A Compositional Semantics for Spatial Perspective-Shifting Adjuncts

Karl Mulligan, Kyle Rawlins

1. Introduction

The meanings of some prepositions are sensitive to perspective. For instance, in the state of affairs described by (1), the ball is ambiguously located either by the back legs of the chair, or opposite the chair from the speaker's point of view.

(1)   The ball is behind the chair.

This ambiguity is due to the meaning of *behind*, and it is driven by the fact that its spatial meaning is dependent on a choice of perspective: whether to align the back direction with the intrinsic orientation of the chair, or with the speaker's line of sight. This choice can be characterized as a choice between either a chair-centered or a speaker-centered *frame of reference* (Levinson 1996), as illustrated in Figure 1a and Figure 1b.

The choice of perspective is usually left implicit, since convention and non-linguistic cues such as eye gaze or object prominence can result in predictably adopting one reference frame over another. However, perspective can be lexically fixed with a left-adjoined *from*-adjunct, as in (2):

(2)   From the door, the ball is behind the chair.

In this sentence, the spatial meaning of *behind* strictly uses the line of sight from the door to the chair, to the exclusion of the object-centered and observer-centered reference frame options presented in (1), effectively ignoring the intrinsic orientation of the chair. Barlew (2016) observes that sentences using a left-adjoined *from* phrase in this way seem to be able to fix perspective, without requiring reference to object orientation or to a particular perspective holder.

These data raise the following questions: how is perspective encoded in the semantics of *behind*, and how is the perspective-shifting behavior with *from* observed in (2) accomplished compositionally?

(a) Object-centered perspective   **(b)** Observer-centered perspective   **(c)** *from*-shifted perspective
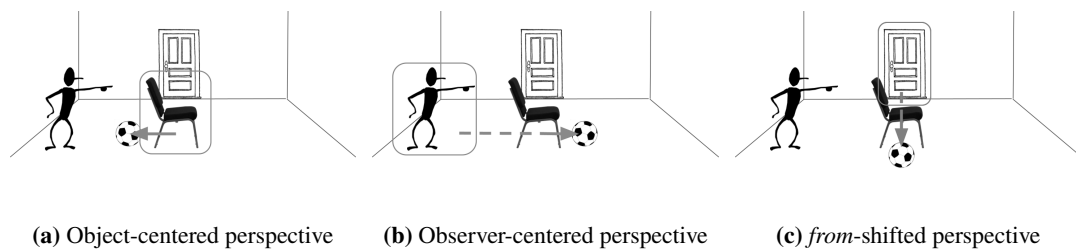
Figure 1: Various spatial configurations for the ball and chair in (1) and (2) owing to the perspective-sensitivity of *behind*.

In this paper, we follow prior work on the semantics of perspectival items in assuming that projective prepositions such as *behind* and *left of* have an implicit argument for perspective holder, which is normally assigned its value anaphorically. In the case of perspective shifting with *from*-adjuncts, we argue that this variable is instead bound, leading to obligatorily shifted perspective. We give a semantics for *from* which

---

* Karl Mulligan, Johns Hopkins University, `karl.mulligan@jhu.edu`

demonstrates this binding mechanism. Furthermore, we claim that the binding operator introduced by left-adjoined *from* is best analyzed as a generic quantifier over individuals, with the locational information in the *from*-adjunct forming the content of the restrictor.

The paper is organized as follows: in Section 2, we give a perspective-sensitive analysis of prepositions like *behind* using Vector Space Semantics (Zwarts 1997), showing how an implicit perspective center variable can be used to set frame of reference compositionally. In Section 3, we give a semantics for left-adjoined *from* which binds this variable with a generic pronoun, building on arguments made by Moltmann (2006, 2010) and Pearson (2013). Lastly, in Section 4, we show how our generic quantifier analysis supports other forms of perspective shifting, such as a conditional paraphrase of (2).

2. Semantics of projective prepositions

2.1. Spatial language and perspective-sensitivity

Spatial prepositions serve to locate one entity, called the figure object, based on the location of another, the ground object. For instance, in (1), *the ball* is the figure and *the chair* is the ground. We focus here on only a particular class of prepositions: the projective prepositions, which are terms like *in front of*, *behind*, (*to the*) *left/right of*. While the meaning of non-projective prepositions (e.g., *in*, *at*, *near*) is dependent only on the location of the two objects, the meaning of projective prepositions is additionally dependent on the direction from the ground object.

