# From roots to trees: Thematic roles and the lexicon-syntax interface

Dissertation prospectus

Eva Neu, UMass Amherst
eneu@umass.edu, eva-neu.github.io
June 10, 2025

**Abstract**

Dowty (1991) proposed that thematic roles (θ-roles) are defined by sets of abstract entailments that determine the syntactic position of an argument. In my dissertation, I argue a seemingly simple point: non-thematic entailments do not determine the syntactic position of an argument. Surveying a wide range of cross-linguistic data, I show that the verb-specific entailments that appear to establish the identity of an argument across alternations do not consistently track this argument's thematic entailments nor its syntactic position. This finding is at odds with the view that the denotation of a verb contains a set of discrete arguments defined by their verb-specific entailments and that for a given verb, each θ-role consistently describes one (or no) member of this set. Instead, I argue that θ-roles describe arguments in the context of a complete event description. To make this view more concrete, I set out to disentangle the semantic contributions that lexical roots and θ-roles, respectively, make to the verbal event.

## 1. Introduction

### *1.1. Overview*

My dissertation is about θ-roles. I begin by outlining what I believe to be some widely held assumptions about θ-roles. I then briefly describe a set of data that challenge these assumptions and that motivate this dissertation. I end this section by making specific the question I set out to solve.

The function of θ-roles is to relate arguments to event descriptions. In (1), Brutus is related to the stabbing event by an agent θ-role, and Caesar by a patient θ-role:

(1) $[\![\text{Brutus stabbed Caesar}]\!] = \lambda e.stab(e) \wedge Agent(e)(Brutus) \wedge Patient(e)(Caesar)$

For concreteness' sake, I assume that θ-roles realize independent syntactic heads, as depicted very schematically in (2). I simply label the relevant heads Agent and Patient, but we might identify them with Voice and *v*. The denotations of *stab*, Agent and Patient are given in (3).

(2)

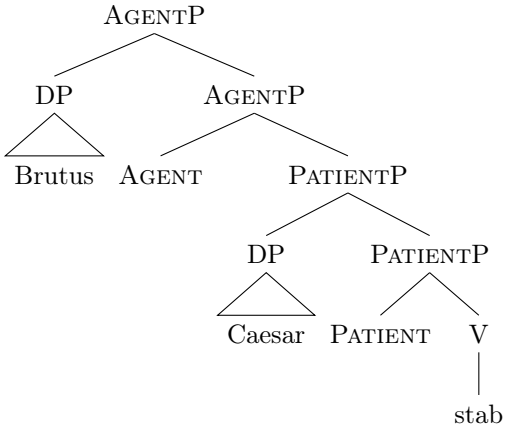

(3) a. $[\![\text{stab}]\!] = \lambda e.stab(e)$
b. $[\![\textsc{Agent}]\!] = \lambda P.\lambda x.\lambda e.P(e) \wedge Agent(e)(x)$
c. $[\![\textsc{Patient}]\!] = \lambda R.\lambda y.\lambda e.R(e) \wedge Patient(e)(y)$

Ever since they made their appearance in Fillmore (1968) and Gruber (1965), θ-roles have been a crucial ingredient of a broad range of theoretical approaches to argument structure. Their most important responsibility is to determine in which syntactic position each argument of a verb is merged. For instance, there is wide agreement that in a transitive verb with an agent and a patient argument, the agent is merged as an external argument surfacing as the subject and the patient as an internal argument surfacing as the object. More controversially, θ-roles have also been argued to play a role for other grammatical phenomena such as binding (Everaert and Anagnostopoulou, 1997; Wilkins, 1988; Williams, 1987), control (Nishigauchi, 1984; Williams, 1987) and noun incorporation (Mithun, 1984). In general, the function of θ-roles as a theoretical construct is to explain why arguments of different verbal roots come to have the same grammatical properties – why, for instance, the subject arguments of the verbs *stab*, *abandon* and *devour* behave identically for a wide range of diagnostics. Syntactic and semantic composition does not operate on, and arguably cannot even 'see,' the concrete interpretation of specific arguments. There is no designated syntactic position for stabbers. What syntax and compositional semantics operate on are abstract clusters of arguments, and these clusters are what are commonly termed θ-roles.

Let us zoom in on the question of what exactly it means for θ-roles to relate arguments and event descriptions. A common intuition, I believe, is that θ-roles identify their argument with a particular kind of participant in the event described. For instance, a stabbing event involves two kinds of participants, a stabber and a stabbee. In *Brutus stabbed Caesar*, the agent role identifies Brutus with the stabber and the patient role identifies Caesar with the stabbee. In general, whenever combining with a stabbing event, Agent will identify its specifier argument with the stabber and Patient its specifier arguments with the stabbee. When combining with a devouring event, Agent and Patient will pick out the person doing the devouring and the thing being devoured, respectively, and so forth. Under this view, to relate arguments and event descriptions means to specify what kind of concrete activity the argument performs in the event described.

The point that this dissertation will try to make is that there is something fundamentally flawed about the intuition outlined in the previous paragraph. In order to make this point, I first need to turn this intuition into a formal theory that can be properly refuted. This formal theory, to my knowledge, has never been explicitly proposed or espoused by anyone, but I do believe that the idea underlying it is present in much work on the syntax-semantic interface, past and present. When discussing some concrete argument-structural constructions in later parts of this prospectus (and the eventual dissertation), I will engage with some competing analyses which I argue are motivated precisely by this idea.

Our formalization will take inspiration from Dowty (1991). Dowty argued that there are only two θ-roles, agent and patient, each of which comes with a set of characteristic properties. Agents are associated, among other things, with sentience and volition, and patients with undergoing a change of state and being causally affected. One key innovation of Dowty's work is to formalize these properties as entailments. For instance, *Brutus stabbed Caesar* entails that Brutus was sentient and acted volitionally. The other innovation is that θ-roles have a prototype structure: not all instantiations of agents and patients need to come with all the entailments associated with these roles. For instance, *Brutus hated Caesar* does not entail that Brutus acted volitionally, but only that he was sentient, which still makes Brutus an agent in Dowty's sense. If an argument has entailments associated with both agent- and patient-hood, it will receive the θ-role whose prototype it matches most closely. This θ-role then determines the syntactic position that the argument occupies.

Dowty states that 'a role type like "Agent" is defined semantically as whatever entailments of verbs about NP referents are shared by the verbal argument-positions that we label with the term "Agent" (and excludes whatever is entailed for those arguments that differs from one verb to the next)' (1991:552). The parenthesis is key for our purposes. Arguments come with a range of other entailments that have nothing to do with proto-properties but that are due to the lexical semantics of the verb. *Brutus stabbed Caesar*, for instance, entails that Brutus wields a sharp weapon like a knife or a dagger, an entailment that is shared by hardly any agent argument of other verbs. Borrowing a term from Dowty (1989), I will define the entailments associated with an argument of a verb (*qua* argument of the verb) at the exclusion of thematic properties as this argument's *individual role*. These non-thematic entailments can then be used to define the hazy notion of

'kind of participant' – such as a stabber or a stabbee – that we saw earlier. The individual role of a stabber, for instance, is defined by the entailment of wielding a sharp weapon in a particular way. In the following, I will often refer to individual roles by handy nominalizations such as stabber and stabbee, but this should be understood as a shorthand for the relevant sets of entailments; the nominalizations themselves have no deeper import.

To recap, θ-roles are defined by sets of entailments that are shared between arguments of different verbs and that determine an argument's syntactic position. Examples include sentience, volition, being causally affected, etc. Individual roles are defined by sets of entailments that are specific to concrete verbs and that are part of the lexical semantics of the verb. Examples include wielding a sharp weapon, being penetrated by a sharp weapon, etc. Let us get back to the idea that the way in which θ-roles relate arguments and event descriptions is by identifying their argument with a particular kind of participant in the event; for instance, that the agent role identifies Brutus with a stabber. We can now reformulate this intuition as follows: the θ-role selects a concrete individual role involved in the event, as defined by non-thematic entailments, and associates its argument with this individual role.

The question this raises is what determines the mapping relationship between thematic roles and individual roles. Clearly, this mapping is restricted in some way. For example, we never find instances of *stab* in which the stabber argument exhibits properties associated with patients – for instance, surfacing in object position –, and vice versa for stabbees. How come that for the verb *stab*, the non-thematic entailment of wielding a sharp weapon is always associated with the agent θ-role? I argue that we do not actually have a theory that would answer this question. What we do have is an intuition that takes the place of a theory, which is that stabbing is inherently a highly agentive thing to do. That is, the θ-role of an argument is somehow directly determined by its individual role: a stabber is by necessity always an agent, and the agent of a stabbing is always a stabber, or in other words, being a stabber is simply what it means to be the agent of a stabbing event. Under this view, then, the θ-roles of a given verb and its individual roles stand in a one-to-one mapping relationship.

So far, I have sketched out a set of common assumptions about θ-roles. I have focused on the question of how concretely θ-roles relate arguments to event descriptions by introducing the notion of individual roles and describing the mapping relationship between θ-roles and individual roles. I believe that assumptions along these lines underlie, explicitly or implicitly, much research on argument structure. What my dissertation will be concerned with is a set of data that give us a more complex picture of the mapping relationship between θ-roles and individual roles. In the following, I briefly summarize two of the six case studies I am planning to work on. Section 2 will discuss all six case studies in more detail.

The first case study concerns variable behavior of intransitive verbs. It is well-known that many verbs can be used both in an unergative structure, in which the subject is base-generated as an external argument, and in an unaccusative structure, with the subject starting out as an internal argument (Borer, 2005; Dowty, 1991; Perlmutter and Postal, 1984; Sorace, 2000, 2004, 2011, a.m.o.). The basic phenomenon is demonstrated in (4) for Hindi-Urdu. In this language, as in many others, reduced relatives have long been observed to only be licensed with unaccusative verbs (e.g., Bhatt and Embick, 2017). In Hindi, the verb *fly* passes this diagnostic only when it takes an inanimate argument such as *kite* (9b), but not with animate arguments such as *bird* (9a):

(4) a. *uṛ-ii (huu-ii) ciṛyaa
fly-PFV.FSG be-PFV.FSG bird.FSG
Intended: 'the flown bird'

b. uṛ-ii (huu-ii) patang
fly-PFV.FSG be-PFV.FSG kite.FSG
'the flown kite' (Ahmed, 2010:8f.)

