs seem to lack. The predominant alternative view takes speech acts to require *uptake* from the audience—but this is typically cached out as the recognition of speaker intention (Lance and Kukla 2013; McDonald 2022)! So, speech acts require (some sort of) speaker intention, we can no longer be confident that this is present for online speech acts, so we can no longer confidently give uptake to emailed speech acts. Speech act theory remains secure, but at the cost of our ability to promise, order, and offer online.

If we adopt the expansive response and allow that LLMs can perform speech acts, we reject a key component of traditional speech act theory regarding speaker intention (and the nature of uptake). Large language models would be no mere 'stochastic parrots' (Bender et al 2021). This may also entail that speech acts can be performed unintentionally by human speakers. We would then need an alternative conception of illocutionary force to explain this.

I will suggest a potential solution, grounding the force of an utterance-token in its contribution to the 'conversational score' (see Lewis 1979), rather than the intention of the speaker or uptake of the audience. This provides another reason to adopt the speech-act theoretic framework I have developed in earlier work (Cousens 2023, 2024). It preserves our ordinary online communicative practices, but at the cost of some received views in speech act theory—a cost I think we should be willing to pay. And it does so without making the concessions of other attempts to allow for AI to perform speech acts (e.g. Green and Michel 2022).

References

Austin, J. L. (1962). *How to do things with words*. Oxford University Press.

Bender, E., Gebru, T., McMillan-Major, A., & Mitchell, M. (2021). On the dangers of stochastic parrots: Can language models be too big? *FAccT '21: Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, 610-623, https://doi.org/10.1145/3442188.3445922.

Cousens, C. (2023). Tweet Acts and Quote-Tweetable Acts. *Synthese, 202*(187).

Cousens, C. (2024). Exercising Illocutionary Power, Or: How To Do Things With Other People's Words. In *Harmful Speech and Contestation*, ed. M. Popa-Wyatt, Palgrave Macmillan, 85-108.

Fisher, S. (2024). Something AI Should Tell You – The Case For Labelling Synthetic Content. *Journal of Applied Philosophy*, https://doi.org/10.1111/japp.12758.

Green, M., & Michel, J. (2022). What might machines mean? *Minds and Machines*, *32*, 323-338.

Lance, M., & Kukla, R. (2013). Leave the gun; take the cannoli! The pragmatic topography of second-person calls. *Ethics*, *123*(3), 456-478.

Lewis, D. (1979). Scorekeeping in a Language Game. *Journal of Philosophical Logic*, *8*, 339-359.

McDonald, L. (2022). Reimagining Illocutionary Force. *The Philosophical Quarterly*, *72*(4), 918-939.

Strawson, P. F. (1964). Intention and convention in speech acts. *Philosophical Review, 73*(4), 439-460.

Degrees of Agential Substantive (Ir)Rationality

Crystal Seet, University of Wisconsin-Madison

Agency and Intentions in Language 5 (AIL5)

Consider the following case:

> Talia intends to count the number of blades of grass in her garden, even though doing so brings her no great pleasure. Since counting the blades of grass in her garden brings Talia no great pleasure, serves no other worthwhile purpose, and comes with a significant opportunity cost, Talia lacks sufficient reason to count the number of blades of grass in her garden. Intending to do so is thus substantively irrational. ...Though Talia knows that, in order to complete the count, she must keep track of how many blades she has counted so far, she can't be bothered to keep track and so doesn't intend to. [This] seems to involve Talia in a second kind of irrationality. [The] problem isn't that she has strong reason to keep track of the number of blades of grass she's counted, but fails to intend to do so. ...Rather, the mistake is the distinctive, structural one of failing to intend the means to her ends.[1]

This passage illustrates the distinction between substantive and structural rationality. Roughly, substantive rationality is about how well one responds to one's reasons, while structural rationality is about how coherent one is with regards to one's attitudes. Contrast Talia with her sister, Tania:

> Tania also intends to count the number of blades of grass in her garden, [but] she intends to keep track of her progress. Tania is more structurally rational than Talia since she exhibits means-end coherence. But she is, if anything, less substantively rational than Talia, since she has two intentions that she lacks adequate reason for whereas Talia only has one.[2]

This interpretation of the case—that Tania is more substantively irrational than Talia—exemplifies the:

> **Counting Intuition**: If an agent S holds a propositional attitude p that goes against S's reasons, and S infers from p to a further attitude q that also goes against S's reasons, then, all else equal, holding q makes S more substantively irrational.

Tania, by some practical inference, comes to have an additional intention unsupported by her reasons compared to Talia, so, according to the Counting Intuition, she is more substantively irrational than Talia.

---

1. Alex Worsnip (2021), *Fitting Things Together: Coherence and the Demands of Structural Rationality* (Oxford University Press), 6.

2. Worsnip, 6.

I argue that the Counting Intuition is false. Insofar as our objects of evaluation are agents—call this *agential* rationality—I propose that degrees of agential substantive rationality should be determined not by one's attitudes that do not conform to one's reasons, but by the quality of one's substantive reasoning that generated those attitudes. On this view, agents are never more substantively irrational simply in virtue of holding an additional attitude that fails to conform to their reasons.

Why might some find the Counting Intuition plausible in the first place? I start by examining how the Counting Intuition might be theoretically supported. I consider what I'll call a Balancing View of agential substantive rationality. Philosophers frequently talk about rationality in terms of one's *balance of reasons*.[3] Agents have (normative) reasons for or against φ-ing, each with their assigned weights, and we can weigh one's reasons for φ-ing on the one hand, and one's reasons against φ-ing on the other. To evaluate an agent's substantive rationality, we look at how closely an agent's decisions match their balance of reasons.

> **The Balancing View**: *S*'s degree of substantive irrationality is determined by the total weight of reasons against each of *S*'s attitudes.

The Balancing View looks at the weight of reasons against each attitude, attitude by attitude. Let's start with Talia's and Tania's initial intention to count blades of grass. Talia and Tania have identical reasons against counting blades of grass, so they have the same total weight of reasons against their intention to count blades of grass. Talia and Tania are even so far. The same reasons against counting blades of grass also oppose keeping track of their progress. Tania intends to keep track of her progress, while Talia lacks such an intention. Tania has non-zero weight of reasons against her subsequent intention to keep track of her progress while Talia does not. Finally, let's combine these weights across attitudes for each agent. The combination of the total weight of reasons against each of Tania's attitudes outweighs the combination of the total weight of reasons against each of Talia's, so Tania is more substantively irrational than Talia on the Balancing View.

Here's a counterexample to the Balancing View. Amy and Bob are two executives in a meeting about their company's marketing decisions. They've been presented with market research that indicates that advertising campaign *A* would be more efficient than campaign *B*. They agree on campaign *A* without much discussion. While Amy made her decision based on the evidence presented by their market research team, Bob made the same decision simply because the colour scheme of the advertisements for campaign *A* contains his favourite colours.

On the Balancing View, Amy and Bob are equally substantively irrational since they have the same reasons available to them and they arrive at the same decision.[4] However, Bob is intuitively more substantively irrational than Amy; Bob's process of reasoning seems to have gone awry whereas Amy's

---

3. See Bratman (1987, 2009), Broome (1999, 2013), Kiesewetter (2017), Lord (2018), Scanlon (1998, 2014), Schroeder (2021).

4. To resolve any lingering worries about this case, we can further stipulate that Bob and Amy have the same taste in colours, and that Bob did not pay attention to whether those colours that happen to be his favourite have been shown to be effective in advertisements.

has not. This suggests that an important feature of agential substantive rationality is the quality of one's reasoning. The Balancing View overlooks this feature by construing agential substantive rationality solely in terms of agents' attitudes and balance reasons. We've identified a desideratum for theories of agential substantive rationality: a plausible theory of agential substantive rationality ought to produce the verdict that agents like Amy and Bob—who have identical reasons available to them, form identical attitudes, but differ in the quality of their processes of reasoning—differ in their degrees of substantive rationality.

