A force-domain analysis of English *have*

Jisu Sheen

Advisor: María Mercedes Piñango

Department of Linguistics

Yale University

This thesis is submitted for the degree of *Bachelor of Arts*, 2019

Acknowledgements

I would like to thank

my advisor: María Mercedes Piñango, for being my favorite architect,

new and old students in the Language and Brain lab: Martín Fuchs for singing with me in lab on bad days, Sara Sánchez-Alonso for teaching me, Nanyan Wu for inspiring me, James Lin for being there for the birth and early childhood of this dataset, and Andy Zhang for changing my life forever,

Hadas Kotek, Raffaella Zanuttini, Jay Jamal Pittman, James Lin (again!), Magda Andrews-Hoke, Noah Macey, Rose Mintzer-Sweeney, Anelisa Fergus, and Will Merrill,

my first TFs in the department: Luke Lindemann, Martín Fuchs (again!), and Matt Tyler,

the House of Fermentation: skate girl Kayley Estoesta, Madalina "'s Got Jokes" Ene, Dr. Patrick Han, Xavier Ruiz, our good friend Derek, and newcomers Stella "Then Do?" Xu and Sofia Bliek—this thesis would not exist if not for tubeworld; it would be nothing without your love—

the House of Plants and Mammals: my brave unni Titania Nguyen and strong-hearted Chairin Kim,

Victoria Wang, Hayun Cho, Brook Hsu, Marina Tinone, Justine Xu, Seth Wallace,

my blood relatives, for doing intergenerational work,

my former advisor in the computer science department for telling me, "Whatever you do, don't major in linguistics," and

my statistically significant other.

Chapters

Abstract

The role of context in what is referred to as locative *have* and possessive *have* in English has been found (Zhang 2018) to be important in comprehension; here we look at a dataset of produced contexts in order to determine the role of context in the production of *have*. This essay is broken down into three chapters: a description of a dataset whose behavior cannot be fully explained by the analyses we currently have available, a discussion of conceptual structure to help us develop an analysis that has the power to fill in these gaps, and a return to the dataset in order to test predictions based on the analysis from the middle section. Finally, we run through examples from the dataset with specific explanations stemming from the analysis and explore a possible task-dependent aspect of the results of this dataset. The main contribution I make with this project is an updated LCS for *have*, developed in section 3.1, and a test of how this plays out in subject-produced elicitations.

Preliminary notes:

Though all of these subject responses were gathered verbally, the transcriber (James Lin, Nanyan Wu, or Jisu Sheen) added punctuation for better readability. For this reason, no punctuation shown here should be taken as an absolute interpretation of the subject's intended utterance.

All elicited contexts are shown in their entirety, from the first thing the subject uttered after being presented with the stimulus to the last thing they uttered before the presentation of the next stimulus item.

In these quotations, P stands for participant (the subject), and E stands for experimenter. Typically, the experimenter did not speak throughout the task except to provide each stimulus item and prompt, but they asked for clarification if there was ambiguity in the subject's response.

Chapter 1: The Data

1.1 What we know about *have*

Cross-linguistic accounts of *have* link the conceptual structure of locative *have* and possessive *have* (Zhang 2018; Deo 2015). These contrast with accounts of *have* that place special restrictions on the acceptability of locative *have*. For example, Myler (2016) writes that, in the case of locative *have*, a "well-known constraint on this construction is that it requires a locative small clause…containing a pronoun anaphoric to the subject of *have*. Absent such a pronoun, the result is either ungrammaticality or a non-locative interpretation for the sentence" (p. 263). As examples of locative *have* with the supposedly required locative small clause, Myler (2016) provides the following (originally numbered (22a-c)):

(1) This tree has nests.
(2) I have a cockroach on my head.
(3) The stadium has two pubs flanking it.

Myler (2016) goes on to state that the following sentence (originally numbered (23a)) is "rejected as anomalous by native speakers" (p.264):

(4) *This tree has nests.

However, according to Zhang (2018), constructions like these with no small clause attached may indeed be deemed acceptable by native speakers, with acceptability modulated by the context as well as individual differences between the speakers themselves. That is, given the right context and a sufficient level of context-sensitivity in the comprehender/hearer, sentences like (4) will be considered acceptable.

As we look at the dataset laid out in the next section (1.2), we will first use it as a way to answer curiosities about how predictions like Myler (2016)'s and Zhang (2018)'s play out in production. Given a sentence like (4), can subjects produce a context that licenses this locative *have*? What kinds of contexts are these?

1.2 The study

55 right-handed native speakers of English (28 women, 24 men, and 3 nonbinary people), between 18-30 years old, were recruited from the Yale community to perform a context elicitation task using English *have*. Given a series of fourteen stimulus items that connected two entities with *have*, in the form:

"The NP_1 **has** a NP_2 that is ADJ,"

where the meanings of the two NPs were chosen such that a containment or ownership relation between them would be unexpected, the subjects were asked to provide a context: Where might they have heard this? What kind of person would say this? And in what situation might someone say it?