Importantly, the direction from the ground object for *behind* is not a hard-and-fast property of the physical environment like, say, the cardinal direction SOUTH, but is instead subject to context-dependence. Which direction exactly is meant by a given preposition is a matter of contextually determining frame of reference, or a coordinate system for representing directions in space. The center or 'origo' of this coordinate system is what defines a frame of reference: in an observer-centered (also called 'relative' or 'deictic') reference frame, what counts as the 'back' axis for the meaning of *behind the chair* depends on the coordinate system centered at an observer viewing the ground object, and is defined by her own front and back; in an object-centered (or 'intrinsic') reference frame, the center is the ground object itself, so that *behind the chair* means 'back' relative to the axes that describe the conventional front and back of a chair. We call the entity defining frame of reference the *perspective center*.

A hallmark of perspectival phenomena is their shiftability (Bylinina et al. 2015, Anderson 2020). While other context-sensitive expressions like the first-person indexical *I* can change their reference depending on who the speaker is, perspective-sensitive expressions exhibit different meanings in a wider variety of contexts, such as when embedded under attitude verbs. In a relative frame of reference, for example, the perspective center is not limited to the speaker:

(3) *Context: Sam and Lisa are standing in a room with a circular table. Antoine was in the room five minutes ago, but is not currently.*

    a. Sam: The ball is behind the table.   SPEAKER (Sam)
    b. Sam (to Lisa): Put the ball behind the table.   LISTENER (Lisa)
    c. Sam (to Lisa): Antoine thinks the ball is behind the table.   ATTITUDE-HOLDER (Antoine)
    d. Sam: Antoine found the ball behind the table.   (salient) SUBJECT (Antoine)

This wide variability in what entities may count as a licit perspective center ultimately motivates an account which treats perspective center as an anaphorically-determined free variable in our analysis of *behind*.

2.2. Denotation for behind

To give our analysis for *behind*, we will use Vector Space Semantics (VSS), a compositional framework for the meaning of spatial prepositions (Zwarts 1997). Originally designed to account for modified PPs such as *two feet above the table*, VSS enriches the standard ontology of semantic types e and t with types derived from vector space: type p for points in space, and type v for vectors. To get in and out of vector space, the following two functions are needed: the *location function* LOC ($D_e \rightarrow D_{\langle p,t \rangle}$), which gives the

set of points in space that an entity occupies[1]; and the *antilocation function* LOC^{-1} ($D_{\langle v,t \rangle} \rightarrow D_{\langle e,t \rangle}$), which lifts a set of vectors from P' into a predicate for the entire PP, returning true of an entity just in case it occupies the space characterized by those vector endpoints. For example, the P'-level denotation for *outside* is:

(4) $[\![\text{outside}]\!] := \lambda A_{pt}.\lambda \mathbf{v}_v. \ \text{ext}(\mathbf{v}, A)$ (Zwarts & Winter 2000)

where A is the area (given by LOC) occupied by *outside*'s DP complement, and 'ext' denotes vectors *externally closest* to A, i.e., the minimal length vectors whose starting points are at the boundary of A and are directed outward from it.

For projective preposition *behind*, we start with the same denotation as *outside*: we are provided the set of points A defining the region of the ground object, and we pick out vectors which extend outward from A. What *behind* further stipulates is that those vectors must have a component in the opposite direction of the positive horizontal (*front*) direction[2]:

(5) $[\![\text{behind}_{\text{zw2000}}]\!] := \lambda A_{pt}.\lambda \mathbf{v}_v. \ \text{ext}(\mathbf{v}, A) \wedge comp(-front, \mathbf{v}) > 0$ (Zwarts & Winter 2000)

The denotations for *in front of*, *right of*, *above*, etc. are defined analogously, but with different directions for the component: *front*, *right*, *up*. How are these directions defined? Zwarts and Winter correctly isolate the relevant portion of the denotation needed to set frame of reference, acknowledging that axes constants like *front* are "pragmatically determined" (p. 181), but leave the discussion at that.