I propose a Reasoning View of agential substantive rationality that meets this desideratum. To sketch the Reasoning View, I'll start by assuming: if the distinction between substantive and structural rationality is prima facie plausible, then a distinction between substantive and structural reasoning is also prima facie plausible. Call these substantive and structural reasoning. Since substantive rationality is about responding to one's reasons, an agent's substantive reasoning is just their process of responding to their reasons: whether the reasons one takes oneself to have to φ really bear on φ-ing, favour φ-ing, and are as weighty as one takes those reasons to be. Structural reasoning, on the other hand, is concerned with the application, in some very loose sense, of structural requirements: e.g., when one deliberates about the means to one's end when one has an end, when one realises that one has contradictory beliefs or cyclical preferences when they are brought to the forefront of one's mind, etc. If agential rationality is about reasoning, and a distinction between substantive and structural reasoning is plausible, then it's worth exploring how we might understand agential substantive rationality in terms of substantive reasoning.

> **The Reasoning View**: One's degree of substantive irrationality is determined by the quality of one's reasoning.

On the Reasoning View, Amy and Bob have different degrees of agential substantive irrationality because they differ in their substantive reasoning: Amy took the evidence presented by the market research team to be a decisive reason to choose campaign A, while Bob commits an error in substantive reasoning by taking his taste in colours as a decisive reason to choose campaign A.

Let's return to the case of Talia and Tania. A corollary of the Reasoning View is that if two agents engage in identical substantive reasoning, then they are equally substantively irrational. Here's an intuitive interpretation of their cases: Talia and Tania engage in identical substantive reasoning to arrive at their initial intentions to count blades of grass, and then diverge in their structural reasoning to wind up in different places about their (lack of) intention to keep track of their progress. If this is right, then Talia and Tania are equally substantively irrational on the Reasoning View. Hence, the Reasoning View opposes the Balancing View and the Counting Intuition; it's possible for two agents to be equally substantively irrational even if one has an additional attitude that goes against their reasons, and agents are never more substantively irrational simply in virtue of holding an additional attitude that fails to conform to their reasons.

References

Bratman, Michael (1987). *Intention, Plans, and Practical Reason*. Cambridge, Mass.: Harvard University Press.

Bratman, Michael (2009). Intention, practical rationality, and self-governance. *Ethics* 119 (3):411-443.

Broome, John (1999). Normative requirements. *Ratio* 12 (4):398–419.

Broome, John (2013). *Rationality Through Reasoning*. Malden, MA: Wiley-Blackwell.

Kiesewetter, Benjamin (2017). *The Normativity of Rationality*. Oxford, UK: Oxford University Press.

Lord, Errol (2018). *The Importance of Being Rational*. Oxford, UK: Oxford University Press.

Scanlon, Thomas (1998). *What We Owe to Each Other*. Cambridge: Harvard University Press.

Scanlon, Thomas (2014). *Being Realistic About Reasons*. Oxford: Oxford University Press.

Schroeder, Mark (2021). *Reasons First*. New York: Oxford University Press.

Worsnip, Alex (2021). *Fitting Things Together: Coherence and the Demands of Structural Rationality*. New York: Oxford University Press.

Agency and Intentions in Language 5 (Online, 29-31 January 2025)

CAN COMPARTMENTALIZATION MAKE SENSE OF AN AGENT HOLDING INCONSISTENT BELIEFS?

Lawrence Lengbeyer

While there is freedom to fantasize or suppose as one wishes, believings are more tightly constrained. It would seem, for one thing, that one cannot believe something that one does not regard as true. And yet people do appear to believe propositions that they know to be highly dubious or even false, where the evidence as they understand it does not to their minds support the propositions' truth. How can this be? How can cognitive states that are epistemically irrational in being insensitive to known evidence still qualify as *beliefs*, given that in disregarding evidence they seem not sufficiently regulated by standards of truth? In this paper, I develop recent accounts of cognitive compartmentalization (by Mandelbaum, and Lengbeyer) in order to explain how such epistemically-irrational believings are possible, in part by responding to skeptics (Evnine, and Bortolotti) who deny that a compartmentalized cognitive architecture can satisfactorily play this explanatory role.

Beliefs, understood as specific propositions or sentences held with a distinctive kind of attitude, reside (along with images) within the individual agent's mental repository of representational resources for cognizing. The beliefs must satisfy a distinctive set of constraints, or they do not function as beliefs. The constraints are commonly taken to include what we might term Stability (beliefs are enduring attributes of the agent), Governance (they govern reasoning whenever they are relevant to the topic of thought), Unification (they afford and constitute a single, unified, coherent outlook, with exceptions being irrational anomalies), and Truth-directedness (they reflect what their possessors take to be true).

This standard "Integrationist" picture of cognition, whereby beliefs are the cognizing instruments of a thoroughly unified self and are all ready to contribute to each cognitive process, has come under challenge by those who have noted familiar psychological phenomena that seem irreconcilable with it. Ordinary language-users are surprisingly tolerant of folk psychologizing that allows for cognitive non-integration and willingly ascribes beliefs even where the Governance, Unification, and/or Truth-directedness conditions are not met. While there may be senses of "belief" that make it infelicitous to ascribe belief in a proposition thought by the agent to be unsupported by the evidence, there is a familiar sense for which such ascriptions pose no problem.

Even those open to the possibility of mental compartmentalization, however, tend to see it as incapable of all the heavy lifting needed to make sense of inconsistent beliefs. For Evnine, for example, "[p]artitions arise to ward off potential conflict between different sides of our lives" (95)—like Geach's Japanese astronomer who treated the sun alternately as a divinity and as a natural body to be scientifically investigated (94)—and can thus be of psychic service in limited life circumstances. But quirky exceptions cannot overcome the fundamental need for that "single, integrated informational system provided by belief," because an agent must react flexibly to a complicated world using "everything a person knows that might be *relevant*" (emphasis added).

Yet Evnine is too quick to assume that relevance means *topic*-relevance, and that an agent must adopt one and the same overall outlook and reasoning toolkit at all times. Such a single cognitive stance might be maximally functional were we to inhabit a homogeneous lifeworld. But that seems not to be our lifeworld. Ours is fragmented into diverse, layered, and quickly-shifting pragmatic situations or immediate mental tasks. In this world, the inconsistent representational resources enabled by a compartmentalized cognitive endowment can be exploited to productive effect. Compartmentalization also facilitates an urgently-needed abbreviation of the process of recruiting cognitive resources for the agent's immediate current project, by storing resources in clusters that are keyed to recognizable contexts of purposive activity that have been experienced as units in the past. All of this cognitive division brings possibilities for resistance to evidence that bears upon our cognitive commitments, and hence for *epistemically*-irrational believing that may or may not be *practically* rational.

Compartmentalization has the edge over Integration in handling conflicting evidence, because discrepancies in our fragmented lifeworld are not always within our powers to reconcile, and sometimes they are not within our interests, either. Compartmentalization allows us to subsist comfortably within these bounds. It lets good use be made of *special-purpose ideas* whose limitations are not exposed (or 'corrected') by the ideas having to interact with certain of their peers (as they would in a thoroughly integrated mind). One such special-purpose cognitive asset is the *simplified idea*: a more elementary version of a representation that also resides in one's cognitive pool, which in its proper context can be more useful than the non-simplified one would be.

Bortolotti, for her part, takes Compartmentalism to lack the resources to make sense of cases where people are aware of their inconsistencies but refuse to correct these. Her 'exhibit A' is Kahneman & Tversky's "Asian Disease Problem," where experimental subjects reversed their preferences between two options based simply upon the linguistic framings of the options (2010: 80-81), and would not alter their answers after debriefing (87). But she seems to presuppose that we inhabit a single, all-encompassing domain within which all propositions can be logically combined, as opposed to a world divided into pragmatic spaces that might call for intellectual stances that are differing and incompatible. A Compartmentalism that indexes cognition to pragmatic task contexts can account for the Asian Disease Problem by showing how Kahneman & Tversky caught their participants between two notions of cognitive relevance—the familiar one defined on sameness of substantive topic, and one defined on pragmatic utility for the immediate practical thinking task at hand.