This contrast indicates that the subject of (9a) is syntactically an external argument and the subject of (9b) an internal argument. Crucially, the arguments also differ with respect to their semantic proto-properties. In (9a), the bird's flying is an effortful, intentional activity, whereas in (9b), the kite's flying is more akin to a passive floating. That is, the first argument has properties associated with proto-agents, the second one properties associated with proto-patients.

I argue that these two facts taken together – the syntactic difference and the difference in proto-properties – warrant the conclusion that the argument is assigned a different θ-role in both sentences. In general, the external argument of unergatives is assumed to be assigned an agent θ-role, and the internal argument of

unaccusatives a patient θ-role. One might of course argue that variable unaccusativity simply deviates from this mapping in that both argument positions in (4) are assigned the same θ-role, or that in general, the mapping between θ-roles and syntactic positions is more complex. However, this would make the semantic contrast a suspicious accident. I believe that if we adopt Dowty's definition of θ-roles as sets of proto-properties that determine the syntactic realization of an argument, then it follows that *the bird* in (9a) receives an agent θ-role and *the kite* in (9b) a patient θ-role.

The implications that variable unaccusativity has for our understanding of θ-roles have arguably remained underexplored in the previous literature. In the terminology introduced earlier, the arguments of (9a) (*the bird*) and (9b) (*the kite*) receive the same individual role which I will refer to as a flier, defined by the entailment of moving through the air. If it is correct that at the same time, *the bird* receives an agent θ-role and *the kite* a patient θ-role, this means that a single individual role can correspond to different θ-roles. The mapping between individual roles and θ-roles is thus not simply one-to-one. More fundamentally, variable unaccusativity demonstrates that the θ-roles and individual roles of a given verb can and should be distinguished: being a flier and being the agent of a flying event are not simply the same thing. Both arguments in (4) come with the entailment of moving through the air, but only one of them also comes with the entailment of being sentient and acting volitionally.

I argue that variable unaccusativity is not the only domain in which θ-roles and individual roles do not stand in a one-to-one mapping relationship. A second set of data with a similar profile are direct causatives formed from unergative verbs, seen for Hindi-Urdu in (5):

(5) a. Rohan naach rahaa hai.
Rohan.M dance PROG.MSG be.PRS.3MSG
'Rohan is dancing.'

b. Shama Rohan-ko nach-aa rahii hai.
Shama.F Rohan-DOM dance-CAUS PROG.F be.PRS.3MSG
'Shama is making Rohan dance/twirling him around (the dance floor).'

(Bhatt and Embick, 2017:124)

In previous work (Neu, to appear), I have shown that the intransitive in (5a) passes unergativity diagnostics, indicating that the sole argument, *Rohan*, is realized as an external argument. In the causative (5b), on the other hand, there is syntactic evidence that *Rohan*, now surfacing as an object, is base-generated as an internal argument. If external arguments are agents and internal arguments patients, this predicts that Rohan is an agent in the intransitive and a patient in the transitive.

Again, one might object that the mapping between syntactic positions and θ-roles could be more complex in this case. But again, the difference in syntactic properties also goes along with a difference in semantic proto-properties. In (5a), Rohan acts purposefully and of his own accord, whereas in (5b), he is at the mercy of Shama, being passively twirled and twisted around the dance floor like a puppet on strings. Therefore, I argue that *Rohan* is assigned an agent θ-role in the intransitive but a patient θ-role in the transitive. In the latter, *Shama* serves as the agent instead.

Let us look at the individual roles. Both versions of Rohan share the salient entailment of moving their feet. I thus attribute them the same individual role which I label a dancer. Shama, on the other hand, is associated with the entailment of twirling someone else around. Thus, as seen previously for variable unaccusativity, a single individual role – that of a dancer – is mapped onto two different θ-roles. In addition, however, we now also find a single θ-role – an agent – being mapped onto two different individual roles. Rohan in (5a) and Shama in (5b), while both being agents, are associated with very different non-thematic entailments. What this suggests is that the mapping relationship between θ-roles and individual roles not only can be one-to-many, but many-to-many.

The upshot of this is that contrary to the picture I outlined at the beginning of this section, not only can the θ-roles and the individual roles of a given verb be dissociated, but individual roles do not seem to systematically determine θ-roles, nor θ-roles to determine individual roles. To know that an argument is interpreted as the agent of an event described by a given verb does not in and of itself tell us anything about the non-thematic entailments associated with the participant; it merely tells us which abstract prototype the argument most resembles. For instance, knowing that a participant is the agent argument of the root *dance*

does not tell us whether this participant moves their feet. If this is correct, it means that θ-roles, in the sense of abstract clusters that syntax and compositional semantics have access to, are not as transparently related to root meaning as we might think they are. It also means that the intuitions that we have about the individual roles of arguments are not reliable evidence for their syntactic position, assuming that the latter are determined by θ-roles. And it raises the question of how the concrete interpretation of an argument is determined, if not as the trivial result of combining a θ-role and a verbal root.

The second part of this prospectus will survey a range of other constructions that have the same profile as variable unaccusativity and direct causatives of unergatives. I label this group of phenomena *variable linking*. What I mean by this term are argument-structural alternations in which a single individual roles is mapped onto different θ-roles in different alternants. A possible response to these data and the complications they pose for a theory of θ-roles is to argue that the analyses I give for them are incorrect and that the examples, if properly analyzed, do not constitute any evidence for variable linking. Specifically, one might either argue that the arguments in question occupy different syntactic positions from what I propose or that the mapping between θ-roles and syntactic positions is more complex. Part of what this dissertation will need to achieve is to convince the reader that this is not a valid strategy. I will try to do so partly by engaging in detail with alternative analyses from the previous literature, partly by construing what I believe to be the strongest version of an alternative analysis and then attempting to refute it.

If θ-roles and individual roles indeed do not need to stand in a one-to-one mapping relationship, then we need a theory of θ-roles that is compatible with this state of affairs. The main goal of this dissertation is to develop such a theory. Specifically, I want to understand how abstract θ-roles, when combining with an argument and an event description, come to tell us which concrete kind of activity the argument performs in the event. This theory should neither under- nor overgenerate: while it needs to allow for variable linking, it must also ensure that *Brutus stabbed Caesar* cannot mean that Caesar stabbed Brutus. It should answer questions like the following: why is the agent argument of *stab* interpreted as performing rather than as undergoing a stabbing, and why can it not be otherwise? Why can the agent argument of *dance* be someone who dances but also someone who makes someone else dance? Why can a participant who flies be variably either an agent or a patient? How should we formalize θ-roles, and what exactly is the status of individual roles? What is the algorithm by which event participants are grouped into the finite set of distinct clusters that we call θ-roles? Where in the grammar does this algorithm apply? It is not the goal of this dissertation to determine the specific set of proto-properties that each θ-role is associated with, but to understand how, given a set of properties, a θ-role can determine the concrete interpretation of an argument.

I believe that besides contributing to our understanding of θ-roles, this project is of broader relevance for research on argument structure. First and foremost, how we think about the mapping between θ-roles and individual roles shapes our analysis of argument-structural configurations. To take the case of direct causatives of unergatives as an example, if the two Rohans in *Rohan is dancing* and *Shama is dancing Rohan* are assumed to receive the same θ-role because both come with the non-thematic entailment of moving their feet, this would entail that either both arguments are base-generated in the same syntactic position or that the same θ-role can be associated with different syntactic positions. The first claim has consequences for the syntax of the verbal domain, the second for the mapping between θ-roles and syntactic positions; neither is trivial. In general, the six case studies should make it evident that whether or not θ-roles and individual roles need to stand in a one-to-one mapping relationship matters for the analysis of a wide range of empirical phenomena. In addition, dissociating θ-roles from individual roles provides a fresh perspective on questions such as whether a single argument can be assigned more than one θ-role and how θ-roles feature in entailments across argument-structural alternants. I will briefly take up both of these points in Section 4.

In Section 1.2, I describe the theoretical and methodological framework that I will assume. Section 2 then summarizes the six case studies of variable linking that form the empirical core of this dissertation, and in Section 3 I lay out some thoughts on what a theory of θ-roles that allows for variable linking could look like. Section 4 addresses some further questions.

### *1.2. Some formalities*

In the following, I first establish some background assumptions that I will make about syntax, θ-roles and their mapping relationship. In many ways, my strategy will be to be as specific as necessary and as agnostic

as possible. The general point about θ-roles that this dissertation tries to get across is quite independent of many concrete issues that are debated in research on argument structure, and on many of them I will take no, or only a tentative, stance. At the end of this section, I then delineate the empirical scope of this work.

To begin with syntax, I will assume the overall architecture developed in the tradition of Distributed Morphology (DM, Halle and Marantz, 1993, 1994), mostly in the interest of connecting this dissertation with current work on argument structure in the generative mainstream. Concretely, I adopt the idea that verbs are built by merging an acategorical root with a verbalizer *v*. I assume that the latter also introduces the patient θ-role assigned to the internal argument, whereas the agent θ-role assigned to the external argument is introduced by Voice (Kratzer, 1996). The general structure is given in (6).

(6)

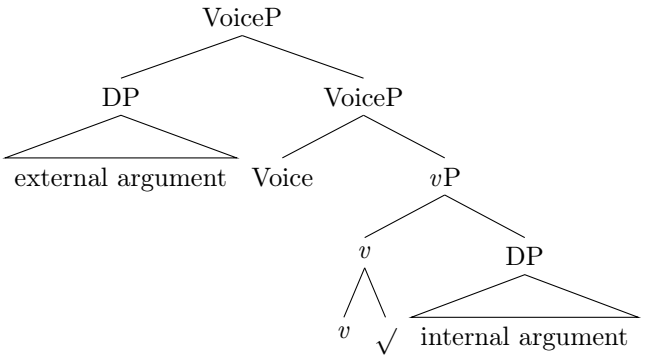


I will analyze dative-marked arguments as being introduced by an applicative head, whether low or high, following Pylkkänen (2008). In the case studies, I will at times discuss whether additional argument positions should be proposed to account for the data; in each case, I will conclude that they should not.

This is of course not the only view on verbal syntax, and it might not even be the best one. It is not the goal of this dissertation to figure out which is. The key question we will always be concerned with in the case studies is whether two arguments occupy the same or different positions, and this question is to a considerable extent independent of what precisely the syntax of any of these positions is. For this purpose, a wide variety of assumptions about the structure of the verbal domain will do, and I choose one that is widely accepted and overall well corroborated. On that note, the case studies employ a variety of syntactic tests, sometimes remaining agnostic as to which positions exactly these tests pick out – let alone why –, as long as they help us to distinguish between two syntactic positions.