To enrich this denotation, Bohnemeyer (2012) posits an indexical account of the axis used by projective prepositions using the VSS framework. Bohnemeyer replaces the placeholder axis *front* with an axis function which takes an indexical anchor entity, something like $front(c_a)$ where c_a is a perspectival anchor fixed by the interpretation function as part of a Kaplan-style context tuple, included among other variables like speaker and utterance time. Here is an indexical version of *behind*:

(6) $[\![\text{behind}_{\text{idxl}}]\!]^c := \lambda A_{pt}.\lambda \mathbf{v}_v. \ \text{ext}(\mathbf{v}, A) \wedge comp(-front(c_a), \mathbf{v}) > 0$ (based on Bohnemeyer (2012))

However, as Barlew (2016) observes, this function incorrectly returns a vector with same direction when the perspectival anchor is co-located with the ground object in an intrinsic FoR (e.g., "behind me" when $c_a = c_{\text{speaker}}$) as when the anchor is used to set a relative FoR ("behind the tree, relative to me"). Naming an indexical anchor entity, then, is insufficient for disambiguating between frame of reference.

2.2.1. Our denotation for behind

As previously stated, the shiftability of the perspective center make an anaphoric analysis of the axis variable in *behind* the most appealing. In our denotation of *behind*, we add an unpronounced, implicit argument z which corresponds to the perspective center, and compute a context-sensitive value for the axis variable in terms of it. To be able to derive directional information from various types of perspective centers in a single, unified denotation, some new vector terminology and helper functions are required.

Vectors can be classified as object-internal *axis vectors* and object-external, projective *place vectors* (Zwarts 2003). These distinct vector types represent distinct kinds of directional information assumed to be available in the environment: the axis vectors are associated with object orientations, whereas the place vectors are associated with viewpoints or projections of axial information. Below, we define a few functions used to extract this information for use in the semantics of projective prepositions.

1. We assume an *orientation* function ORIENT ($D_e \rightarrow D_{\langle v,t \rangle}$) which supplies the primary or *front* axis vector of any contextually salient oriented entity, assumed to be the part of how the human visual system recognizes and represents objects[3] (Marr & Nishihara 1978). If the entity lacks internal axes

---

[1] For simplicity, in our denotations we include LOC directly in the denotations of prepositions rather than breaking out the region variable A as a separate argument, allowing the preposition to compose directly with its entity arguments.

[2] This too is a simplification from the actual denotation in Zwarts & Winter (2000), which compares the component of a vector to its projection on a perpendicular axis, giving a region with more tolerance for entities which are not directly above A.

[3] See Landau & Jackendoff (1993) for more discussion of how the use of axes is reflected in language.

(e.g., a tree), ORIENT returns the set of all possible axis vectors for that entity.

2. We also assume an *axis-to-place* function A2P ($D_v \rightarrow D_v$), which gives an (object-external) place vector extending outward from an object in the direction of its (object-internal) axis vector (this operation essentially amounts to positive scalar multiplication). For an oriented entity, the resulting vector may be interpreted as the "viewpoint" associated with its orientation.

3. Lastly, we assume a function AXIZE which, given perspectival center z and a ground object x_e:
   - if $z = x$ (i.e., in a ground object-centered reference frame), returns the object's internal axis ORIENT(x) unmodified;
   - otherwise (e.g., in a speaker-centered reference frame where $z = g(i)$), returns a new, axis vector in the opposite direction of the "point of view" $g(i)$, effectively (re)setting the internal *front* axis vector of x.

This function accounts for the difference in interpretation of *behind me* and *behind...*(relative to me)[4], as explained when discussing the empirical gaps of (6). In Figure 1, the boldface arrows correspond to the internal axis vectors returned by AXIZE.

Putting it all together, we arrive at the following denotation; z will receive a value anaphorically via the assignment function g (see Fig 2):

(7) $[\![behind]\!]^g = \lambda z_e \, . \, \lambda x_e \, . \, \lambda \vec{v}_{\langle v,t \rangle} \, . \, \text{ext}(\vec{v}, \text{LOC}(x)) \wedge comp(\vec{v}, \text{A2P}(-\text{AXIZE}(\text{ORIENT}(z), x))) > 0$,
         defined only if end-point(A2P($g(i)$)) ∈ LOC(x) ∧ ORIENT(z) returns a unique vector.

We have also added definedness conditions which ensure that only vectors from entities that end at the ground object are able to serve as perspective center, which limits the oriented entities in the discourse that can serve as a reference frame-setter to only those whose viewpoint actively includes the ground object. As such, the definedness condition predicts an infelicity for an utterance of *The ball is behind the tree* when neither the speaker nor addressee is oriented toward the tree, since for no i does $g(i)$ give a unique vector ending at the intrinsically non-oriented ground object.

Finally, we also assume that prepositions like *behind* that have a free perspectival variable license predicate abstraction, making the clause a type-$\langle e,t \rangle$ property, thus allowing the perspective center variable to be bound at different scopes (though our proposal is compatible with other means of binding such variables). This provides, for instance, a technical means for obtaining shifted readings when *behind* is under an attitude verb, as in (3c), but also ultimately for perspective-shifting with *from*-adjuncts, as we explain below.