Compartmentalism thus brings with it a metaphysics, one inflected by phenomenology—i.e., a humanly-experienced reality that is not susceptible to one and only one complete, coherent, and true account, but comprises diverse and complexly-interrelated pragmatic situations or locations. While it is a conception of reality that lacks the perfection of the Integrationists' ideal—e.g., because pragmatic contexts that we ordinarily keep distinct can on occasion merge and cause us indeterminacy in cognitive processing (as was artificially induced, ingeniously, by Kahneman & Tversky)—this conception might be the best-available means for explaining realistically how humans (and perhaps non-humans, too?) can cope successfully with life.

Bortolotti's critique of Compartmentalism is also hampered by a too-rough conception of cognitive synchronicity or simultaneity. She grants that compartmentalization permits *diachronic* inconsistency between cognitive states that are kept from being brought together in the mind, but she insists that when *synchronic* dissonance threatens, there must inevitably be some sort of reconciling action to prevent or eliminate the inconsistency. Yet she does not slice the stream of consciousness finely enough. Her examples of simultaneous inconsistent cognitive commitments may actually be rapid, unnoticed diachronic shifts as the specific micro-purposes that agents pursue from moment to moment evolve.

If cognition is compartmentalized, then Unification and Truth-directedness cannot be absolute requirements for belief attribution. Governance is the core component, but it too needs qualification. Consider an agent who supposedly believes *p* and believes *not-p*, and who, according to Compartmentalism, reasons at times only with *p*. How, Evnine asks, can we say that *not-p* is a *belief* that Governs mental life if it is not used when it is relevant? Were Governance based on *topic*-relevance, this critique would land: surely whenever *p* is topic-relevant, so too is *not-p*, and to ignore the latter would be to exclude it from Governance, hence belief.

But with the shift to pragmatic situation/task-relevance, *p* and *not-p* are no longer peas in a pod, destined both to be employed in identical cognizing situations. So *not-p*'s Governance credentials need not be impugned by it lying dormant while one utilizes *p*. Beliefs must indeed be involved in governing reasoning, but they need not *always* be involved, in particular where they are topic-relevant but not situation-relevant.

The dividedness of our cognitive systems thus in effect reduces our reliance upon evidence in making, retaining, and revising our cognitive commitments. Yet despite this, routine folk psychologizing, including reasoning about minds that incorporates attributions of belief, productively (if imperfectly) goes on.

And if one's life is not a coherent unity, but is divided among recurring pragmatic task-types that elicit differing personas from one's repertoire, then cognitive assets like beliefs start looking like *special-purpose tools* and not (or not always) general *all-things-considered convictions* or commitments. The agent, rather than identified with a single coherent, principled perspective, comes to look like a more-or-less integrated collectivity of practically-oriented agents whose mental activity is motivated more by pragmatic efficacy, and less by truth, than we are accustomed to think.

CORE REFERENCES

Bendaña, Joseph, & Eric Mandelbaum. 2021. "The Fragmentation of Belief." In C. Borgoni, D. Kindermann, & A. Onofri, eds., *The Fragmented Mind*, pp. 78–107. Oxford University Press.

Bortolotti, Lisa. 2022. "Delusion." *The Stanford Encyclopedia of Philosophy* (Summer 2022 Edition), Edward N. Zalta (ed.) https://plato.stanford.edu/archives/sum2022/entries/delusion/.

Bortolotti, Lisa. 2010. *Delusions and Other Irrational Beliefs*, Oxford University Press.

Evnine, Simon. 2008. *Epistemic Dimensions of Personhood*. Oxford University Press.

Lengbeyer, Lawrence. 2007. "Situated Cognition: The Perspect Model." In David Spurrett, Don Ross, Harold Kincaid, & Lynn Stephens, eds., *Distributed Cognition and the Will: Individual Volition and Social Context*. MIT Press.

On the interaction of aspect and ability in two Hindi/Urdu constructions

Prerna Nadathur
The Ohio State University & New York University

Butt (1997) describes an unexpected **dispositional** reading for certain Hindi/Urdu complex predicates under imperfective marking: while (1) simply indicates that Yusuf habitually drives a car, (2a)—in which *calaa* ('drive') is modified by the light verb *le* ('take')—indicates that Yusuf not only can (has the ability) to drive, but also chooses regularly to exercise this ability. The dispositional reading suggests that *le* introduces modal semantics, but this is difficult to reconcile with its apparent role under perfective marking: *le* in (2b) is often described as an 'aspectual' light verb, contributing a completive interpretation (Singh 1998, a.o.).

(1) *Yusuf gaaṛii calaa-taa (hai).*
Yusuf car drive-IMPF.M (be.PRS)
'Yusuf drives a car.'

(2) a. *Yusuf gaaṛii calaa le-taa (hai).*
Yusuf car drive take-IMPF.M (be.PRS)
'Yusuf will (can and does) drive a car.'

b. *Yusuf-ne gaaṛii calaa li-yaa.*
Yusuf-ERG car drive take-PFV.M.
'Yusuf completed driving a/the car.'

In pursuit of a unified analysis, I compare the pattern of interpretation in (2) to another well-known case where modality is seemingly erased under perfective marking: the **actuality entailments** of ability modals (Bhatt 1999). While imperfectively-marked *sak* ('can') in (3a) describes a 'pure' (potentially-unrealized) ability, its perfective counterpart (3b) is well-paraphrased by the implicative verb *manage* (Karttunen 1971), and requires that the ability-target actually occurs.

(3) a. *Yusuf gaaṛii calaa sak-taa hai (lekin us-ne gaaṛii nahĩĩ calaa-yii).*
Yusuf car drive can-IMPF.M be.PRS (but he-ERG car NEG drive-PFV.F).
'Yusuf can drive a car (but he did not drive a/the car).'

b. *Yusuf gaaṛii calaa sak-aa (#lekin us-ne gaaṛii nahĩĩ calaa-yii).*
Yusuf car drive can-PFV.M (#but he-ERG car NEG drive-PFV.F).
'Yusuf managed to drive a/the car (#but he didn't drive a/the car).'

Drawing on a recent causal analysis of implicative verbs (Nadathur 2023a,b), I suggest that the semantic structure of *manage* offers a path towards unifying the dispositional reading in (2a) with the 'completed' reading in (2b), as well as accounting for the ability/actuality alternation in (3). I propose that what *manage*, *le* and *sak* share is presuppositional reference to a causal background structure in which their subjects must take action to *bring about* (causally ensure) the realization of some event or state associated with the main/embedded verb. *Manage* and *le* both require that this causing action is realized (producing habitual readings in the imperfective and episodic readings under the perfective), but differ in the relationship that the causing action has to the event structure described by the embedded verb. Ability modals differ from both *manage* and *le* in the asserted dimension, establishing only their subjects' capacity to realize the relevant causing action: this produces stative (pure ability) readings under imperfective marking, but triggers an eventivizing operation of *aspectual coercion* (de Swart 1998, Homer 2021, a.o.) when composed with the perfective, resulting in the actualized (implicative) reading in (3b). This approach points to an underlying link between the aspectual properties of event types and the structure of *causal models* (language-independent representations of contextual causal information), which supposedly 'aspectual' light verbs in Hindi/Urdu and other languages are well-positioned to probe.

Korean change of state predicates
Non-culminating readings across scale and causative structures

Paola Fritz-Huechante & Elisabeth Verhoeven
Humboldt-Universität zu Berlin

Research across various languages has demonstrated that a set of change of state (henceforth CoS) predicates in the perfective can yield two distinct readings: (i) a culminating reading, in which the result state specified by the verb's semantic core is achieved, and (ii) a non-culminating (hereafter NC) reading, in which the result state does not obtain (cf. Tatevosov & Ivanov 2009; Demirdache & Martin 2015; Martin & Schäfer 2017, a.o.). Korean causative CoS predicates also exhibit this phenomenon (cf. Beavers & Lee 2020; Fritz-Huechante et al. 2020). Central to the availability of NC readings is the agenthood properties of the external argument, as captured by the Agent Control Hypothesis (ACH, cf. Demirdache & Martin 2015). For instance, under the ACH, negating the result state is possible in the presence of an agentive subject as in (1a), while the result state cannot be defeated in the presence of an inanimate causer as in (1b).