A quick word on event structure. There is broad agreement that even a single verbal domain can correspond to a complex event structure containing two or more eventualities. For instance, direct causatives are standardly analyzed as involving, at the very least, a causing event and a result state (Alexiadou et al., 2015; Levin and Rappaport Hovav, 1995; Ramchand, 2008; von Stechow, 1996, a.m.o.). However, I assume that within a single verbal domain, these are merely subevents that taken together constitute an event that serves as the domain for thematic uniqueness (Carlson, 1998). In the case studies sketched out below, I will set out the analyses in terms of this larger superevent, glossing over the internal event-structural make-up without denying its reality. In the final dissertation, I plan to argue more in detail why a more fine-grained perspective on event structure does not do away with variable linking.

Next, let us turn to θ-roles. As outlined in the previous section, my notion of θ-roles is largely Dowty's. I borrow two ideas in particular: a) θ-roles are defined by sets of entailments, and b) θ-roles are prototypes that arguments can resemble to a greater or lesser degree. Dowty posited only two such proto-roles, agent and patient, aiming to trim down the endlessly proliferating inventory of θ-roles that had been proposed by researchers over the years. Their respective proto-properties are listed in (7) and (8). Dowty parenthesizes the last two items on each list, arguing that they might also rather be attributed to the discourse-related properties of subjecthood.

(7) *Contributing properties for the agent proto-role*
- a. volitional involvement in the event or state
- b. sentience (and/or perception)

c. causing an event or change of state in another participant
d. movement (relative to the position of another participant)
e. (exists independently of the event named by the verb)

(8) *Contributing properties for the patient proto-role*

a. undergoes change of state
b. incremental theme
c. causally affected by another participant
d. stationary relative to movement of another participant
e. (does not exist independently of the event named by the verb) (Dowty, 1991:572)

I diverge from Dowty only minimally by adding to his list the proto-role of goal, which he regarded merely as an intermediate point between agents and patients. I do so for two main reasons. First, goals do seem to have a distinctive profile of proto-properties not reducible to medium-agent-medium-patient-hood, such as receiving a transfer of possession, whether physical or mental, being emotionally affected and suffering positive or negative consequences as a result of the event. Secondly, both psycholinguistic and typological work has offered evidence for the cognitive reality of agent, patient and goal, but not of other roles such as instrument (see Rissman and Majid, 2019 for an overview). The role of goal, covering more fine-grained labels such as recipient, location, affectee, benefactive, experiencer, etc., is a broad one. I leave it as an open question whether it ultimately needs to be differentiated further to account for various grammatical phenomena (e.g., Beavers and Koontz-Garboden, 2020), but for our purposes, I do not believe that it does.

In addition, another strong candidate for a distinct θ-role are state holders. Holders are unusual in that they are defined not so much in terms of the properties attributed to them but based on the kind of eventuality they are a part of, suggesting that they are a somewhat different concept than agents, patients and goals. I am unsure about their status; we will return to this matter when discussing stative passives in Section 2.5.

In sum, the θ-role inventory I will assume consists of agent, patient and goal. It is by no means a settled question which entailments exactly are associated with them, but we should have a sufficiently clear intuition. Dowty argued that to determine the θ-roles of an argument that has properties associated with several roles, the grammar simply computes with which role it shares the most properties. Such a counting algorithm seems too simplistic; rather, the comparison most likely operates on a more holistic *gestalt* level.

While I will largely operate under the assumption that each of the three roles comes with its own set of proto-properties, I also plan to discuss in later parts of the dissertation the possibility that one of these three roles does not have any proto-properties itself but is defined merely in relation to the other two. This includes Dowty's claim that goals are defined as mid-points between agents and patients, but also the view maintained by Kratzer (2002) that patients are merely a default role assigned to any argument that does not meet the specifications of any other role. For the purposes of the theoretical puzzle that I will try to solve, this is a relevant question; for the purposes of the empirical case studies, it should not matter too much.

A crucial question to address next is how to determine whether two arguments are assigned the same or different θ-roles, and what role syntactic evidence plays in this. Recall that the puzzle this dissertation is concerned with are cases of variable linking, where two arguments that receive the same individual role are assigned different θ-roles. In the above sketch of variable unaccusativity and direct causatives of unergatives, I have motivated the existence of variable linking in part by arguing that the arguments in question occupy different syntactic positions. Of course, this line of argument only goes through if different syntactic positions are expected to be associated with different θ-roles, and the reader might object that this is not necessarily the case.

I agree with this objection, and I do not think that a difference in the syntactic position of two arguments strictly entails that they receive different θ-roles. We know too little about the mapping between θ-roles and syntactic positions to warrant this conclusion. But I do think we know something. We know that it is very common for external arguments to be associated with agent proto-properties, for internal arguments with patient proto-properties and for applicative arguments (indirect objects) with goal proto-properties. It should be noted that unlike Dowty, I thus describe the syntactic repercussions of θ-roles in terms of underlying argument position, not in terms of surface subject- and object-hood. The mapping between θ-roles and syntactic positions might very well be more complex than that. But operating under the assumption that any syntactic position can be associated with any θ-role is too weak a theory.

Hence, I rely on the following strategy for determining whether we are dealing with the same or different θ-roles. We can conclude that two arguments receive different θ-roles if and only if a) they differ in their syntactic position (independently of the verb), b) they also differ with respect to their proto-properties, and c) those proto-properties match those standardly associated with their respective syntactic positions. By this I mean that external arguments should be associated with agent proto-properties, internal arguments with patient proto-properties and applicative arguments with goal proto-properties. For instance, in the above discussion of variable unaccusativity, I have shown that in the unergative variant, the argument not only behaves syntactically as an external argument but is also associated with proto-agent properties such as sentience, and vice versa for the unaccusative variant. I believe that taken together, these two pieces of evidence warrant the conclusion that we are indeed dealing with different θ-roles.

In sum, in the case studies, I will diagnose θ-roles based on proto-properties and syntactic evidence, based on Dowty's definition of θ-roles as clusters of proto-properties that determine the syntactic position of an argument. This is certainly not the only way of defining θ-roles. θ-roles are not things in the world that we could straightforwardly observe but theoretical constructs whose validity depends on the work they do in explaining the world to us. Whether the way of operationalizing θ-roles adopted here is a good one will need to be judged based on the overall success of this work as a whole.

Finally, I give a brief overview over the range of empirical data that we will see in the case studies. In general, the goal is to confine ourselves to structures for which the mapping between θ-roles and syntactic positions is comparatively well established, namely, that can be analyzed in terms of external arguments, internal arguments and some sort of indirect object position. (This is not to say that more complex analyses have not been devised for the constructions surveyed in the case studies, but only that I argue that they do not need to be devised.) Hence, we will not see indirect causatives or any biclausal constructions, small clauses, figure/ground alternations, the spray/load alternation, the dative alternation, etc. A full theory of θ-roles should of course eventually cover all of these constructions, but they fall outside the scope of this dissertation. With all of this being said, let us do some syntax.

## 2. Six variations on variable linking

In this section, I discuss six case studies of variable linking. I do so mainly in order to make the case that variable linking is a widespread and robust phenomenon that any theory of θ-roles needs to be able to account for. In addition, the constructions surveyed in the following are interesting in their own right and, for the most part, controversial. My hope is that the perspective from variable linking will contribute to our understanding of them.

The six case studies are arranged in order of increasing complexity. For the purposes of this prospectus, I merely provide an overview over the basic data, sketch out my analysis, hint at some evidence for it, explain why under this analysis, the alternation constitutes an instance of variable linking, and briefly summarize the relevant previous literature. Overall, my goal will not be to provide a comprehensive analysis of each case study but to focus on the mapping relationship between θ-roles and individual roles. This section ends with a brief typological summary of the six case studies discussed.

### *2.1. Variable unaccusativity*

I begin with the most basic case, variable unaccusativity. Many intransitive verbs can be used in either an unergative or an unaccusative structure, depending on the wider context of the sentence (Borer, 2005; Dowty, 1991; Perlmutter and Postal, 1984; Sorace, 2000, 2004, 2011, a.m.o.). In the Hindi example (4) discussed above in Section 1.1 and repeated below as (9), the verb 'fly' can appear in a reduced relative in (9b) but not in (9a):

(9) a. *uṛ-ii (huu-ii) ciṛyaa
fly-PFV.FSG be-PFV.FSG bird.FSG
Intended: 'the flown bird'

b. uṛ-ii (huu-ii) patang
fly-PFV.FSG be-PFV.FSG kite.FSG
'the flown kite' (Ahmed, 2010:8f.)

This indicates that the same individual role, which is defined by the entailment of moving through the air and which I label a flier, is realized as an agent in the external argument position in (9a) but as a patient in

the internal argument position in (9b). As argued earlier, in (9a), flying is a purposeful activity that requires effort on part of the bird, in line with the agentive construal of the argument, whereas in (9b), the kite merely passively floats in the air and is accordingly construed as a patient. In terms of Dowty's proto-properties, the bird but not the kite can be described as volitional and sentient.

A similar effect of animacy/sentience can be observed in example (10). In Russian, only unaccusative structures allow for closest conjunct agreement. In example (10a) containing an animate argument, closest conjunct agreement is ruled out, which is evidence for an unergative syntax. Example (10b) containing an inanimate argument, on the other hand, passes this unaccusativity diagnostic:

(10) a. *Na lestničnoj ploščadke stojal sosed i ego brat.
on stairway landing stood.MSG neighbor.MSG.NOM and his brother.MSG.NOM
Intended: 'My neighbor and his brother were standing on the stairway landing.'
b. Na stole stojal stakan i kuvšin.
on table stood.MSG glass.MSG.NOM and jug.MSG.NOM
'On the table stood a glass and a jug.' (Krejci, 2020:126f.)

In the final example for variable unaccusativity in (11), the relevant proto-property is not sentience, but volition. In Tsova-Tush, unergative verbs are realized with ergative case and unaccusative verbs with nominative case. In (11a), the verb 'fall' is construed as an intentional activity and thus surfaces with an unergative syntax in which the agent receives ergative case. In (11b), on the other hand, the argument receiving the same individual role is interpreted as a patient instead, in line with the nominative case marking indicating an unaccusative syntax.