Subjects' responses were audio-recorded and transcribed. The following are three examples of stimuli and responses from the data:

(5) "*The maple tree has a car that is red.*"

P: Oh, ok, um...it could be like, there's a car under the maple tree. When you say like, it's, um...that's the...that's the...when you're trying to like, point out that car so you'll say like it's under the maple tree. The maple tree has that red car.

(6) "*The chair has a box that is cardboard.*"

P: Um, someone left their cardboard box on top of the chair.

(7) "*The honeydew has a bag of gala apples.*"

P: I'm envisioning a giant honeydew melon walking with a bag of gala apples, back from Stop & Shop.

As you can see from these examples, subjects' responses varied in how they related the two entities to each other, and it looks like they did manage to generate locative contexts for the stimuli. The adjectival section, or last part, of the stimulus is simply there to create distance between the second NP and the end of the sentence, which was not necessary in this context elicitation task, but was necessary for paired tasks that had matching stimuli. The task discussed here was performed as part of a multi-experiment project from the Language and Brain Lab at Yale University, connecting the locative and possessive conceptual structures for *have*. The other experiments involved in this project collected various neurolinguistic measures (acceptability ratings, self-paced reading, electroencephalography, and functional magnetic resonance imaging) of how *have* works in comprehension (Zhang 2018). In the tasks associated with these experiments, subjects were given different kinds of context—or sometimes no context—and behavioral data measures were collected in order to look at the effect of context on comprehension of have.

The context elicitation task was included in order to answer questions about *have* in production: Given no information but two entities and the lexical item *have*, could subjects create a context that could relate these entities? What kinds of requirements would arise in these contexts? In this way, we could use subjects' responses to determine what is important or necessary to include in a context that could license a non-ownership, non-containment *have*.

The subjects' responses were coded by type ("locative," "possessive," or neither) and counted. The labels for each category are not to be taken too definitively, and should be interpreted as follows: "Locative" responses placed the entities in a context where they were

related in some spatial way (next to each other, on top of each others, etc.), while "possessive" responses bypassed the non-containment, non-ownership task constraints by changing the semantic nature of the entities. The total number of responses of each type, as well as the number of subjects who provided at least one response of that type, are as follows:

	Number of responses	Number of subjects
Locative	424	55
Possessive	324	53
Neither	52	33
Total	800	55

What happened in these "possessive" responses was that subjects would place the entities in a cartoon world or would construct their context such that the first noun phrase would not correspond to its canonical meaning but instead served as a placeholder for a metonymic or metaphorical meaning. Examples of these are shown below.

(8) *"The maple tree has a car that is red."*

P: Maybe it's, like, a super tall dude, and his nickname's the Maple Tree. I know I already used that answer once. That'll be the last time I use it.

E: It's alright. You can use any answer you—

P: But yup, he's a big basketball player, and his nickname is the Maple Tree. And uh, he's got a car. They're talking about his car.

(9) *"The signpost has a scooter that is pink."*

P: The signpost…oh. Okay, so…there's a school where all the signs have to go to learn how to be good signs, and at the end of school, like, they get to pick which sign to become, but this sign doesn't, like…he's not…cause like there are different signs.

There's like the yield sign and the stop sign. But like, when they go to school they don't have like, a name on them yet, so they're all blank, and so this sign rides his scooter, pink scooter to school, but we don't know what kind of sign he is yet.

(10) *"The saucepan has a cookie sheet that is rectangular."*

P: I guess The Saucepan is, like, a store that sells cookware, and it would have a cookie sheet in it.

(11) *"The Picasso has a painting by Vincent van Gogh."*

P: I guess it's like an art museum that's dedicated to Picasso, or like funded by Picasso, or just, like, really appreciates Picasso, so named it the Picasso, and they bought a painting by that other person.

(12) *"The muffin has a bagel with sesame seeds."*

P: Um. Maybe when two children are making make believe. Because the sentence doesn't really make sense otherwise? [laughs]

E: And could you tell me uh, what they're doing when they're playing make-believe?

P: Uh. Playing at having a bakery.

E: And what do you think someone would mean when they said this sentence?

P: [to self] muffin had a bagel with sesame seeds...um...I don't really know. [laughs] I guess it would be...it would be sort of attributing the muffin [laughs] humanoid characteristics. so talking about the muffin as a character, I guess.

E: Mhm. And could you tell me about the bagel, then?

P: Um. I think the bagel would just be a bagel. [laughs]

(13) *"The honeydew has a bag with gala apples."*

P: The honeydew has a bag with gala apples.

E: Mhm.

P: I guess honeydew's referring to something which is either like, adjacent to a bag with apples, or contains the bag? So maybe the honeydew could be like a...no, I—I wouldn't call it, like, an establishment, though. Honeydew has a bag...? Honeydew has a bag...yeah this one's hard [laughs].

E: [laughs]

P: Maybe just something that happens to contain the bag? Maybe someone named their backpack Honeydew or something. Agh.

(14) *"The notebook has a cup that is white."*

P: A cup?

E: Yeah.

P: *The Notebook*, the movie, has a cup in it that is white.