(1) a. Yuli-ka ipwul-ul mal-ly-ess-ta. haciman ipwul-i malu-ci anh-ass-ta.
Yuri-NOM blanket-ACC be.dry-CAUS-PST-DECL but blanket-NOM be.dry-CONN NEG-PST-DECL
'Yuri dried the blanket, but the blanket was not dry.'
b. hayspyeth-i ipwul-ul mal-ly-ess-ta. #haciman ipwul-i malu-ci anh-ass-ta.
sun-NOM blanket-ACC be.dry-CAUS-PST-DECL but blanket-NOM be.dry-CONN NEG-PST-DECL
'The sun dried the blanket, but the blanket was not dry.'

This study extends the analysis of NC readings in Korean by investigating two additional, less-explored factors: scale structure and causative structure. Drawing from scalar classification frameworks (cf. Hay et al. 1999; Kennedy & McNally 2005; Kennedy & Levin 2008), we focus on Korean causative upper-bounded predicates (e.g. *mallita* 'to dry') and lower-bounded predicates (e.g. *ceksita* 'to wet'). Korean data reveal that NC readings occur with the former and not so with the latter. Sentence (1) is an instantiation of an upper-bounded predicate. In order for (1) to hold true, the blanket needs to be maximally dry. This maximum degree corresponds to the standard degree of comparison (or bound) that manifests the property specified by the semantic core of the verb. A NC reading arises in the presence of an agentive subject as in (1a), yielding an interpretation that Yuri acted upon the blanket to dry it without necessarily causing the crucial CoS (i.e. the blanket being completely dry). This is not the case with lower-bounded predicates. Lower-bounded predicates hold true at the presence of a minimal amount of change. In (2), a minimum CoS occurs as soon as the subject referent acts upon the object, hence negating the result state generates a contradiction irrespectively of the type of subject.

(2) Yuri-ka / pipalam-i ipwul-ul ceks-y-essta. #haciman ipwul-i
Yuri-NOM / rainstorm-NOM blanket-ACC be.wet-CAUS-PST-DECL but blanket-NOM
cec-ci anh-ass-ta.
be.wet-CONN NEG-PST-DECL
'Yuri / the rainstorm wetted the blanket, but the blanket was not wet.'

We further observe that Korean causative structures have an impact on NC readings. Morphological causatives, formed by attaching the causative morpheme *-i* (or its allomorphs) to the stem of the stative verb (e.g. *mal-li-ta* 'be.dry-CAUS-DECL'), are able to yield NC readings with upper-bounded predicates in the presence of an agentive subject, as in (1a). In contrast, periphrastic causatives, built by attaching the light verb *hata* 'do' to the stative verb stem (e.g. *malu-key ha-ta* 'be.dry-ADVR do-DECL'), are able to produce NC readings across both lower- and upper-bounded predicates in agentive contexts. This distinction aligns with the morphosyntactic features of *hata* 'do', which can emphasize either an activity or result state reading (Fritz-Huechante et al., 2020; Choe, 2022).

(3) Yuri-ka ipwul-ul malu-key / cec-key hay-ss-ta. haciman ipwul-i
Yuri-NOM blanket-ACC be.dry-ADVR / be.wet-ADVR do-PST-DECL but blanket-NOM
malu-ci / cec-ci anh-ass-ta.
be.dry-CONN / be.wet-CONN NEG-PST-DECL
'Yuri made the blanket dry / wet, but the blanket was not dry / wet.'

Two 1–5 Likert scale experiments were conducted with native Korean speakers, in which participants evaluated sentences canceling the result state with *but*-clauses, as in (1). The experiments were built w.r.t the verbs' availability to form both morphological and periphrastic causative structures, testing the factors: (a)

subject type (agent vs. causer) and (b) scale structure (lower vs. upper). 8 target items (4 lower-bounded and 4 upper-bounded) were selected plus 24 fillers per experiment. Sentences were presented online on IBEX HU. Expectations were: (a) the factor subject type has an impact on the acceptability of a sentence to the extent that NC readings are allowed in the presence of an agent in contrast to causers, and (b) an interaction between scale structure and causative structure to the extent that NC readings are available for upper-bounded predicates in a morphological causative construction, whereas (c) NC readings are available for both lower- and upper-bounded predicates in a periphrastic construction. 32 Korean native speakers participated in the surveys (N = 16 per list). One participant was excluded due to missing data points (N = 31, 15 female, 15 male, 1 no gender. Age: 22–42. M = 29,90). Fig. 1 for morphological causatives shows an interaction of the factors subject type and scale structure to the extent that the acceptability of NC readings is higher in the condition with upper-bounded predicates constructed with an agentive subject. In the case of periphrastic causatives, Fig. 2 shows that this interaction is not present, i.e. both lower- and upper-bounded predicates behave similarly to the extent that the acceptability of NC readings is higher in the presence of an agentive subject than with a causer. The data was fitted with a cumulative link mixed models (cf. Taylor et al. 2023). The random effects structure contained intercepts for items (different verbs) and participants. Results showed a significant effect of subject type (p=.02), scale structure (p=.02), a significant interaction of causative structure^scale structure (p=.01), and a marginal significant interaction of causative structure^scale structure^subject type (p=.05).

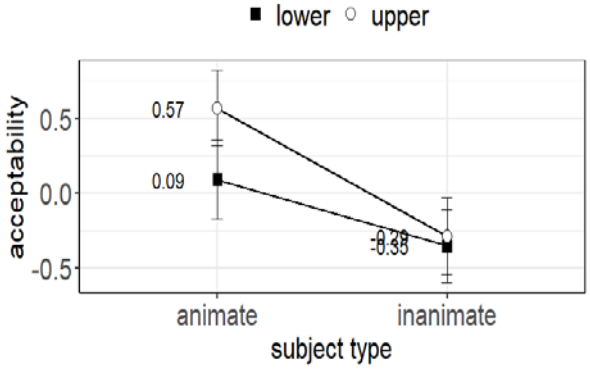

Fig. 1: Acceptability NC readings morphological causatives (95% C.I.)

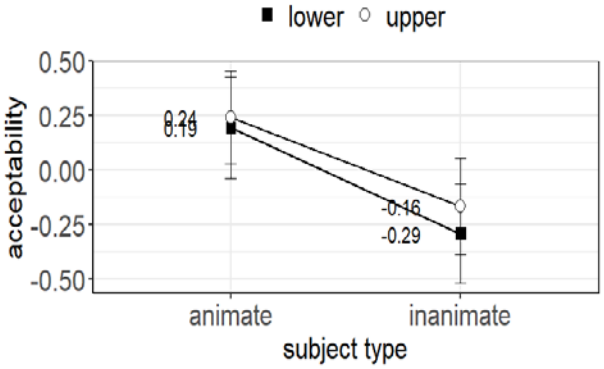

Fig. 2: Acceptability NC readings periphrastic causatives (95% C.I.)

This work underscores the nuanced interplay of agentivity, scale properties, and morphosyntactic causative structures in shaping NC readings, contributing to broader cross-linguistic studies of event semantics and causation.