(11) a. (as) vuiž-n-as.
1SG.ERG fell-AOR.1SG-ERG
'I fell down, on purpose.'
b. so vož-en-sO.
1SG.NOM fell-AOR.1SG-NOM
'I fell down, by accident.' (Holisky, 1987:105)

In the examples surveyed so far, the syntactic realization of the intransitive was sensitive to how well the individual role of the single argument matched the proto-role of an agent in terms of sentience, intentionality, volition, etc. The second factor that has been argued to play a role for variable unaccusativity is the telicity status of the event (e.g., Sorace, 2000, 2004, 2011, a.m.o.). Undergoing a change of state is a property associated with patients (Dowty, 1991); hence, typically unergative verbs can allow for an unaccusative construal when appearing with an overt result state. This effect is demonstrated below for Dutch, where unergatives appear with the auxiliary 'have,' but unaccusatives take 'be.' According to this diagnostic, in (12a), the bare intransitive 'roll' behaves as an unergative, whereas 'roll downstairs' in (12b), in which the ball ends up in the state of being downstairs, passes as unaccusative:

(12) a. De bal heeft/*is gerold.
the ball has/is rolled
'The ball rolled.'
b. De bal is/*heeft naar beneden gerold.
the ball is/has to down rolled
'The ball rolled downstairs.' (Sorace, 2000:876)

Overall, the unergative and unaccusative variants differ from each other semantically precisely with respect to properties that have been associated with agent and patient θ-roles, namely, volition and sentience on the one hand side and undergoing a change of state on the other. This strongly supports the claim that there are indeed two different θ-roles involved. Variable unaccusativity is widely attested cross-linguistically, and their analysis is largely uncontested (but see Chen, 2024). These cases thus provide a solid empirical basis for the claim that θ-roles and individual roles must be dissociated.

*2.2. Direct causatives of unergatives*

Turning to a more complex case, several languages allow verbs that as intransitives pass unergativity diagnostics to also appear in direct causative constructions, seen for Hindi-Urdu in (13), for Turkish in (14) and for Sason Arabic in (15):

(13) a. Rohan naach rahaa hai.
Rohan.M dance PROG.MSG be.PRS.3MSG
'Rohan is dancing.'
b. Shama Rohan-ko nach-aa rahii hai.
Shama.F Rohan-DOM dance-CAUS PROG.F be.PRS.3MSG
'Shama is making Rohan dance/twirling him around (the dance floor).'
(Bhatt and Embick, 2017:124)

(14) a. Bebek uyu-du.
baby sleep-PST
'The baby slept.'

b. (Ben) bebeğ-i uyu-t-tu-m.
I baby-ACC sleep-CAUS-PST-1SG
'I put the baby to sleep.'

(15) a. i-zak.
3M-laugh
'He laughs.'

b. a-zakkiy-u.
1SG-laugh.CAUS-him
'I make him laugh.' (Akkuş, 2021:175)

These data pose a puzzle if it is assumed that the subject of the intransitive and the causee of the transitive are assigned the same θ-role and occupy the same syntactic position. Given the evidence from unergativity diagnostics, the intransitive subject is expected to occupy the external argument position. However, direct causatives – unlike indirect causatives – are widely assumed to only make one external argument position available (but see Nie, 2020), and this position should already be occupied by the causer.

As outlined in Section 1.1, I have argued in previous work that direct causatives of unergatives in Hindi-Urdu, Turkish and Sason Arabic are standard transitives, in line with previous work by Harris (1981) on Georgian and Legate (2014) on Acehnese. The subject is realized as an agent argument in the external argument position and the object as a patient argument in the internal argument position (16):

(16) a. *Intransitive*

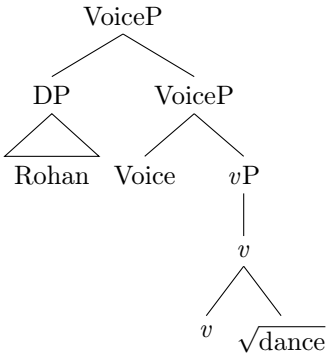


b. *Transitive*

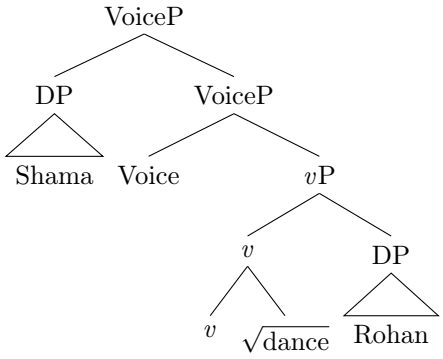


Semantically, as described earlier, this is reflected in the difference in thematic properties ascribed to the object of the causative and the subject of the intransitive, respectively. That is, the latter is depicted as acting volitionally, whereas the former is not. Syntactically, I have shown that the object of the causativized predicate passes internal argument diagnostics, namely, being able to appear in reduced relatives (Hindi-Urdu), induce a telic interpretation (Turkish) and take a resultative secondary predicate (Sason Arabic), contrasting in all of these respects with the subject of the intransitive. The upshot of this analysis is that an argument receiving the same individual role in the intransitive and transitive variant – a dancer, a sleeper

or a laugher – is associated with two different θ-roles. In addition, a single θ-role – that of an agent – is associated with two different individual roles. The mapping relation between individual and θ-roles is thus not only, as in variable unaccusativity, one-to-many, but many-to-many.

Direct causatives of unergatives have been attested in a variety of languages, and different alternative analyses that avoid variable linking have been proposed (e.g., Nash, 2021; Ramchand, 2008). One approach, which in Neu (to appear) I have labeled the low subject analysis, has posited that subjects of unergatives are base-generated in Spec$v$P, but subjects of transitives in SpecVoiceP (17) (Kouneli, 2021; Massam, 2009; Myler, 2022; Tollan, 2018; Tollan and Massam, 2022; Tollan and Oxford, 2018; see also Kumaran, 2021; Pineda and Berro, 2020).

(17)

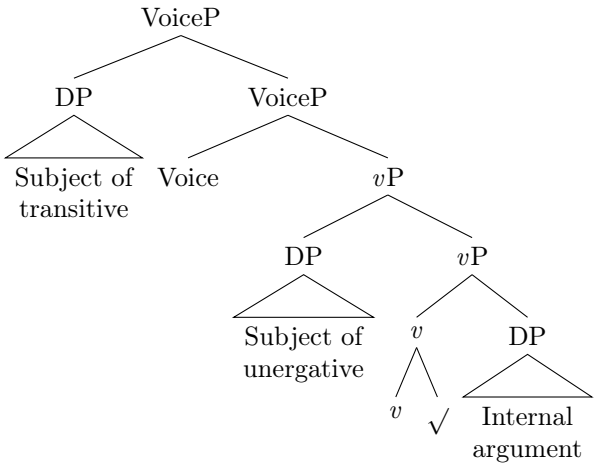


As a result, unergatives can causativize by merging an external causer in SpecVoiceP, whereas in transitives, this position is already occupied.

The goal of the low subject analysis is to ensure that the subject of the intransitive and the object of the transitive are realized in the same syntactic position and assigned the same θ-role in both alternants. That is, the shared non-thematic entailments between the two arguments – such as moving one's feet – are attributed to a shared θ-role. The nature of this θ-role assigned to the Spec$v$P position is not always clarified, but Massam (2009) argues that Spec$v$P receives a 'doer' $\theta$-role, as opposed to the 'agent' $\theta$-role assigned to SpecVoiceP, and Tollan (2018) describes unergative subjects as 'non-effortful and/or non-affecting agents or experiencers. Neither of those two categories can be clearly distinguished from agents. More acutely, I have shown that the low subject analysis fails to capture the difference between the subject of the intransitive and the object of the transitive in terms of both thematic properties and syntactic behavior.

An open question is how direct causatives of unergatives compare to English transitives of manner of motion verbs (18):

(18) a. Amir danced around the hallway.
b. Sepideh danced Amir *(around the hallway).

Contra previous work (Biggs, 2019; Folli and Harley, 2006; Ramchand, 2008), I suspect that transitives of manner of motion verbs are regular agent/patient transitives and thus fall into the broader class of direct causatives of unergatives. However, transitives of manner of motion verbs are peculiar in that, for instance, (18b) requires a PP to be felicitous. I currently have no explanation for this fact.

### *2.3. Direct causatives of ingestives*

Direct causatives of ingestives are a natural extension of direct causatives of unergatives. Ingestives, or ingesto-reflexives, are a class of transitives such as 'eat,' 'drink,' 'see,' 'hear,' 'learn,' etc. In Hindi-Urdu (19) and Sason Arabic (20), ingestives are able to form direct causatives:

(19) a. Mina-ne angrezii siikh-ii
Mina-ERG English.F learn-PFV.F
'Mina learned English.'
b. Tina-ne Mina-ko angrezii sikh-aa-yii.
Tina-ERG Mina-DAT English.F learn-CAUS-PFV.F
'Tina taught Mina English.' (lit. 'Tina learned Mina.DAT English.')

(Bhatt and Embick, 2017:128)

(20) a. şarab-e mayn.
drank-F water
'She drank water.'
b. şarrıp-to-lla mayn.
drank.CAUS-1SG-her.DAT water
'I gave her water to drink.' (lit. 'I drank her.DAT water.')

In both examples, the subject of the transitive bears the same individual role as the dative-marked argument of the causative. The latter can be shown to occupy a position identical to that of indirect objects in standard ditransitives (Bhatt and Embick, 2017; Neu, to appear); in (21), I tentatively depict it as a high applicative argument.

(21) *Transitive*

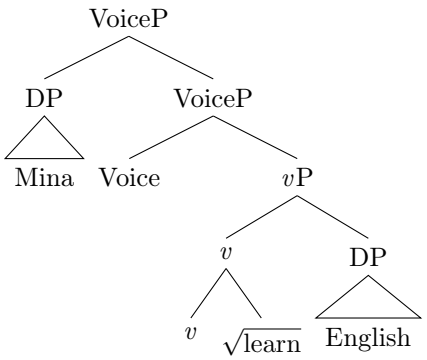


(22) *Ditransitive*

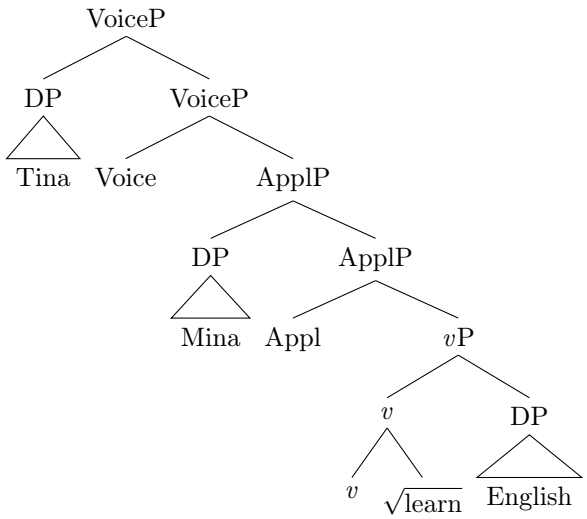


Again, this difference in syntactic realization goes along with a difference in semantic interpretation: while the subject of the transitive receives the role of an agent, acting intentionally, the indirect object of the ditransitive is construed as a goal, receiving English lessons or water, respectively.