3. Shifting perspective with *from*-adjuncts

We turn now to the semantics of perspective-shifting *from*, which we claim introduces a quantifier that binds the perspective center implicit argument of prepositions like *behind* with a generic pronoun.

3.1. Perspective-shifting from-adjuncts

The earliest observation of the perspective-shifting properties of *from*-adjuncts to our knowledge comes from Barlew (2016). Based on an analysis of locative morphemes in the Bantu language Mushunguli, Barlew suggests that, to set frame of reference, projective prepositions need only use an abstract Point of View (POV), or ⟨location, orientation⟩ pair, rather than make reference to a particular perspective-holding entity. As an example of a construction which instantiates an abstract POV, Barlew gives (8):

(8)   A: Is the pond just in front of the the large oak tree?
        B: Yes, that's right. And from that hill over there, the pond is to the right of the large oak tree.
                                                                              (Barlew 2016: ex. (45))

---

[4] This function may be parameterized to account for cross-linguistic variability in relative reference frame processing; for instance, relative reference frames for *front* and *back* are reflected relative to the perspective holder in English, but in Hausa, they are transposed (Hill 1982, Shusterman & Li 2016).

For the interpretation of *to the right* in (8), neither of the interlocutor's perspectives are used; nor is an intrinsic orientation for either the hill or the tree, because these entities lack an intrinsic orientation. Rather, the *from*-adjunct *from that hill over there* introduces a direction which is used to set a relative frame of reference for *to the right of the tree*. How is this frame of reference-switching maneuver accomplished compositionally?

3.2. A solution with generic quantifiers

While we agree with Barlew's observation that a location and orientation are sufficient for establishing a relative frame of reference, we believe that this idea is not incompatible with positing a hypothetical observer entity as an implicit argument for projective prepositions. There is no hint of implicit personhood anywhere in (8), and yet, it is very difficult for one to compute the intended meaning without first imagining what the viewpoint *would be* if, hypothetically, one were standing on the hill. Moreover, computing that the relevant orientation in (8) is a vector pointing from the hill out toward the tree requires the non-trivial assumption that the POV orientation should always "face" the ground object, a requirement which is commonsensical for a hypothetical observer but does not transparently arise from the purely geometric VSS denotations for either *from* or *behind* alone. These observations motivate an account of perspective-shifting *from*-adjuncts which involves (at least at a hypothetical, technical level) entities which supply a contextually-restricted observer-centered frame of reference, to the exclusion of other anaphorically determined reference frames.

Given these considerations, we propose that *from*-adjuncts interact with the perspectival component of projective prepositions by introducing an operator which binds the perspectival center variable of a projective preposition in its scope. We claim that this operator is the generic quantifier GEN, taking the familiar tripartite quantificational structure [OPERATOR *var* : RESTRICTOR] SCOPE. The variable in question is a *simulated* entity bound by the generic quantifier, in the sense of Moltmann (2006), where an agent identifies with the properties of the simulated individual as if it were first-personal experience. As such, the simulated individual functions as a hypothetical observer, with the DP of the *from*-adjunct *the door* supplying that observer's location as part of the restriction. The *from*-adjunct then binds the perspective center variable of any projective prepositions within its scope, fixing frame of reference to that of the simulated observer. (9) gives our proposed denotation for left-adjoined perspective-shifting *from*:

(9) $[\![\text{from}]\!]^{g,c} = \lambda x_e \, . \, \lambda p_{\langle e,t \rangle} \, . \, [\text{GEN } z : [\![\text{at}]\!](z,x)] \, p(z)$
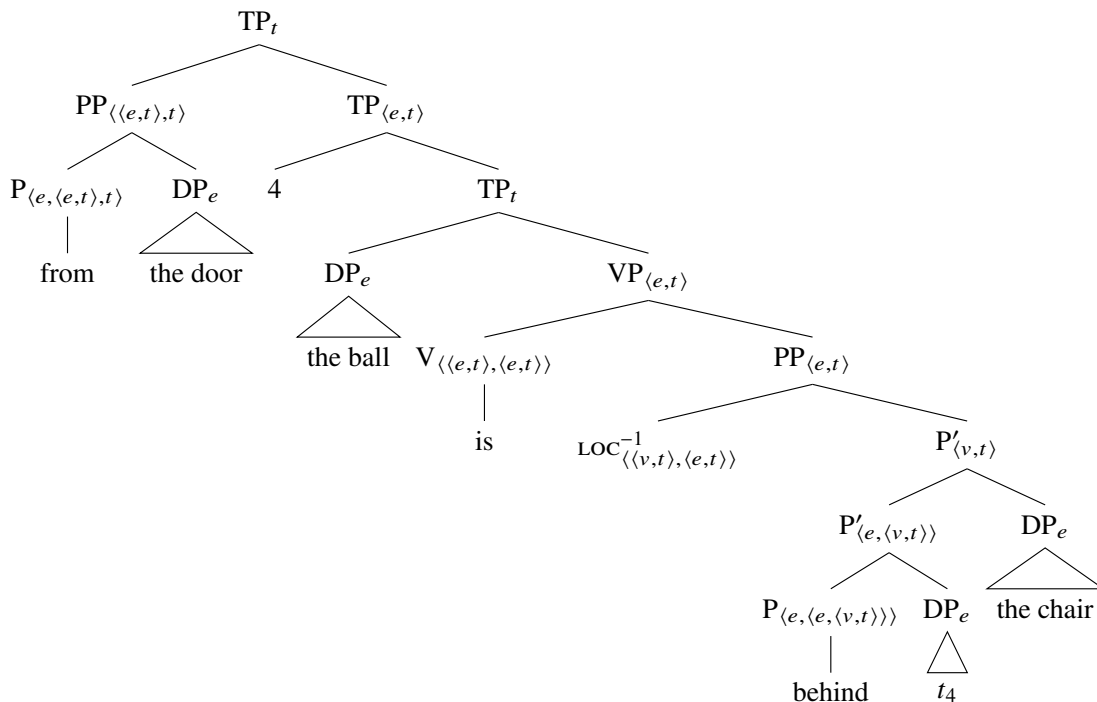
Applying this to our motivating sentence (2) from the introduction, repeated below as (10), we get the following:

(10)   From the door, the ball is behind the chair.
        a.   [GEN z: z is at the door] λz [the ball is behind(z) the chair].
        b.   [$\forall z : \text{at}(z, \text{the-door}) \wedge C(z) \wedge I(c_{\text{speaker}}, z)$] behind(the-ball, the-chair, z)

We assume following Krifka et al. (1995) that GEN is a quasi-universal operator; that is, it generally has universal force, but can handle exceptional individuals with the use of a contextual restriction function C. In our analysis, C ensures that only generic individuals who are oriented toward the ground object, have an unobstructed, normal line of sight, etc. are able to bind the perspective center variable for projective prepositions. The other restriction is the identification relation I, based on Pearson (2013), which is in turn adapted from Moltmann (2006). This relation holds if an agent, such as the speaker, identifies with the bound generic individual, thus limiting perspectives for z to those which can be simulated as first-person experience by the agent.

Why appeal to a generic quantifier? Our motivating idea is that computing the spatial meaning of a sentence with *from*-shifted perspective involves putting oneself in the shoes of some observer—*any* non-particular observer[5]—who is present at the shifted location and able to set their own frame of reference. Reading off the logical form in (10b) seems to capture this intuition nicely: "for any given individual z at the door, the ball is behind the chair from z's perspective." As long as the interlocutors are able to imagine

---

[5] In fact, generic quantifiers have also been used to give a convincing treatment of free-choice *any* by Dayal (1998).

TP$_t$

PP$_{\langle\langle e,t\rangle,t\rangle}$ TP$_{\langle e,t\rangle}$

P$_{\langle e,\langle e,t\rangle,t\rangle}$ DP$_e$ 4   TP$_t$

from   the door   DP$_e$ VP$_{\langle e,t\rangle}$

the ball   V$_{\langle\langle e,t\rangle,\langle e,t\rangle\rangle}$ PP$_{\langle e,t\rangle}$

is   LOC$^{-1}_{\langle\langle v,t\rangle,\langle e,t\rangle\rangle}$ P$'_{\langle v,t\rangle}$

P$'_{\langle e,\langle v,t\rangle\rangle}$ DP$_e$

P$_{\langle e,\langle e,\langle v,t\rangle\rangle\rangle}$ DP$_e$ the chair

behind t_4

Figure 2: A tree for (10). The perspective center variable for *behind* is bound by the generic quantifier introduced by *from*-adjunct *from the door*.

or simulate being such an individual, a frame of reference centered at that individual's location is made accessible.