References: Beavers, J. & J. Lee. 2020. Intentionality, scalar change, and non-culmination in Korean caused change-of-state predicates. *Linguistics* 58(5). 1233–1283. • Choe, J. S. 2022. Structure and meaning of Korean causatives: Evidence from structural priming. *Discourse and Cognition* 29(2). 195–214. • Demirdache, H. & F. Martin. 2015. Agent control over non-culminating events. In E. Barrajón, J. L. Cifuentes & S. Rodríguez (eds.), *Verb classes and aspect*, 185–217. Amsterdam: John Benjamins. • Fritz-Huechante, P., E. Verhoeven & J. Rott. 2020. Agentivity and non-culminating causation in the psych domain: Cross-linguistic evidence from Spanish and Korean. *Glossa* 5(1). 1–35. • Hay, J., C. Kennedy & B. Levin. 1999. Scale structure underlies telicity in degree achievements. In T. Matthews & D. Strolovitch (eds.), *Proceedings of SALT*, 127–144. Ithaca, NY: CLC Publications. • Kennedy, C. & B. Levin. 2008. Measure of change: The adjectival core of degree achievements. In L. McNally & C. Kennedy (eds.), *Adjectives and adverbs: Syntax, semantics and discourse*, 156–182. Oxford: Oxford University Press. • Kennedy, C. & L. McNally. 2005. Scale structure, degree modification, and the semantics of gradable predicates. *Language* 81(2). 345–381. • Martin, F. & F. Schäfer. 2017. Sublexical modality in defeasible causative verbs. In A. Arregui, M. Rivero & A. Salanova (eds.), *Modality across syntactic categories*, 87–108. Oxford: Oxford University Press. • Tatevosov, S. & M. Ivanov. 2009. Event structure of non-culminating accomplishments. In L. Hogeweg, H. de Hoop & A. L. Malchukov (eds.), *Crosslinguistic semantics of tense, aspect, and modality*, 83–130. Amsterdam: John Benjamins. • Taylor, J.E., Rousselet, G.A., Scheepers, C., Sereno, S.C. 2023. Rating norms should be calculated from cumulative link mixed effects models. *Behavior Research Methods* 55, 2175–2196.

Causative invitations and types of causation **AIL5, Jan 2025**

Patricia Cabredo Hofherr (UMR 7023 – SFL, CNRS & U. Paris 8)

Causal constructions express causal relationships between two events. Causal connections can be of different types, including factitive (*make X V*), permissive (*let X V*) among others (see Nadathur & Lauer 2020 for references and discussion). The present study contrasts three causative auxiliaries from Haitian Creole *fè* `make' (1a), *kite* `let, allow' (1b) and *ba(y)* `give' (1c).

(1) a. Manman mwen **fè** Rito fouye twou a.
mother 1SG **make** R. dig hole DET
'My mother made Rito dig the hole.' (Govain 2022:38, ex 4b)

b. M **kite** timoun nan jwe ak chat la.
1SG **let** child DET play with cat DET
'I let the child play with the cat.'

c. Jan **bay** Mari kondwi vwati a.
Jean give Marie drive car DET
'Jean invited Mari to drive the car.' (Glaude 2012:170, ex 21b)

Claim The causative with *bay* `give' expresses a relationship of *causative invitations* contrasting with both factitive *fè* `make' and permissive *kite* `let, allow' causatives. In causative invitations two conditions have to be fulfilled: (i) the embedded predicate has to describe an action under the causee's control and (ii) the action of the causee is part of an interaction with the causer. The causative invitation reading is translated here by *invite*. Of the Haitian causative auxiliaries only *fè* `make' is implicative, *bay* `give' and *kite* `let' causatives imply that the embedded event is caused but do not entail it.

Causee controls the action Both factitive *fè* `make' causatives (1) and *kite* `let' permissives (2) allow caused events that are not under the control of the causee. In contrast, causatives with *bay* `give' are infelicitous with actions that are not under the causees control (3).

(2) a. Pwofesè a **fè** [timoun yo] renmen literati.
professor DET make children DET.PL love literature
'The professor makes the children love literature.'

b. Jann **fè** Mari ri.
J. make M. laugh
'J. made Mari laugh'

(3) a. Jann kite Mari mouri.
J. let M. die.'

b. Jean kite [mi an] tonbe.
Jean let wall DET fall
'J. let the wall fall down (i.e. did nothing to prevent it happening).'

(4) a. #M ap **bay** [sè mwen an] ri.
1SG ASP give sister 1SG DET laugh
#'I invite my sister to laugh.'

b. #Jan **bay** [mi an] tonbe.
Jan give wall DET fall
'J. invites the wall to fall.'

Interaction with the causer *fè* `make' allows a natural force like *van an* `the wind' as a causer (5a), *kite* `let' and *bay* `give' are infelicitous in this type of context (5b/c)

(5) a. Van an **fè** mi an tonbe.
wind DET make wall DETfall
'The wind made the wall come down.' (Govain 2022:40, ex 10b)

b. [Van an] **#kite** [mi an] tonbe.

c. [Van an] **#bay** [mi an] tonbe.
wind DET let /give wall DET fall

#'The wind let the wall fall/invited the wall to fall.'

The causee of BAY-causatives has to be capable of interaction: either animate (6a/b) or an interactive machine (7a/b) (e.g. a computer, ticket machine, automatic door).

(6) a. M **ba** li benyen avan nou sòti
1SG give 3SG take.bath before 1PL go.out
'I invited him/her to take a bath before we went out.'

b. M **bay** [sè mwen an] chwazi [mizik la]
1SG give sister 1SG DET choose music DET
'I invited my sister choose the music.'

(7) a. [Òdinatè a] **ba** w met modpas la.
computer DET give 2SG enter password DET
'The ticket machine invites you to enter your password.'

b. [Machin nan] **ba** w chwazi [kantite tikè w vle a]
machine DET give 2SG choose number tickets 2SG want DET
'The ticket machine invites you to choose the number of tickets you want.'

The examples (6) and (7) are also grammatical with *kite* `let' as a causative. However, the *bay* and *kite* causatives differ in their interpretation. Permissive *kite* `let' has permission readings (allowing the causee to perform an action) or non-interference readings (where the causer does not stop the causee from performing an action) – in either case the causee has the intention of performing the action independently of the causer. In contrast, in *bay*-causatives the causer causes the intention of the causee to perform the action as well as the action itself. With bay-causatives the caused action is performed in a reaction to the causer.

Analysis *Bay*-causatives allow invitation by humans and interactional causation with an interactive machine. This interpretation of *bay* doesn't correspond to an intermediate enabling condition for direct causation (as in Wolff 2003). In *bay*-causatives we have two layers of causation: (i) causation between the causing event and the caused event and (ii) causation of the intention of the causee to carry out the event.

(8) e1 = causing event / e2 = caused event
BAY-causative CAUSE(e1,e2) & CAUSE(e1, intention(CAUSEE,e2))
The causing event e1 causes the event e2 & the intention of the causee to carry out e2

The three types of causatives differ with respect to the second layer of causation. The double causation of an event and an intention in *bay*-causatives contrasts with factitive and permissive causatives. Factitive causatives like *fè* `make' do not impose any conditions on the intentions of the causee wrt to the event e2, while in permissive readings of causatives like *kite* the intention of the causee to carry out e2 (intention(CAUSEE,e2)) pre-exists the permission relation and in non-interference readings the causer refrains from initiating an event e1 that could stop e2 whether it is intentional (*play with the cat* in 1b) or non-intentional (*die* 3a). Future work has to show whether causees can be inanimate in examples like "The phone GIVES the computer connect to the hotspot" and whether *intention* as postulated here has to be generalised to also cover a *reaction* from an interactional machine. This analysis including a separate causation of intention is supported by the fact that *invite* can grammaticalize as an auxiliary or an adjective expressing causation of intention:

(9) a. Kunst im Parlament lädt zum Nachdenken über Demokratie ein. (German)
art in-the parliament invites to-the reflect.INF about democracy PRT
`Art in parliament invites reflecting about democracy.'

b. an inviting prospect / The room is very inviting
inviting: attractive in a way that makes you want to do something, go somewhere, be near someone, etc. https://www.britannica.com/dictionary/inviting

References: **Glaude 2012**. *Aspects de la syntaxe de l'haïtien*. Editions Anibwé. **Govain 2022**. *La question linguistique haïtienne : histoire, usages et description*. Mémoire HDR, U. Paris 8. **Nadathur & Lauer 2020** Causal necessity, causal sufficiency, and the implications of causative verbs", *Glossa* 5(1): 49. **Wolff 2003**. Direct causation in the linguistic coding and individuation of causal events. *Cognition* 88 : 1–48.