As in direct causatives of unergatives, the mapping between individual roles and θ-roles is thus many-to-many: the individual roles of the learner and drinker are mapped onto different θ-roles in the two variants, and the agent role associated with the external argument position corresponds to two different individual roles (learner and teacher, and drinker and making-drinker). Cross-linguistically, it is highly common for ingestives to be the only transitives that can form direct causatives in a language (Krejci, 2020). The analysis given above has been anticipated by Harris (1981) for Georgian and Legate (2014) for Acehnese; a competing account that remains to be discussed is Krejci (2020).

### *2.4. Agent/affectee alternations*

Agent/affectee alternations are the non-valency changing counterpart of direct causatives of ingestives. One of the Hindi ingestives that can causativize to form ditransitives, discussed in the previous section, is the

verb 'see.' Example (23) demonstrates that 'see' allows for a dative-marked argument even in the absence of an external causer:

(23) a. Ram-*ne* Sita-ko dekh-aa.
Ram-ERG Sita-DOM see-PFV
'Ram saw Sita.'
b. Ram-*ko* Sita dikh-ii.
Ram-DAT Sita see-PFV.F
'Ram saw Sita.' (lit. Sita appeared to Ram) (Bhatt and Embick, 2017:130f.)

The object 'Sita' in (23b) cannot bear DOM marking or be passivized, providing further evidence that the two variants in (23) differ in their underlying syntax. Semantically, the interpretation of Ram in the two variants differs along the same lines as seen for direct causatives of ingestives: in (23a), Ram intentionally observes Sita, whereas in (23b), she merely happens to have stumbled into his field of vision. The evidence thus indicates that Ram, bearing the individual role of a seer, is realized as an agent in the external argument position in (23a) and as a goal in the indirect object position in (23b), as shown in (24):

(24) a. *Agentive reading*

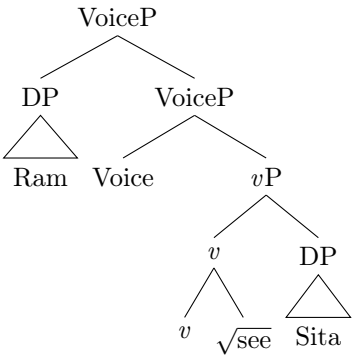


b. *Affectee reading*

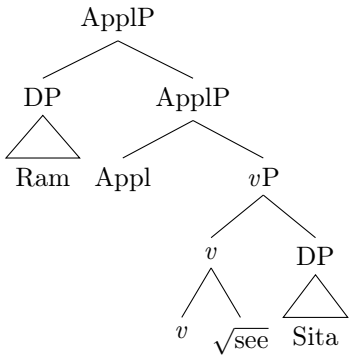


To keep up with conventions, I will describe the dative marked argument as an affectee, but I do not commit to the idea that affectee is a θ-role distinct from goal.

While the agent/affectee alternation in (23) is highly restricted in Hindi, languages such as Spanish (25) and German (26) allow it much more productively:

(25) a. El tintorero quemó los pantalones de Carolina (a propósito/para vengarse).
the dry-cleaner burnt.SG the trousers of Carolina (on purpose/to take revenge)
'The dry-cleaner burnt Carolina's trousers (on purpose/to take.revenge).'
b. Al tintorero se le quemaron los pantalones de Carolina (*a propósito/*para vengarse).
to.the dry-cleaner REFL her.DAT burnt.PL the trousers of Carolina (on purpose/to take.revenge)
'The dry-cleaner (accidentally) burnt Carolina's trousers (*on purpose/*to take revenge).'
(Cuervo, 2003:186–187)

(26) a. Ich habe (absichtlich) die Vase zerbrochen.
I have (on.purpose) the vase broken
'I broke the vase (on purpose).'
b. Mir ist (*absichtlich) die Vase zerbrochen.
me.DAT is (on.purpose) the vase broken
'I (accidentally) broke the vase (*on purpose).'

In both sets of examples, a single individual role – a burner and a breaker, respectively – is variably mapped onto an agent or an affectee role. Under the agent interpretation, the burning or breaking could be caused

either intentionally or unintentionally. Under the affectee interpretation, the burning or breaking is caused unintentionally, and the focus is on the fact that the event happens to the participant's detriment.

Similar cases are reported cross-linguistically (see Chen, 2024 for an overview). However, affectee arguments are a murkily defined group which might also be syntactically heterogeneous (Bosse et al., 2012), and not all alternations between a high- and a low-agentive reading of an argument need to involve a difference in the underlying syntax. Hence, I plan to focus on agent/affectee alternations in one or two languages – perhaps Spanish and German – and carefully establish their semantic and syntactic properties.

### *2.5. Stative passives*

Stative passives are another area where we see potential cases of variable linking. As is well known, statives can normally take an internal, but not an external argument (27):

(27) a. The door is opened. (i.e., the door is in the state of having been opened)
b. *The girl is opened. (i.e., the girl is in the state of having opened something)

One possible analysis suggested by this contrast is for a stativizing head to only be allowed to merge with a *v*P that contains an internal argument (28):

(28)

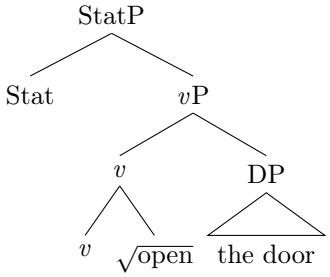


However, in some languages, statives formed from ingestive verbs can surface with an argument that we would expect to be realized in the external position and to be interpreted as an agent. The phenomenon is shown for Greek in (29). Example (29a) is a regular stative containing, as expected, an internal argument interpreted as a patient. In contrast, (29b) is ambiguous between two interpretations. Besides the syntactically expected, albeit pragmatically odd reading that Mary is in the state of having been eaten, the sentence can also mean that Mary is in the state of having eaten – apparently exactly what was ruled out in (27b).

(29) a. I zoni ine asfalis-men-i.
the.NOM belt.NOM be.3SG secure-PTCP-F.NOM
'The seat belt is fastened.'
b. I Maria ine faɣo-men-i.
the.NOM Mary.NOM be.3SG eat-PTCP-F.NOM
'Mary is eaten.' / 'Mary has eaten.' (Paparounas, 2025)

An argument bearing the individual role of an eater, such as *Mary* in *Mary eats an apple*, is normally realized as an agent in the external argument position, making its appearance in a stative passive unexpected. Such pseudo-agentive statives are surprisingly common, having been attested in English (30), Hebrew (31), Italian (32) and Latin (Haspelmath, 1994), among others:

(30) Mary is drunk/learnéd/well-read. (Paparounas, 2025:24)

(31) a. Ha tapux haya axul.
the apple be.PST.3SG eat.PTCP
'The apple was in an eaten state.'
b. Lazet la derex axulim ve stuyim.
get under way eat.PTCP.PL and drink.PTCP.PL
'Set off having eaten and drunk enough.' (Arad, 1998, cited in Paparounas, 2025:24)

(32) Venite giá mangiati e bevuti.
come.IMP.2PL already eat.PTCP.PL and drink.PTCP.PL
'Come over after having eaten and drunk.' (Arad, 1998, cited in Paparounas, 2025:25)

I propose to analyze these examples as cases of variable linking. Focusing on (29b), the individual role of an eater is mapped onto a patient rather than an agent θ-role. Someone who eats undergoes a change of state, from being hungry to being full, and Paparounas (2025) notes that (29b) can only be used if the eating event has been completed and has reached its result state. The same is true for other ingestive verbs that show the unexpected stativization pattern in (29b). As pointed out earlier in the context of variable unaccusativity, undergoing a change of state is a property associated with patients. Accordingly, in the absence of an eatee – which would outshine the eater with respect to patient-like properties –, an eater qualifies as a patient realized in the internal argument position and can thus surface in the stative.

The analysis sketched out so far glosses over some complications. For reasons that are largely orthogonal to our purposes, a number of recent papers has argued that statives are formed by merging the stativizing head not with a full *v*P, as shown in (28), but with a *v* head, and that the DP argument is merged in SpecStatP where it receives the θ-role of a holder (33) (Biggs and Embick, 2025; Embick, 2023; Hamo and Meihami, 2024; Paparounas, 2025). The relevant denotation for Stat is given in (34).

(33)

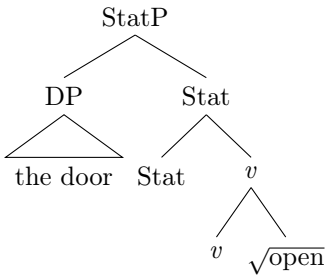


(34) $[\![\text{Stat}]\!] = \lambda P.\lambda x.\lambda s.\exists e'.P(e') \wedge CAUSE(e', s) \wedge STATE(s) \wedge HOLDER(s) = x$
(Paparounas, 2025:27)

My above explanation for why seemingly agentive arguments like eaters can appear in statives relied on the idea that statives take a patient argument, and that arguments of ingestives can qualify as patients by virtue of undergoing a change of state. If statives in fact take a holder argument, this explanation seems to lose its ground.

However, the analysis in (33) cannot be the full picture. If the DP argument is a holder of the state resulting from the verbal event but is not actually thematically integrated into the verbal event itself, nothing so far blocks **The girl is opened*. In other words, it is not clear why *the girl* does not qualify as a holder of the state resulting from the opening. Paparounas (2025), recognizing this conundrum, adds to the above analysis the meaning postulate in (35):

(35) a. $\forall x.\forall e.\forall s.[event(e) \wedge CAUSE(e, s) \wedge STATE(s) \wedge HOLDER(s) = x] \rightarrow THEME(e) = x$
b. OR $\forall x.\forall e.\forall s.[event(e) \wedge CAUSE(e, s) \wedge STATE(s) \wedge HOLDER(s) = x] \rightarrow AGENT(e) = x$ in the context of $\sqrt{\text{eat}}$, $\sqrt{\text{learn}}$... (Paparounas, 2025:28f.)

Under this analysis, the eater argument in *Mary is eaten* is thus assigned an agent θ-role after all, contrasting with regular statives in which the argument is assigned a patient θ-role. In this way, Paparounas ensures that the non-thematic entailment of ingesting food is consistently associated with the same θ-role.