Other authors have explored the apparent connection between generics and perspectival phenomena. In her analysis of generic *one* (and its empty counterpart, arbitrary PRO), Moltmann (2006) notes properties attributed to generic individuals are interpreted using first-personal inference. This is a desirable property of the generic quantifier, since the frame of reference provided by *from*-adjuncts behaves like a relative frame of reference, fully compatible with a strategy that uses a simulated, displaced first person as the perspective center.

More immediately relevant to our account, Pearson (2013) gives an analysis of predicate of personal taste *tasty* involving generic quantification over individuals, inspired by elements of Moltmann (2010). For Pearson, as in our account, the implicit "judge" (our perspective center) parameter is not directly given its value indexically, but is first bound by an internal generic quantifier, deploying the use of the identification function I to restrict the perspectival component to one which is obtained through first-personal experience indirectly via simulation of a generic individual. Our analysis shares many of the same features, further supporting the notion that genericity is an important ingredient in the interpretation of shiftable perspectival items.

3.3. Syntactic properties of perspective-shifting from-adjuncts

Based on the relatively simple [P DP] surface syntax of *from*-adjuncts, it is worth stopping to consider whether this quantificational analysis is unnecessarily complex. However, there are a couple properties of the special way in which *from* is used (i.e., left-adjoined, serving as a restrictor for the domain a perspectival item) which support a quantificational proposal.

For one, the *from*-adjunct must appear in the left periphery to act as a perspective shifter. This appears to be a privileged position, as far as PP modifiers are concerned, in that only in the left periphery can they restrict semantically unspecified domains. To illustrate, consider the following pair from Maienborn (2001):

(11) a. Eva signed the contract in Argentina.   (event modifier)
      b. In Argentina, Eva is still very popular.   (frame-setting modifier)

Whereas (11a) implies *Eva signed the contract*, (11b) does not imply the general proposition *Eva is still very popular*. The latter use of the locative *in Argentina* serves as a domain restrictor, showing that PPs used in this position can often have quantificational force.

In addition, the syntax and semantics of left-adjoined *from* bear some resemblance to a correlative structure, in which a subordinate clause is left-adjoined to the matrix clause. Unlike noun modification with relative clauses, in correlatives, the left-adjoined clause acts behaves as a quantifier, binding variables within its scope (Srivastav 1991, Dayal 1995). For example, in Hindi, (12) shows how a left-adjoined free relative binds the proform demonstrative DEM in the main clause:

(12) **jo   laRkii khaRii hai** vo   lambii hai .
      [REL girl   standing is] DEM tall   is .
      'The girl who is standing is tall.'   (Srivastav 1991: ex. 3a)

While Hindi also allows embedded and right-adjoined relatives, these are argued to be syntactically and semantically distinct from the correlative; only in the left-adjoined position does the free relative take on quantificational force, allowing it to bind the demonstrative pronoun in the matrix clause. Similarly, the *from*-adjunct introduces a restriction on the location of possible generic individuals, binding the implicit perspective center pronoun in *behind*.

4. Conditional perspective shifters

So far, we have shown that left-adjoined *from*-adjuncts can serve as generic quantifiers, shifting perspective of any projective prepositions in their scope. Are there other kinds of adjuncts that serve the same perspective-shifting purpose?

One place to start is the following general property of quantified expressions: elements in the restrictor, such as domain or contextual restrictions, can often be rephrased as *if*-clauses (Partee 1995). Generic sentences are no exception:

(13) a. Jays chirp.
      b. If something is a jay, it chirps.

We can apply this pattern to our running *from*-adjunct example sentence (2), substituting American English generic/impersonal *you*[6] for the variable in the antecedent of the conditional. This results in the following, fairly natural utterance:

(14)   If you're at the door, the ball is behind the chair.

We claim that (14) has the same truth conditions as its generic equivalent, (2). This is expected on our proposal, since the elements of the restrictor and scope (on a tripartite analysis of conditionals), assuming an equivalence between generic *you* and the bound generic variable. At a discourse level, it also appears to serve the same perspective-shifting purpose.