Intentionality, control, and the semantic notion of agent

Tatjana Ilic

(Keio University, Japan; National Defense Academy, Japan)

The notion of agency is a matter of controversy among linguists. While most theories assume that agent is a discrete and semantically invariable category, some linguists have argued for a more granular approach to agency based on the idea of the agent prototype (DeLancey, 1984; 1990; Dowty, 1991). On this view certain agents are more prototypical than others depending on the number of agent properties they involve. This difference is often reflected in morphosyntax, as less prototypical agents commonly receive a different morphosyntactic coding compared to the prototypical ones. Intentionality, causality, and control over the predicated event are some of the semantic properties commonly associated with the agent. A prototypical agent causes an event or a change of state in another participant; this causality is intentional and proceeds towards the intended outcome; the agent has control over the entire course of the event – the agent can see to that the event happens and is fully responsible for the intended outcome.

In what follows, I present evidence from dative reflexives in Slavic languages which supports the prototype approach to the notion of agent and furthermore implicates control over the event and its outcome as the fundamental property of agency. Examples below contrast prototypical nominative agent with the non-prototypical dative agent in dative reflexives in Slavic (Ivanova, 2014).

1) Ja rabotaju. *Russian*
   I_{NOM} work.$_{PRES.1.SG}$
   'I'm working.'

2) Segodnja mne **otlično** raboaet -sja.
   today I_{DAT} splendidly work.$_{PRS.3.SG}$ $_{REFL}$
   Lit: 'It works itself to me splendidly today.'
   'Today my working goes splendidly.'

3) Az rabotja. *Bulgarian*
   I_{NOM} work.$_{PRES.1.SG}$
   'I'm working.'

4) Raboti mi se.
   work.$_{PRS.3.SG}$ I_{DAT} $_{REFL}$
   Lit: 'It works itself to me.'
   'I feel like working.'

Dative reflexive construction is widely spread across all Slavic languages with the same basic morphosyntactic properties – the dative agent, the reflexive (anticausative) morphology, and a nonagreeing verb – and the same shared basic meaning that the noted eventuality is due to some factors which are beyond control of the dative agent (e.g., Rivero and Arregui, 2012 and the references therein; Ilic, 2013; 2014; Ivanova, 2014 among others). However, the specific interpretations differ. Dative reflexives in West Slavic (Polish, Czech, Slovak) and Russian represent subjective evaluations of actualized events with the meaning of "inexplicable ability to do something well" (Ivanova, 2014). "The agent perceives his own action as proceeding well (or not well) for reasons independent of him" (Wierzbicka, 1988). "The goal is achieved through luck and external conditions" (Dąbrowska, 1997), and the success is commonly attributed to "the environment in which the action takes place" (Wierzbicka,

1988). Examples (5) from Czech (Franks, 1995) and (6) from Russian (Ivanova, 2014) illustrate this point. Because they are subjective evaluations of the noted events, these constructions occur with an obligatory adverbial modifier that qualifies the event as proceeding excellently (5), or with difficulties (6) and (7) – as perceived by the speaker (the last example is from Polish, as cited in Franks, 1995).

5) Sestře se **<u>tam</u>** pracuje ***výborně***. *Czech*
sister$_{DAT}$ $_{REFL}$ there work.PRS.3.SG excellently.
'(My) sister is working excellently there.'

6) Mne **<u>pri zakrytoj fortočke</u>** ***ploxo*** spit-sja. *Russian*
I$_{DAT}$ with closed window badly sleep-$_{REFL}$
'My sleeping goes badly with the window closed.'

7) Ten artyku-ł **<u>nie</u>** pisze mi się dobrze. *Polish*
this.NOM article.NOM NEG write.3.SG.PRS me$_{DAT}$ $_{REFL}$ well
'I just can't write this article well.'

Dative reflexives in South Slavic, on the other hand, have the meaning of dispositions towards potential (non-actualized) events. They speak of uncontrollable "urges …and impulses of the dative [referent], not actions in the 'real world'… They report "dispositions that need not materialize" (Rivero and Arregui, 2012) and are "inexplicable" in their nature (Ivanova, 2014), as in (8) for Slovenian.

(8) Janezu se spi. *Slovenian*
Janez$_{DAT}$ $_{REFL}$ sleep.PRS.3.SG
'Janez is sleepy/Janez feels like sleeping.'

However, regardless of the interpretational differences illustrated above, eventualities expressed by dative reflexives in Slavic languages are conceptualized across the board as caused by some factors which are beyond control of the dative agent. Even when the event is interpreted as actualized and caused intentionally, as it is the case in West Slavic and Russian examples, the dative agent is not seen as the ultimate causal factor responsible for the event actualization and its outcome, but rather as a mediating cause. I provide evidence for the existence of a presupposed causal relation between the dative agents and the primary causal factors that facilitate (or prevent) event actualization in West Slavic and Russian and are responsible for the dative agent's dispositions.

References:

DeLancey, Scott. 1984. Notes on Agentivity and Causation, *Studies in Language* 8, 181 213.

Dowty, David. 1991. Thematic Proto-Roles and Argument Selection,' *Language* 67: 547-619.

Franks, Steven. 1995. *Parameters of Slavic Morphosyntax*. New York: Oxford University Press.

Ilic, Tatjana. 2013. *Modality and Causation in Serbian Dative Anticausatives: a Crosslinguistic Perspective.* Ph.D. Dissertation.

Ilic, Tatjana. 2014. Modality and Causation: Two Sides of the Same Coin. In B. Copley and F. Martin (eds.), *Causation in Grammatical Structures*, 152-175. Oxford: Oxford University Press.

Ivanova, Elena. 2014. Reflexive-with-Dative Construction in Russian and Bulgarian. International Conference on Education, Art and Intercultural Communication (ICELAIC 2014).

Rivero, María Luisa, and Ana Arregui. 2012. Building involuntary states in Slavic. In *Telicity, change, and state: A cross-categorial view of event structure*, eds. Violeta Demonte and Louise McNally, 300–332. New York: Oxford University Press.

Agents are anti-telic

Mingjiang Chen, University of Connecticut
mingjiang.chen@uconn.edu

There are roughly three ways of defining Agents in the literature. The ***first*** one treats all external arguments as bearing a semantically underspecified θ-role, defined as a participant in a *causing* event. *Volitionality* is used to further classify this underspecified θ-role into more specific types. Agents are defined as volitional participants in causing events (Levin and Rappaport Hovav 2005: 70).

The ***second*** one suggests that Agents are simply participants in *causing* events (while Causers are *causing* events themselves) (e.g., Pylkkänen 2002). *Volitionality* has been argued as a property common to all Agents. Therefore, it is an *identifying* rather than a *defining* property of Agents.

The ***third*** one has never been explicitly formulated in the literature but is implicit in many works (e.g., Folli & Harley 2005). It suggests that Agents are participants in *act* events (while Causers are participants in *causing* events). This is illustrated in (1a), which represents the neo-Davidsonian adaptation of Rappaport Hovav and Levin's (1998) lexical semantic template for activities in (1b).

(1) a. λxλe.[Agent(e, x) & ACT(e) …] b. [x ACT$_{<MANNER>}$]

This definition predicts that Agents may correlate with some identifying properties of act events. This prediction is borne out. I propose that one such identifying property of act events is *anti-telicity* (the impossibility of being licensed in telic predicates). This implies that the third definition should be preferred. I argue that Agents, as arguments licensed by act events, are also *anti-telic* and cannot appear with telic predicates. After reviewing arguments for using volitionality to identify Agents, I discuss its problems and argue that *anti-telicity* is a better identifying property of Agents than *volitionality*.