Thus, the original puzzle would still remain – namely, why some verbs seem to suddenly license an agent role where we expect a patient – and the solution from variable linking would still apply. Overall, there are a number of open (and intriguing) questions surrounding stative passives, concerning their status as phrases (28) or complex heads (33), their event structure, the nature of holder arguments, and certain contrasts between predicative and attributive uses that I have not touched on here. For the purposes of this dissertation, I do not hope to solve all, or even most of these questions. My goal is rather to show that whichever reasonable analysis is adopted, pseudo-agentive stative passives of ingestives are evidence for variable linking.

### *2.6. Symmetric predicates*

The final case study has a somewhat different flavor than the constructions surveyed so far, and it is in the most infant state. I use the term symmetric predicates as an umbrella term for strictly symmetric predicates (e.g., *marry*, *meet*, *rhyme*) and pseudo-symmetric predicates (e.g., *hug*, *talk to*, *kiss*) (Winter, 2016, 2018). Both verb classes allow both for an intransitive use with a conjoined DP as a subject and for a transitive use. In the latter case, either argument can surface as subject or object, with either little or no effect on the truth conditions of the sentence. With strictly symmetric predicates, the intransitive and both transitives are truth-conditionally equivalent; i.e., in (36), each entails the other two:

(36) a. Joel and Clementine married.
b. Joel married Clementine.
c. Clementine married Joel.

With pseudo-symmetric predicates, these entailments need not hold. In (37), all three sentences could be true at the same time, but it might also be the case that Clementine hugged Joel against his will, making (37c) true, (37b) false and (37a) at least heavily infelicitous:

(37) a. Joel and Clementine hugged.
b. Joel hugged Clementine.
c. Clementine hugged Joel.

I will make two assumptions for now, both of which we will need to revisit later on. First, I assume that we can take the morphosyntax of these sentences at face value. That is, the intransitive uses have only an external argument, while the transitive uses have an external and an internal argument. Secondly, I posit that symmetric predicates come with two identical individual roles, such as two marriers or two huggers.

Against this background, the alternations in (36) and (37) present themselves as instances of variable linking. The two individual roles can either both be realized as agents, or one of them as an agent and the other as a patient. In the pseudo-symmetric example (37), this is reflected in the interpretation of the sentences. In (37a), Joel and Clementine qualify as equally agentive for the purposes of the hugging, resulting in a reading where the hugging is mutual and symmetric. In the transitives (37b) and (37c), either one or the other is depicted as more agentive. This is compatible with a broad range of readings: if Clementine hugs Joe, Joe might resist the hugging, passively tolerate it, cautiously return it, or even engage in hugging behavior just as enthusiastically as Clementine. In the latter case, with both participants being equals with respect to thematic properties, θ-role assignment is random and subject to information-structural pressures only.

With strictly symmetric predicates, mapping the individual roles onto different θ-roles comes with no interpretative changes other than information-structural. I assume that there is no more or less agentive way of marrying someone, or more in general, that the two participants of strictly symmetric predicates are necessarily equals with respect to thematic properties (Dowty, 1991; Winter, 2016, 2018). Hence, with these verbs, θ-roles is always purely information-structurally determined.

This is a broad sketch of an analysis at best, and as flagged earlier, the assumptions that it relies on are not innocuous. One complication comes from VP ellipsis. In (38), Clementine is obligatorily interpreted as being unable or unwilling to marry Joel; the sentences cannot mean that she is unable or unwilling to marry, say, Patrick:

(38) a. Joel and Clementine should marry, but Clementine can't.
b. Joel wants to marry Clementine, but she doesn't.

If the conjuncts are syntactically regular transitives or intransitives, it is not clear what forces this interpretation. Rather, the data in (38) appear to be more in line with an analysis according to which the transitive and intransitive variants are derived from a single underlying structure. I suspect that the ellipsis facts might fall out from pragmatic pressures, but I currently have no worked-out solution to this challenge.

What is more, symmetric predicates push the notion of individual roles that I have relied on so far beyond its comfort zone. Do verbs like *hug* really have two identical individual roles? Shouldn't we rather distinguish

between someone who hugs actively and someone who is being hugged? What about someone who hugs, but in a fairly restrained fashion? Someone who welcomes the hugging as opposed to someone who resists it? Should we postulate two, three, four different individual roles? The discomfort that the reader might feel at this point is a healthy one. A key upshot of symmetric predicates is that they make us question the idea that each verb comes with a finite set of clearly defined, contentful individual roles as part of its semantic core. And as discussed more in detail in Section 3, I believe that abandoning this idea is key to solving the puzzle of variable linking.

### *2.7. A typological summary*

The six cases of variable linking surveyed in this section fall into different categories. To begin with, variable unaccusativity and agent/affectee alternations are the least offensive kind. In both, the same individual role can be associated with different θ-roles, but a single θ-role does not associate with different individual roles. The main difference between variable unaccusativity and agent/affectee alternations is that the latter contain an additional argument whose θ-role and syntactic position remains stable. In a sense, it is not surprising that variable unaccusativity and agent/affectee alternations should be possible. Naively speaking, the activity described by the non-thematic entailments can simply be construed in different ways, with the relevant participant either resembling an agent more than a patient, or vice versa.

Direct causatives of unergatives, direct causatives of ingestives and stative passives are more radical. Here, a single θ-role – agent in direct causatives of unergatives and ingestives, patient in stative passives – is associated with different individual roles. Direct causatives of ingestives differ from direct causatives of unergatives only by the presence of an additional argument whose θ-role and syntactic position remains stable. Crucially, all these three alternations are valency-changing; that is, in one of the alternants, one individual role does not compete for a θ-role. For instance, let us compare the Greek stative passive *Mary is eaten* to the transitive *Mary ate an apple*. In the stative, *an apple* – or some other suitable eatee – is not explicitly represented in the event; hence, the question whether it resembles the patient prototype more closely than Mary does not even apply. This seems a relevant observation to explain why these alternations are possible.

This leaves us with symmetric predicates, the most mysterious case of all. The alternation is not valency-changing; in all alternants, all individual roles do compete. But of course, symmetric predicates are not just any predicates. With pseudo-symmetric predicates such as *hug* and *kiss*, the individual roles of hugger and huggee, kisser and kissee can be held to the same degree and at the same time by a single individual, and cannot even always be told apart. Divorce lawyers arguably deal with this predicament on a regular basis. With strictly symmetric predicates such as *marry*, the two individual roles involved in such events are identical. This sharply contrasts with the individual roles of predicates like *stab*, which are clearly distinct from each other. It does not seem accidental that the complete swap of θ-roles that we see with symmetric predicates should only be possible when the individual roles involved are highly similar or identical.

## 3. Thematic roles operate at the level of the event description

Let's recap. Following Dowty (1991), I have defined θ-roles as clusters of proto-properties that determine the syntactic position of an argument. I have then shown that these θ-roles do not stand in a one-to-one mapping relationship to individual roles, defined as the non-thematic entailments associated with an argument. This poses a problem if it is assumed that θ-roles consistently pick out the same individual role for a given verb. If this assumption is not made, then the question arises what constrains the mapping relationship between θ-roles and individual roles – why, for instance, the patient role of a stabbing event is never associated with the non-thematic entailment of wielding a sharp weapon.

To answer this question, we need to clarify the concept of individual roles and what their status in the grammar should be. The idea that θ-roles pick out individual roles is, I believe, informed by a very particular view on verbal semantics. This view holds that the denotation of a verb is a structured event description that contains a finite number of discrete participant types defined in terms of verb-specific entailments. The verb *stab*, for instance, would describe a relation between someone wielding a sharp weapon and someone or something being penetrated by a sharp weapon. In other words, to know what *stab* means is to know that

there must be a stabber and a stabbee. Under this view, each θ-role then describes one (or none) of these discrete participant types: an agent of a stabbing is a stabber, a patient of a stabbing is a stabbee. This is a picture that I suspect we have inherited from lexicalism. It attributes individual roles a well-defined grammatical status by making them an integral part of the lexical semantics of verbs.

But this picture is not what is explicitly endorsed by much current work on argument structure. In DM, roots start out as acategorical conceptual blobs, and verbalized roots are simply predicates of events of type ⟨s,t⟩. Since individual roles are not explicitly represented as part of these denotations nor anywhere else in the semantics, they cannot be formally mapped onto θ-roles in the first place. Overall, in DM and similar architectures, non-thematic entailments have no well-defined grammatical status, and it is not clear why we should expect them to have any relevance for syntax.

I believe that this is a positive development. Instead of going back to structured verbal denotations that would allow us to explicitly represent the mapping between individual roles and θ-roles, we should abandon the lexicalist view for real and for good. Dowty taught us that θ-roles have a prototype structure and that the argument they describe is the one that best matches their prototype. But the arguments that θ-roles compare to find their best match is not a set of individual roles encoded on the verb. It is a set of arguments defined in the context of a complete event description, an idea anticipated in Marantz (1981).

To make this concrete, let us go back to the example of direct causatives of unergatives, specifically the case of Hindi-Urdu *dance*. The puzzle we faced here was that the argument with proto-agent properties in the external argument position could be associated with different non-thematic entailments, namely, either the entailment of moving one's feet or of twisting someone else around. That is, the agent role appears to select different individual roles. If we assume that θ-roles compare the various individual roles that are defined as part of the meaning of the verb and pick out the one that best matches their proto-properties, this behavior seems erratic.

But if we assume that θ-roles describe an argument in the context of an entire event description at the level of its existential closure or beyond, they are perfectly well behaved. One thing that a θ-role does is add information to the event description, namely, that the event involves a participant with a particular set of properties. Thus, the intransitive and the transitive use of *dance* describe different events, one of which has one, the other of which has two participants. What is more, the two events might differ in their internal event-structural make-up. In the context of the intransitive event, the sole argument matches the prototype of an agent better than that of a patient. In the context of the transitive event, one argument – the one eventually surfacing as the subject – equally matches the agent prototype well, and more so than the other argument. There is no reason to expect these two agent arguments of two different events to come with the same set of non-thematic entailments. The thematic comparison operates at the level of the event description, not at the level of the verbal root.

The goal of this dissertation is to flesh out and make precise this intuition. A main part of the challenge will be to work out a denotation for θ-roles that contributes the necessary kind of information to the event description. An agent θ-role does not simply predicate agenthood of its argument in the way that *tree* predicates treehood of whatever entity it describes. θ-roles are gradient and relative. To say that a participant is an agent means that this argument resembles the agent prototype more than it resembles any other prototype, but also more than any other participant in the same event resembles this prototype. How to formalize this calculus is an open question.