Indeed, the conditional perspective shifter phrase *if you're at the door* in (14) seems to serve an *entirely* discourse-level function. Unlike normal, hypothetical conditionals, for which the truth of the consequent is dependent on the antecedent (e.g., *If it's raining, I bring an umbrella*), (14) unconditionally asserts the consequent—that there is a particular spatial configuration involving the ball and the chair—with the antecedent only serving to fix the relevant perspective. Conditionals having this property are typically called biscuit (or speech act) conditionals[7] (Austin 1958). For instance, uttering *If you're thirsty, there's beer in the fridge* has the force of asserting that there is beer in the fridge, regardless of whether or not you are thirsty. To make this contrast clear, compare (14) with a similar conditional involving Frida and her shy cat, Gnocchi:

---

[6] While Moltmann (2010) focuses on generic *one*, generic *you* is assumed to follow the same "semantic strategy," making everything that has been said so far about her account perfectly applicable here as well.

[7] See Rawlins (2020) for a recent overview.

(15)   If Frida is at the door, Gnocchi is behind the chair.

Here, there *is* a conditional dependence between the consequent and the antecedent, so (15) generates a hypothetical reading. In fact, the proper name Frida can be replaced by all types of other NPs: definite, indefinite, second-personal *you*, even *no one*; in every case, the reading that results is hypothetical. Of all these quantificational environments, only the use of a generic pronoun produces a biscuit reading.

Why is the biscuit interpretation exclusively available with generic *you* in the antecedent? The fact that GEN is an type-unselective binder of variables may be useful here. For perspective shifting with *from*-adjuncts, we claimed that the generic quantifier ranged over individuals of type *e*. So far, however, we have made no mention of world or situation variables, which are often deployed for many fairly quotidian generic sentences:

(16)   a.   Mary smokes when she comes home.   (Krifka et al. 1995)
        b.   [GEN $s, x : x$ = Mary $\land$ x comes home in s] x smokes in s

The conditional paraphrase of (16), given below, produces an unambiguously hypothetical reading. This is to be expected, assuming a standard possible worlds semantics for conditionals, with f as an accessibility flavor function:

(17)   a.   If Mary comes home, she smokes.
        b. $\forall w' : f(w_0)(w') \land [\![\text{Mary comes home}]\!](w') \rightarrow [\![\text{Mary smokes}]\!](w')$

The conditional dependence between the two clauses is made transparent due to quantification over the possible world variable w'. We assume that quantification over worlds, even when using a generic operator, is what is most directly responsible for generating hypothetical readings. Take the following intensional analysis for (15), which uses GEN to capture its habitual interpretation:

(18)   [GEN $w', x : x$ = Frida $\land$ x is at the door in w'] Gnocchi is behind(x) the chair in w'

In contrast, on our account of left-adjoined (9), generic quantification for perspective shifting ranges strictly over individuals. We propose that likewise, for conditional perspective shifting with generic *you*, GEN does not have to bind world/situation variables, and when it does not, the relevant conditionals receive biscuit-like readings where only generic entities are bound.

(19)   a.   Hypothetical conditional:
              $\forall w', z : f(w_0)(w') \land [\![z \text{ is at the door}]\!](w') \rightarrow [\![\text{the ball is behind}(z) \text{ the chair}]\!](w')$
        b.   Perspective-shifting conditional:
              $\forall z : [\![z \text{ is at the door}]\!](w_0) \rightarrow [\![\text{the ball is behind}(z) \text{ the chair}]\!](w_0)$

As with biscuit conditionals more generally, conditional perspective shifting with generic *you* has a distinctly discourse-level purpose: here, to explicitly set an alternative frame of reference for one's interlocutor. As a consequence, the truth conditions regarding the actual spatial configuration of the objects in the scope remain unaffected by the presence of the *if*-clause.[8]

In some ways, this separation of possible worlds and individuals is reminiscent of Lewis's centered worlds, which at their simplest consist of a ⟨world, individual⟩ pair, and are used to capture *de se* attitudes (Lewis 1979). A perspective-shifting conditional is one that shifts the perspectival center via the GEN operator by ranging only over individuals.

5. Conclusion

We have given an analysis for projective prepositions like *behind* that shows how perspectival content is encoded, and have explored two compositional strategies used to shift perspective: left-adjoined *from*-

---

[8] We remain neutral here on what the right general account of biscuit conditionals is: on the proposal as it stands, it remains open whether perspective-shifting conditionals should be treated uniformly with other biscuit conditionals (and the lack of world-binding follows from the fact that they are biscuit conditionals), or as "genuine" conditionals that entail their consequent by a quirk of the generic operator.

adjuncts, and conditional perspective shifting. For both strategies, we claim that a generic quantifier binds an implicit perspectival center variable that is used to set frame of reference. Our account demonstrates that elements of other theories of genericity, such as first-personal simulation, are useful for characterizing the truth conditions of shifted perspective in projective prepositions. A number of open questions remain about this means of perspectival shifting and the connection to biscuit conditionals more broadly: generic/impersonal *you* shows up routinely in all sorts of biscuit conditionals, raising the possibility of generalizing this analysis beyond spatial perspective. We leave this question for future work.