Agent/Telicity Connections. The first two definitions of Agents both use volitionality to identify Agents. This is motivated by the observation that the syntactic behavior of certain predicates varies depending on the volitionality of their external arguments. According to Demirdache and Martin's (2015) *Agent Control Hypothesis* (ACH), for a class of **telic** predicates that select Agents or Causers, they have a *zero change-of-state* reading only if their external arguments are **Agents**. In the Mandarin data (2a–b), the predicate *zhe na gen shuzhi* 'snap that branch' is telic (an achievement), as indicated by the adverbial 'in an instant,' and encodes the result that "the branch broke." (2c) shows that when the subject is *volitional* (*Xiaohong*), the predicate can be *atelic*, as indicated by the adverbial "for a long time," and the *zero change-of-state* reading is available. However, when the subject is *non-volitional* ("the big wind"), the predicate cannot be *atelic*, and only the *change-of-state* reading is available. (2) thus supports the ACH *if* Agents are identified with volitionality.

(2) a. *Xiaohong* yizhayandegongfu jiu zhe-le na gen shuzhi. *telic*
Xiaohong in.an.instant at.once snap-PFV that CL branch
'Xiaohong snapped that branch in an instant.'

b. *dafeng* yizhayandegongfu jiu zhe-le na gen shuzhi, *telic*
big.wind in.an.instant at.once snap-PFV that CL branch
'The big wind snapped that branch in an instant.'

c. *Xiaohong/*dafeng* zhe-le na gen shuzhi haojiu *atelic*
Xiaohong/*big.wind snap- PFV that CL branch very.long
Literally: 'Xiaohong/The big wind snapped that branch for a long time.'

Problems. Identifying Agents as *volitional* external arguments faces a few problems. ***First***, Folli & Harley (2008) point out that although the sole arguments of unergatives are often analyzed as Agents, they are not always volitional (e.g., unergatives involving involuntary bodily processes such as *sneeze*; also see Levin and Rappaport Hovav's (1995) sound emission verbs such as *whistle*). ***Second***, if Agents are all volitional, the sole arguments of unergatives should behave differently, especially with respect to *telicity*, depending on their *volitionality*. However, both *volitional* and *non-volitional* arguments of unergatives can only appear in *atelic* predicates. In Italian, unergatives select the auxiliary 'have', while

unaccusatives select 'be'. In (3a), the verb 'roll' behaves like an unergative, and since the PP 'under the table' can only be interpreted as locative in this case, and since only the *for X time* adverb is allowed, the predicate must be atelic. In (3b), the same verb behaves like an unaccusative, and since it takes a directed motion PP 'into the goal' and is only compatible with the *in X time* adverb, it must be telic.

(3)	a.	La	palla	ha	rotolato	**sotto**	**il**	**tavolo**	per/*in	un	secondo.	*atelic*
		the	ball	has	rolled	**under**	**the**	**table**	for/*in	a	second	
	b.	La	palla	è	rotolata	**nella**	**rete**		*per/in	un	secondo.	*telic*
		the	ball	is	rolled	**into.the**	**goal**		*for/in	a	second	

(4) patterns exactly with (3). While the subjects in (3) are non-volitional, the ones in (4) are volitional.

(4)	a.	Luisa	ha	corso	**nel**	**parco**	per/*in	un'ora.	*atelic*
		Luisa	has	run	**in.the**	**park**	for/*in	an hour	
	b.	Luisa	è	corsa	**a**	**casa**	*per/in	un'ora.	*telic*
		Luisa	is	run	**to**	**house**	*for/in	an hour	

An Alternative. In (2), we see two factors at play: the **volitionality** of an external argument and the **telicity** of a predicate. We can follow the first two definitions of Agents in identifying *volitional* external arguments as Agents and explore their connections to telicity, which has been argued to be problematic. Or we can identify Agents with *anti-telicity* (the impossibility of being licensed in telic predicates) and explore their connections to volitionality.

Now a caveat on the definition of anti-telicity is in order. Strictly speaking, *anti-telic* arguments only appear in *atelic* predicates, while *non-anti-telic* arguments can appear in *atelic* or *telic* predicates. However, (2–4) suggests that the relevant concept of anti-telicity is a special case in the sense that *non-anti-telic* arguments should only appear in *telic* predicates, with *anti-telic* arguments still only appearing in *atelic* predicates.

The table below illustrates the difference between the two ways of identifying Agents. [±vol], [±a-t], and [±t] stand for volitionality, anti-telicity, and telicity. Notably, volitionality and anti-telicity are doubly dissociated, with an external argument capable of being volitional without necessarily being anti-telic (see (2a) and (4b)), and vice versa (see (3a)). This implies that volitionality and anti-telicity are independent of each other, and the latter can by no means be reduced to the former.

	[–vol], [–a-t] (2b), (3b)	[+vol], [+a-t] (2c), (4a)	[+vol], [–a-t] (2a), (4b)	[–vol], [+a-t] (3a)
Ag:[+vol]	Non-Agent	Agent	**Agent**	**Non-Agent**
Ag:[+a-t]	Non-Agent	Agent	**Non-Agent**	**Agent**

Consider now the relation between Agents and volitionality under the new way of identifying Agents. (2) demonstrates that volitionality is *necessary* but *insufficient* for identifying an argument as an Agent. The ACH can then be reformulated as (5). However, the unergative pattern illustrated in (3) and (4) implies that there is no significant correlation between Agents and volitionality. (5) is then subject to lexical variation, as only certain verbs require their Agent subjects to be volitional.

(5) Only volitional arguments are allowed to be Agents.

Therefore, the pattern in (2) is also captured if Agents are identified with anti-telicity. Crucially, under this view, the problems with unergatives disappear. Both volitional and non-volitional arguments of unergatives are now Agents, and the connections between Agents and telicity are also explained.

Implications. I argue that Agents are anti-telic because they are participants in *act* events (following the third definition of Agents) and *act* events are anti-telic. In other words, Agents inherit anti-telicity from *act* events. Since anti-telicity is aspectual in nature, it would be hard to justify why Agents have this property under the first two definitions of Agents, which do not aspectually distinguish Agents from other θ-roles. This indicates that the third definition of Agents is more promising.

Conclusion. Agents are often identified with volitionality. However, this approach does not account for why volitional and non-volitional arguments of unergatives behave exactly the same and are both

typically analyzed as Agents. In contrast, identifying Agents with anti-telicity avoids this issue and readily accommodates the observed volitionality restriction with some verbs.

References

Demirdache, Hamida, and Fabienne Martin. 2015. Agent control over non-culminating events. In *Verbal classes and aspects*, eds. Elisa Barrajón López, José Luis Cifuentes Honrubia, and Susana Rodríguez Rosique, 185–217. Amsterdam: Benjamins.

Folli, Raffaella, and Heidi Harley. 2005. Consuming results in Italian and English: Flavors of v. In *Aspectual inquiries*, eds. Paula Kempchinsky and Roumyana Slabakova, 95–120. Dordrecht: Springer.

Folli, Raffaella, and Heidi Harley. 2008. Teleology and animacy in external arguments. *Lingua* 118: 190–202.

Levin, Beth, and Malka Rappaport Hovav. 1995. *Unaccusativity: At the syntax-lexical semantics interface*. Cambridge: MIT Press.

Levin, Beth, and Malka Rappaport Hovav. 2005. *Argument realization*. Cambridge: Cambridge University Press.

Pylkkänen, Liina. 2002. Introducing arguments. Ph.D. dissertation, MIT.

Rappaport Hovav, Malka, and Beth Levin. 1998. Building verb meanings. In *The projection of arguments: Lexical and compositional factors*, eds. Miriam Butt and Willi Geuder, 97–134. Stanford: CSLI Publications.