Answering this question does not have to start from scratch. I hope to be able to explore analogies to similar linguistic phenomena. One of them are gradable predicates like *tall* and *warm*, which can equally be context-sensitive and receive either a restrictive or a non-restrictive reading depending on whether there are other salient entities to which they could apply. Another are uniqueness presuppositions as encoded on definite determiners and similar expressions, which could do a lot of heavy lifting for θ-roles. In general, the strategy I plan to pursue is to think of θ-roles not as picking out an argument, already defined elsewhere, that has a certain set of properties, but to add to the event description the information that there is an argument with a certain set of properties.

To end on a speculative note, what does all of this mean for the status of individual roles? There are plenty of independent reasons to be skeptical about the idea that verbs come with a finite set of distinct and contentful participant types such as stabbers and stabbees as part of their denotation. In the discussion of symmetric predicates, I have pointed out that it is difficult to determine what exactly would have to form part

of the individual role of a hugger. Moreover, Marantz (1981) has shown that for many verbs, the activities that the event participants are said to perform or undergo vary widely depending on the choice of direct object, with no clear line separating literal and idiomatic uses (39). A single individual role for all killers, or a single role for all killees, would have to be semantically bleached to the point of becoming vacuous.

(39) a. kill a cockroach
b. kill a conversation
c. kill an evening watching TV
d. kill a bottle (i.e., empty it)
e. kill an audience (i.e., wow it)
(Marantz, 1981:49)

In a similar vein, Borer (2005) has argued extensively that the number and kinds of arguments that a verb can combine with is too malleable to fit into a single, rigid lexical entry, as demonstrated for the verb *siren* in (40):

(40) a. The police car sirened the Porsche to a stop.
b. The police car sirened up to the accident site.
c. The police car sirened the daylight out of me. (Borer, 2005:26)

All of this does not sit well with the notion that each verb comes with a finite number of well-defined individual roles.

I suspect that we need to regard individual roles as emergent properties of event descriptions. They are not always already contained on the root, but the result of combining a root with a particular set of θ-roles. The verb *stab* does not, in isolation, mean that there is a stabber and a stabbee. A verb in itself doesn't mean anything. Rather, a stabber and a stabbee is what emerges when the verb *stab* is inserted into a transitive syntax, which in this case happens to be the only syntax it can inhabit. The reader might of course object that we can have intuitions about the meaning of *stab* in isolation, and that this meaning always involves a stabber and a stabbee. I believe that this is not true. What we are actually accessing when we try to think about a verb in isolation is a prototypical usage of it, and in the case of *stab*, most of its possible usages are fairly close to this prototypical usage. Things are different for *kill*, and even more so for verbs like *take*. But all of this would need to be worked out further.

To summarize, θ-roles are sets of entailments that determine the syntactic realization of an argument. Having a prototype structure, they describe the argument that best matches their profile. This comparison does not operate on a set of individual roles stored on the verb, but on a set of arguments defined in the context of a complete event description. At the end of the day, all of this is just Dowty. Thematic entailments at the level of the event description matter for syntax. Non-thematic entailments do not, and the more one thinks about it, the harder it becomes to see why they should. There is no dedicated syntactic position for participants who wield a sharp weapon, ingest food or move their feet. These are not things that syntax cares about.

## 4. Further questions

This final section briefly discusses two further questions that this dissertation might bear on, namely, multiple θ-role assignment and entailments across alternations.

### *4.1. Multiple θ-role assignment*

A recurrent idea in the literature on θ-roles is that a single argument can be assigned several θ-roles (Broadwell, 1988; Jackendoff, 1987; Williams, 1994). This claim has been motivated on various grounds, but one of them has been the existence of variable linking. Here I briefly review two proposals for variable unaccusativity and direct causatives of unergatives, respectively, that rely on multiple θ-role assignment.

First, Chen (2024) has argued recently that variable unaccusativity constitutes evidence for an underlyingly transitive structure of unergatives (see also Hale and Keyser, 1993). Roughly, the sole argument of unergatives starts out in the internal argument position and receives a theme θ-role, but then moves to the external argument position where it additionally acquires an agent θ-role. The argument of unaccusatives, on the other hand, is a straightforward internal argument and is only assigned a theme θ-role. In verbs that vary

between unergative and unaccusative behavior, the single argument is thus assigned the same theme θ-role in both alternants, in addition to the agent θ-role in the unergative.

A similar proposal has been made by Ramchand (2008) for Hindi-Urdu causatives of unergatives. I gloss over the details of Ramchand's idiosyncratic approach to verbal syntax here, but in broad strokes, she argues that base unergatives have an initiator and an undergoer θ-role, which are obligatorily co-indexed. In the direct causative, the initiator role is instead associated with the external causer, while the causee retains the undergoer θ-role only. The subject of the intransitive and the object of the transitive thus receive the same undergoer θ-role, and the former additionally receives an initiator θ-role. I have argued against Ramchand's analysis of Hindi-Urdu on empirical grounds in Neu (to appear).

In a nutshell, both analyses aim to ensure that an argument receiving the same individual role is assigned the same θ-role. Specifically, Ramchand argues that the subject of the intransitive and the object of the transitive both undergo the dancing, and that this should be captured by a shared θ-role. But undergoing a dancing is not a thematic entailment. The 'dancing' part is a verb-specific entailment that should be irrelevant to syntax, and the 'undergoing' part is a vacuous placeholder creating the illusion that there is anything that undergoing a dancing, undergoing a breaking and undergoing a pushing have in common. On that note, a θ-role that has occasionally been proposed in the literature on argument structure is that of a doer (e.g., Cuervo, 2003; Massam, 2009; Ótott-Kovács, 2024). This is at best an unfortunate terminological choice. There is nothing intrinsically wrong about positing a separate role to describe a subset of the agent role – perhaps for arguments associated with sentience but not with volition –, but doing is not a thematic entailment, but a placeholder for a verb-specific entailment.

This is not a rebuttal of multiple θ-role assignment as a whole. Secondary predication, for instance, is a phenomenon that might very well require it. But often, multiple θ-roles are used specifically to account for shared non-thematic entailments across alternants, and I have argued that tracking non-thematic entailments is not what θ-roles do. None of this is even to say that a theory of roles that encode verb-specific entailments could not be successful. But it would be a very different theory, and it would need to be explicit.

### *4.2. Entailments across alternations*

Part of the job that θ-roles do for semanticists is to explain certain entailments between sentences containing the same verbs. To begin with the most simple case, (41a) entails (41b):

(41) a. Brutus stabbed Caesar yesterday.
b. Brutus stabbed Caesar.

This fact can be accounted for by positing the denotations in (42) (Parsons, 1990), which make the meaning of (41a) a proper subset of (41b):

(42) a. $\lambda e.stab(e) \wedge Subject(e)(Brutus) \wedge Object(e)(Caesar) \wedge yesterday(e)$
b. $\lambda e.stab(e) \wedge Subject(e)(Brutus) \wedge Object(e)(Caesar)$

By describing Brutus and Caesar as the subject and the object of the stabbing, respectively, we can track the identity of the stabber and the stabbee across the two sentences and ensure that (41a) does not entail that Caesar stabbed Brutus. So far, θ-roles are not needed.

However, surface-level descriptions such as subject and object are insufficient to capture entailments across valency-changing argument-structural alternations. In causative-inchoative alternations such as (43), it seems that the transitive entails the intransitive. However, this entailment does not go through if the arguments are individuated by referring to them as subjects and objects (44):

(43) a. Ahmad broke the glass.
b. The glass broke.

(44) a. $\lambda e.break(e) \wedge Subject(e)(Ahmad) \wedge Object(e)(the\ glass)$
b. $\lambda e.break(e) \wedge Subject(e)(the\ glass)$

Since *the glass* is the object of the transitive but the subject of the intransitive, the entailment in (44) is not valid. However, θ-roles do the trick (45):

(45) a. $\lambda e.break(e) \wedge Agent(e)(Ahmad) \wedge Patient(e)(the\ glass)$
b. $\lambda e.break(e) \wedge Patient(e)(the\ glass)$

If the arguments are individuated by their θ-roles, the entailment goes through as expected.

Cases of variable linking pose a problem for the idea that θ-roles can account for entailments across alternations. For Hindi-Urdu direct causatives of unergatives, does *Shama danced Rohan* entail *Rohan danced* (46)? For Greek stative passives, does *Mary is eaten* entail *Mary has eaten* (47)? For English symmetric predicates, does *Clementine kissed Joel* entail *Clementine and Joel kissed* (48)?

(46) a. $\lambda e.dance(e) \wedge Agent(e)(Shama) \wedge Patient(e)(Rohan)$
b. $\lambda e.dance(e) \wedge Agent(e)(Rohan)$

(47) a. $\lambda e.eat(e) \wedge Patient(e)(Mary)$
b. $\lambda e.eat(e) \wedge Agent(e)(Mary)$

(48) a. $\lambda e.kiss(e) \wedge Agent(e)(Clementine) \wedge Patient(e)(Joel)$
b. $\lambda e.kiss(e) \wedge Agent(e)(Clementine\ and\ Joel)$

According to the denotations given, they should not. Whether speakers feel they do is an empirical question, and I suspect that the answer is quite nuanced. If the alternation attributes largely the same non-thematic, but different thematic properties to an argument, which could be the case for (47), speakers might feel that one alternant does logically follow from the other, but is not pragmatically felicitous in the same contexts. If the alternation attributes somewhat different non-thematic properties to the argument, as arguably in (48), the entailment might be perceived as straightforwardly false (see Winter et al., 2016).

In general, using θ-roles to capture entailments across alternations presupposes that θ-roles establish the identity of an argument by ensuring that the same set of properties is predicated of it in each alternant. If variable linking is attested, this is not what θ-roles do. I do not know whether there is anything else that does. If there isn't, we might have to give up on the idea that entailments across alternations, to the extent that they are valid, easily and automatically fall out from the denotation of sentences. They might become a much hazier matter of world knowledge.

## References


Ahmed, T. (2010). The unaccusativity/unergativity distinction in Urdu. *Journal of South Asian Linguistics*, *3*(1), 3–22.

Akkuş, F. (2021). *(Implicit) argument introduction, Voice and causatives* [Doctoral dissertation, University of Pennsylvania].