References

Anderson, Carolyn Jane. 2020. *Shifting the Perspectival Landscape*. University of Massachusetts Amherst dissertation.

Austin, John. 1958. Ifs and cans. *Journal of Symbolic Logic* 23(1). 74–75.

Barlew, Jefferson. 2016. Point of View and the Semantics of Spatial Expressions in Mushunguli and English. *Journal of Semantics* 33(2). 215–267.

Bohnemeyer, Jürgen. 2012. A vector space semantics for reference frames in Yucatec. In *SULA 6 and SULA-Bar*. Amherst: Graduate Linguistics Student Association, University of Massachusetts.

Bylinina, Lisa, Eric McCready & Yasutada Sudo. 2015. Notes on perspective-sensitivity. *Donum Semanticum*. 67–79.

Dayal, Veneeta. 1995. Quantification in correlatives. In *Quantification in natural languages*, 179–205. Springer.

Dayal, Veneeta. 1998. Any as Inherently Modal. *Linguistics and Philosophy* 21(5). 433–476.

Hill, Clifford. 1982. Up/down, front/back, left/right. A contrastive study of Hausa and English. *Here and There*. 13–42.

Krifka, Manfred, Francis Jeffry Pelletier, Gregory Carlson, Alice Ter Meulen, Gennaro Chierchia & Godehard Link. 1995. Genericity: an introduction. In Gregory Carlson & Francis Jeffry Pelletier (eds.), *The generic book*. Chicago.

Landau, Barbara & Ray Jackendoff. 1993. "What" and "where" in spatial language and spatial cognition. *Behavioral and Brain Sciences* 16(02). 217.

Levinson, Stephen. 1996. Frames of reference and Molyneux's question: Crosslinguistic evidence. In *Language and space*, 109–169.

Lewis, David. 1979. Attitudes De Dicto and De Se. *The Philosophical Review* 88(4). 513.

Maienborn, Claudia. 2001. On the Position and Interpretation of Locative Modifiers. *Natural Language Semantics* 9(2). 191–240.

Marr, David & Herbert Keith Nishihara. 1978. Representation and recognition of the spatial organization of three-dimensional shapes. *Proceedings of the Royal Society of London. Series B. Biological Sciences* 200(1140). 269–294.

Moltmann, Friederike. 2006. Generic one, arbitrary PRO, and the first person. *Natural Language Semantics* 14(3). 257–281.

Moltmann, Friederike. 2010. Generalizing Detached Self-Reference and the Semantics of Generic One. *Mind & Language* 25(4). 440–473.

Partee, Barbara H. 1995. Quantificational Structures and Compositionality. In Emmon Bach, Eloise Jelinek, Angelika Kratzer & Barbara H. Partee (eds.), *Quantification in Natural Languages* (Studies in Linguistics and Philosophy), 541–601. Dordrecht: Springer Netherlands.

Pearson, Hazel. 2013. A Judge-Free Semantics for Predicates of Personal Taste. *Journal of Semantics* 30(1). 103–154.

Rawlins, Kyle. 2020. Biscuit Conditionals. In Daniel Gutzmann, Lisa Matthewson, Cécile Meier, Hotze Rullmann & Thomas Zimmermann (eds.), *The Wiley Blackwell Companion to Semantics*, 1–26. John Wiley & Sons, Ltd.

Shusterman, Anna & Peggy Li. 2016. Frames of reference in spatial language acquisition. *Cognitive Psychology* 88. 115–161.

Srivastav, Veneeta. 1991. The Syntax and Semantics of Correlatives. *Natural Language & Linguistic Theory* 9(4). 637–686.

Zwarts, Joost. 1997. Vectors as relative positions: a compositional semantics of modified PPs. *Journal of Semantics* 14(1). 57–86.

Zwarts, Joost. 2003. Vectors across Spatial Domains: From Place to Size, Orientation, Shape, and Parts. In Emile van der Zee & Jon Slack (eds.), *Representing Direction in Language and Space*, 39–68. Oxford University Press.

Zwarts, Joost & Yoad Winter. 2000. Vector Space Semantics: A Model-Theoretic Analysis of Locative Prepositions. *Journal of Logic, Language and Information* 9(2). 169–211.