(Non-)agentivity and Subjecthood

Beth Levin
Stanford University

Studies of the semantic underpinnings of transitivity associate subjecthood with agentivity, and indeed the first examples of transitive verbs that come to mind have agentive subjects (e.g., *The toddler broke the toy; The gardener swept the path; The cook pounded the meat*). However, transitive verbs do not always have agentive subjects, and the question that arises is whether there are any restrictions on their distribution. There seems to be a consensus that causative transitive verbs, such as those denoting a change of state, allow various ontological types of subjects, but there is controversy about whether transitive manner verbs allow such subjects. There are claims that such verbs are found with non-agentive subjects, including natural forces and other inanimate causes, only in the context of an overtly expressed result, i.e. in an explicit bi-eventive (or causative) structure, and not when they describe simple events (Folli & Harley 2005, 2008, Martin & Schäfer 2014:233, Schäfer 2012). Folli & Harley (2005:96) exemplify this claim with the pair *?*The wind carved the beach* vs. *The wind carved away the beach*.

This so-called "result restriction" is motivated by a small number examples, and its empirical basis has been questioned (Bruening 2010; Mateu & Rigau 2010:264–265; Rappaport Hovav 2016:472). This talk presents English corpus data from verbs of contact that show that the empirical picture is more complex: transitive manner verbs with non-agentive subjects need not co-occur with result phrases (e.g., *Dust particles battered the spacecraft; He felt the sand scrape his skin*). These data show that the restriction is at best a tendency, and in the talk I propose an alternate approach to the range of observed non-agentive subjects. Drawing on Rappaport Hovav & Levin (2024), I introduce conceptual representations of the core meanings of the relevant verbs together with principles of argument realization to account for their argument realization options, including the (non-)agentivity of their subjects.

References

Bruening, B. (2010) "Subject Restrictions, Particle Verbs, and Consumption Verbs", Linguistics Commentary, blog. http://lingcomm.blogspot.com/2010/05/subject-restrictions-particle-verbs-and.html)

Folli, R. & H. Harley (2005) "Consuming Results in Italian and English: Flavors of V", in P. Kempchinsky and R. Slabakova, eds., *Aspectual Inquiries*, Springer, Dordrecht, 95-120.

Folli, R. & H. Harley (2008) "Teleology and Animacy in External Arguments", *Lingua* 118, 190-202.

Martin, F. & F. Schäfer (2014) "Causation at the Syntax-Semantics Interface", in B. Copley and F. Martin, eds., *Causation in Grammatical Structures*, Oxford University Press, 209-244.

Mateu, J. & G. Rigau (2010) "Verb-Particle Constructions in Romance: A Lexical-Syntactic Account", *Probus* 22, 241-269.

Rappaport Hovav, M. (2016) "Review of A. Alexiadou, E. Anagnostopoulou, and F. Schäfer: *External Arguments in Transitivity Alternations*", *Language* 92, 471-474.

Rappaport Hovav, M. & B. Levin (2024) "Variable Agentivity: Polysemy or Underspecification?", *Glossa* 9, Article 10774, 1-39.

Schäfer, F. (2012) "Two Types of External Argument Licensing — The Case of Causers", *Studia Linguistica* 66, 128-180.

Learnable Control

Elena Garadja

The main question that is the subject of this paper is: What is the genus of human action? Formally, an answer to that question might point to paradigm cases of human action by reference to which the broader genus should be understood, and/or provide a list of necessary and sufficient conditions for something to count as a human action. This paper analyzes one dominant way of identifying the genus of human action in contemporary action theory literature (which can for ease of reference be called the *Consciousness* view), provides reasons for rejecting it and offers an alternative.

According to the *Consciousness* view paradigmatic species of human action are intentional, conscious actions, i.e. those to which the special sense of G.E.M. Anscombe's famous question 'Why?' has application. Actions are paradigmatically understood as events that are materially connected to the agent and which the agent is necessarily conscious of as her actions. The *Consciousness* view can acknowledge that there are other action-like events, which lack the element of consciousness, and understand them as deficient or derivatively intelligible phenomena - as pseudo-actions (i.e. as something which has some of the features of the paradigm case, but lacks some other crucial features.).

The *Consciousness* view is problematic for two reasons. First, the criteria espoused by the *Consciousness* view are insufficient. An event can be causally downstream of an agent and the agent could have the conscious belief that this event is her action and yet the connection between the agent and the event may be too accidental, too much of a matter of luck in order for others to genuinely attribute the action to her. Consider the following example: a (flashy) gambler exhales onto the dice as he throws them, disturbing their trajectory somewhat. If he could only answer at the moment he does this he would say "I am blowing the dice double-six-wards". Imagine further that the dice do land on a double-six, and that the contraction of both his diaphragm and cheeks, the shape of his tongue, the direction, and force of the puff of air thereby propelled towards the dice all played a causal role in their flight path, which in turn played a role in their tumbling, which settled on a double-six. In other words, blowing on the dice was materially efficacious in the causal history of two dice landing double six. My contention is that "blowing the dice double-six-wards" is not in fact an action that the gambler is performing, though we might call it a pseudo-action.

Second, consciousness is not necessary for it being an action. Not only do we frequently take one another to be performing actions unconsciously, some actions are arguably performed best when agents lack consciousness of them. Consider the following example of self-deception: A is asking B for forgiveness, thinking to herself "I am asking for forgiveness. I am trying to reestablish a good relationship between us. I am a good person." But A's request for forgiveness shows some defects: A is low on detail about what she has done wrong, she is quick to come up with vague excuses that minimize her wrongdoing, she is impatient about listening to B's side of the story etc. Is she really asking for forgiveness or merely saying the words, while trying to escape an uneasy situation? In some situations we might plausibly think it's the latter, even if the possibility is not on A's mind. This is a banal example of a case where one pursues something shameful and so masks the action from

oneself in order to be able to pursue it. According to the *Consciousness* view, the action just described is not a real action, it is at best a pseudo-action, whereas I would contend that this is a proper action.

According to the alternative view we should understand actions as exercises of *learnable control*. I propose to view action as a unity of attitudinal, material and modal aspects: an embodied imperative attitude with modal properties. On the attitudinal side, to act is to have a commitment to some standard of success, which can be incorporated into the form of a practical argument. This attitude could have been consciously or unconsciously computed by the agent. On the material side, there is a material process causally connected to the agent. On the modal side, certain counterfactuals are true of the agent. These come into view when we consider that an action is an exercise of learnable control and as such has the following two properties:

(i) *Degree of control*: The degree of control can be understood as a ratio of successes over attempts within a set of parameters. In other words, attributing an action (e.g. B-ing) to an agent we judge that in a number of counterfactual scenarios she will take (potentially corrective) measures to successfully accomplish B.

(ii) Learnability: The degree of control must be improvable through a category of actions called learning. I distinguish between two main kinds of learning, the potential for which is necessary for control: first, brute-force, trial and error, learning through repetition (a human kind of machine-learning); and second, learning by means of symbolic representation.

This view can accommodate unconscious actions, such as the self-deception case, described above. Self-deception is something that human beings can learn to get better at. For the most part this skill is acquired in a non-symbolically mediated way, but rather through brute-force trial and error, as well as by osmosis. This view also explains why "blowing the dice double-six-wards" is not a proper action. Throwing a six is not an action, since you cannot increase your control over throwing a six above 16.7%. It can however be understood as a pseudo-action.

The model of action as learnable control is something that I think is implicit in common competent talk about action, though it has not so far been made explicit as a theory. In my view the *Consciousness* model is a useful idealization, which works well enough most of the time and may even be practically indispensable in certain contexts. Action interpretation, especially when it involves suspicious analysis requires time, effort, emotional resources etc. It may simply be unconducive for a variety of purposes. Nevertheless, the *Consciousness* model is severely limited. The model of action as learnable control helps make explicit the *non-accidental* nature of human action and can facilitate more intelligent discourse about a greater variety of actions.

The speaker's attitude to her own words

Olivia Sultanescu
Concordia University

In the second part of *Philosophical Investigations*, Ludwig Wittgenstein writes: "My attitude to my own words is wholly different from that of others… If I listened to the words issuing from my mouth, then I could say that someone else was speaking out of it" (§§103-104). Unlike my interlocutors, I do not listen to the words issuing from my mouth. It seems that I can nevertheless hear those words, but the way in which I hear them is bound to be different from the way in which others hear them. Why does this make the attitude to my own words *wholly different* from that of others? In this talk, I try to address this question.