Alexiadou, A., Anagnostopoulou, E., & Schäfer, F. (2015). *External arguments in transitivity alternations: A layering approach*. Oxford UP. https://doi.org/10.1093/acprof:oso/9780199571949.001.0001

Arad, M. (1998). Are unaccusatives aspectually characterized? (And other related questions). *MIT Working Papers in Linguistics*, *32*, 1–20.

Beavers, J., & Koontz-Garboden, A. (2020). *The roots of verbal meaning*. Oxford UP. https://doi.org/10.1093/oso/9780198855781.001.0001

Bhatt, R., & Embick, D. (2017). Causative derivations in Hindi-Urdu. *Indian Linguistics*, *78*(1–2), 93–151.

Biggs, A. (2019). Objects in motion verb phrases. *Glossa: a journal of general linguistics*, *4*(1), 98. https://doi.org/10.5334/gjgl.820

Biggs, A., & Embick, D. (2025). *English (stative) passive participles: Category, argument introduction, and interpretation* [Ms., McMaster University/University of Pennsylvania].

Borer, H. (2005). *Structuring sense vol. II: The normal course of events*. Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199263929.001.0001

Bosse, S., Bruening, B., & Yamada, M. (2012). Affected experiencers. *Natural Language & Linguistic Theory*, *30*, 1185–1230. https://doi.org/10.1007/s11049-012-9177-1

Broadwell, G. A. (1988). Multiple θ-role assignment in Choctaw. In *Syntax and semantics 21: Thematic relations* (pp. 113–127). Academic Press. https://doi.org/10.1163/9789004373211_008

Carlson, G. (1998). Thematic roles and the individuation of events. In S. Rothstein (Ed.), *Events and grammar* (pp. 35–51). Springer.

Chen, M. (2024). A decomposition analysis of agents. *Glossa: a journal of general linguistics*, *9*. https://doi.org/10.16995/glossa.10624

Cuervo, M. C. (2003). *Datives at large* [Doctoral dissertation, Massachusetts Institute of Technology].

Dowty, D. (1989). On the semantic content of the notion of 'thematic role'. In G. Chierchia, B. H. Partee, & R. Turner (Eds.), *Properties, types and meaning* (pp. 69–130). Kluwer.

Dowty, D. (1991). Thematic proto-roles and argument selection. *Language*, *67*, 547–619. https://doi.org/10.2307/415037

Embick, D. (2023). Smaller syntax for English stative passives: A first report. *Acta Linguistica Academica*, *70*, 285–316. https://doi.org/10.1556/2062.2023.00644

Everaert, M., & Anagnostopoulou, E. (1997). Thematic hierarchies and Binding Theory. Evidence from Greek. In *Empirical issues in formal syntax and semantics. Selected papers from the Colloque de Syntaxe et de Sémantique de Paris (CSSP 1995)* (pp. 43–59). Lang.

Fillmore, C. J. (1968). The case for case. In *Universals of linguistic theory* (pp. 1–88). Holt, Rinehart & Winston.

Folli, R., & Harley, H. (2006). On the licensing of causatives of directed motion: Waltzing Matilda all over. *Studia Linguistica*, *60*, 121–155. https://doi.org/10.1111/j.1467-9582.2006.00135.x

Gruber, J. S. (1965). *Studies in lexical relations* [Doctoral dissertation, Massachusetts Institute of Technology].

Hale, K., & Keyser, S. J. (1993). On argument structure and the lexical expression of syntactic relations. In *The view from building 20* (pp. 53–110). MIT Press.

Halle, M., & Marantz, A. (1993). Distributed Morphology and the pieces of inflection. In K. Hale & S. J. Keyser (Eds.), *The view from building 20* (pp. 111–176). MIT Press.

Halle, M., & Marantz, A. (1994). Some key features of Distributed Morphology. In *MIT working papers in linguistics* (pp. 275–288).

Hamo, A., & Meihami, S. (2024). *Stative passives in Ardalani Kurdish* [Handout from talk given at NELS 55, Yale].

Harris, A. (1981). *Georgian syntax*. Cambridge UP.

Haspelmath, M. (1994). Passive participles across languages. In P. J. Hopper (Ed.), *Voice: Form and function* (pp. 151–177). John Benjamins Publishing Company. https://doi.org/10.1075/tsl.27.08has

Holisky, D. A. (1987). The case of the intransitive subject in Tsova-Tush (Batsbi). *Lingua*, *71*, 102–132. https://doi.org/10.1016/0024-3841(87)90069-6

Jackendoff, R. (1987). The status of thematic relations in linguistic theory. *Linguistic Inquiry*, *18*, 369–411.

Kouneli, M. (2021). *High vs. Low external arguments: Evidence from Kipsigis* [Handout from talk given at SAIAL, Potsdam].

Kratzer, A. (1996). Severing the external argument from its verb. In J. Rooryck & L. Zaring (Eds.), *Phrase structure and the lexicon* (pp. 109–137). Kluwer/Springer. https://doi.org/10.1007/978-94-015-8617-7_5

Kratzer, A. (2002). *The event argument* [Ms., UMass Amherst].

Krejci, B. (2020). *Syntactic and semantic perspectives on first conjunct agreement in Russian* [Doctoral dissertation, Stanford University].

Kumaran, E. (2021). Two types of external argument in Karuk. In S. Kim, N. Kurtz, M. Hewett, & C. Kasper (Eds.), *Proceedings of the 56th Annual Meeting of the Chicago Linguistic Society* (pp. 257–268).

Legate, J. A. (2014). *Voice and v: Lessons from Acehnese*. MIT Press. https://doi.org/10.7551/mitpress/9780262028141.001.0001

Levin, B., & Rappaport Hovav, M. (1995). *Unaccusativity. At the syntax-lexical semantics interface*. MIT Press.

Marantz, A. (1981). *On the nature of grammatical relations* [Doctoral dissertation, Massachusetts Institute of Technology].

Massam, D. (2009). The structure of (un)ergatives. In S. Chung, D. Finer, I. Paul, & E. Potsdam (Eds.), *Proceedings of the 16th Meeting of the Austronesian Formal Linguistics Association* (pp. 125–135).

Mithun, M. (1984). The evolution of noun incorporation. *Language*, *60*, 847–894. https://doi.org/10.2307/413800

Myler, N. (2022). *Argument structure and morphology in Cochabamba Quechua (with occasional comparison with other Quechua varieties)* [Ms., Boston University].

Nash, L. (2021). Nonunitary structure of unergative verbs in Georgian. *Natural Language and Linguistic Theory*, *40*, 1201–1267. https://doi.org/10.1007/s11049-021-09529-8

Neu, E. (To appear). Direct causatives of unergatives: A view from variable unaccusativity. *Natural Language & Linguistic Theory*.

Nie, Y. (2020). *Licensing arguments* [Doctoral dissertation]. New York University.

Nishigauchi, T. (1984). Control and the thematic domain. *Language*, *60*, 215–250. https://doi.org/10.2307/413640

Ótott-Kovács, E. (2024). Argument-introducing pluractionals: An investigation of Kyrgyz and Kazakh assistives. *Natural Language & Linguistic Theory*, *42*, 1–65. https://doi.org/10.1007/s11049-023-09610-4

Paparounas, L. (2025). *Greek stative passives as small non-passives* [Ms., Université du Québec à Montréal].

Parsons, T. (1990). *Events in the semantics of English: A study in subatomic semantics*. MIT Press.

Perlmutter, D., & Postal, P. (1984). The 1-Advancement Exclusiveness Law. In D. Perlmutter & C. Rosen (Eds.), *Studies in Relational Grammar* (pp. 81–125). University of Chicago Press.

Pineda, A., & Berro, A. (2020). Hybrid intransitives in Basque. *Glossa: a journal of general linguistics*, *5*(1), 22. https://doi.org/10.5334/gjgl.824

Pylkkänen, L. (2008). *Introducing arguments*. MIT Press. https://doi.org/10.7551/mitpress/9780262162548.001.0001

Ramchand, G. C. (2008). *Verb meaning and the lexicon*. Cambridge UP. https://doi.org/10.1017/cbo9780511486319

Rissman, L., & Majid, A. (2019). Thematic roles: Core knowledge or linguistic construct? *26*, 1850–1869. https://doi.org/10.3758/s13423-019-01634-5

Sorace, A. (2000). Gradients in auxiliary selection with intransitive verbs. *Language*, *76*, 859–890. https://doi.org/10.2307/417202

Sorace, A. (2004). Gradience at the lexicon-syntax interface: Evidence from auxiliary selection and implications for unaccusativity. In A. Alexiadou, E. Anagnostopoulou, & M. Everaert (Eds.), *The unaccusativity puzzle* (pp. 243–268). Oxford UP. https://doi.org/10.1093/acprof:oso/9780199257652.003.0010

Sorace, A. (2011). Gradience in split intransitivity: The end of the Unaccusative Hypothesis? *Archivio Glottologico Italiano*, *96*(1), 67–86.

Tollan, R. (2018). Unergatives are different: Two types of transitivity in Samoan. *Glossa: a journal of general linguistics*, *35*, 1–41. https://doi.org/10.5334/gjgl.223

Tollan, R., & Massam, D. (2022). Licensing unergative objects in ergative languages: The view from Polynesian. *Syntax*, *25*, 242–275. https://doi.org/10.1111/synt.12232

Tollan, R., & Oxford, W. (2018). Voice-less unergatives: Evidence from Algonquian. In W. G. Bennett, L. Hracs, & D. R. Storoshenko (Eds.), *Proceedings of the 35th West Coast Conference on Formal Linguistics* (pp. 399–408). Cascadilla Proceedings Project.

von Stechow, A. (1996). The different readings of *wieder* 'again': A structural account. *Journal of Semantics*, *13*, 87–138.

Wilkins, W. (1988). Thematic structure and reflexivization. *Syntax and semantics*, *21*, 191–213. https://doi.org/10.1163/9789004373211_012

Williams, E. (1987). Implicit arguments, the binding theory, and control. *Natural Language & Linguistic Theory*, 151–180. https://doi.org/10.1007/bf00166581

Williams, E. (1994). *Thematic structure in syntax*. MIT Press.

Winter, Y. (2016). The reciprocity-symmetry generalization: Proto-roles and the organization of lexical meanings. *Empirical issues in syntax and semantics*, *11*, 245–274.

Winter, Y. (2018). Symmetric predicates and the semantics of reciprocal alternations. *Semantics and pragmatics*, *11*, 1–52. https://doi.org/10.3765/sp.11.1

Winter, Y., Westenbroek, C., Poortman, E. B., Berghoff, R., & Bischetti, L. (2016). *Lexical reciprocity as a typicality preference: Experimental evidence* [Ms., Utrecht University].