We can see from these examples that subjects employ a wide variety of strategies in order to change something about the semantic information originally presented to them. They might use nicknames (8), a cartoon world (9, 12), or metonymy in the form of a proper name (10, 11, 13, 14). We also see in all of these examples, especially in (12), that when the subject determines that circumstances of the situation must be changed in order to license *have*, the second entity can stay as it was; its semantic information does not have to be changed in some imaginative way. As the subject says when asked about the second entity after creating an elaborate circumstance to justify a make-believe scenario, "the bagel would just be a bagel."

Why would there be pressure to change something about the first entity only? The answer may lie in subjects' perception of control. All entities chosen for this context elicitation task were objects with little expected control. This was a side effect of the constraint of selecting entity pairs where a containment or ownership relation would be

dispreferred or unexpected. If subjects needed to imbue one of the entities with a level of control, they would need to change the semantic information associated with that entity and possibly provide a justification for doing so ("we are in a cartoon world," "this is a nickname," etc.). There may be reason to believe, as broken down throughout the rest of this essay, that subjects felt a pressure to imbue the first entity, but not the second, with a level of control. In doing so, they would create a power differential. The importance of this power differential will become clear in section 2.2, when the concept of *control asymmetry* in the multidimensional space is introduced.

As shown in Figure (1), the absence of an expected containment or ownership relation did not always lead to a locative relation. In fact, the number of possible locative responses given ranged from one to fourteen, and we saw this full spectrum of number of locative responses in the data. the average number of locative responses per subject was 7.4. This tells us two things: first, that every subject tested had the ability to interpret *have* as a locative relation between two entities—everyone did this at least once—and second, that we cannot simply say that the blocking of a containment or ownership relation based on shape, size, and function of the objects will result in a noncanonical analysis of the *have* relation. There must be other factors contributing to the even spread in number of locative responses produced by each subject. What are these factors, and how can they modulate the kinds of contexts subjects are able to construe for a *have* relation? This question is what motivates the rest of this project.

Inter-subject variability in number of locative responses out of 14, sorted from low to high:

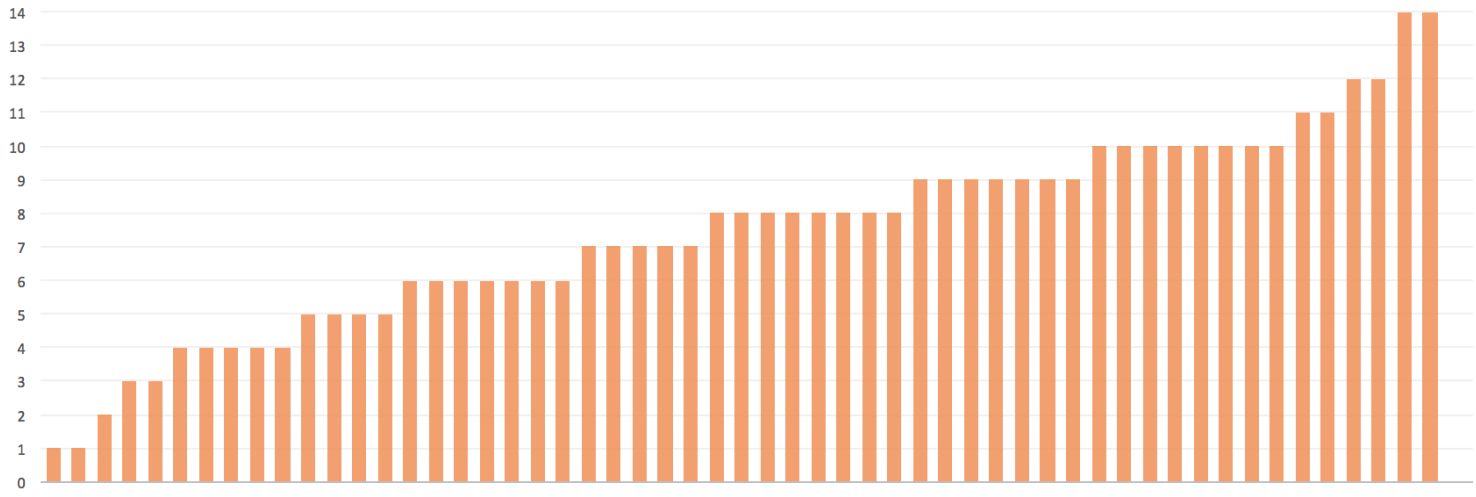

Figure (1): Inter-subject variability in number of locative responses. Each vertical bar represents one subject (out of 55 total).

Chapter 2: Structure

2.1 Semantic and episodic memory

In order to describe the kind of context necessary to license the conceptual structure of the lexical item *have*, we can draw upon two types of memory: semantic and episodic. While semantic memory describes our storage of generalized knowledge about the world, episodic memory refers to memories of specific events. For example, knowledge that dogs bark would be stored in semantic memory, while information about a specific instance of your neighbor's dog barking one night would be stored in episodic memory.

Piñango (2019) writes that semantic and episodic memory are not distinct from each other, but exist at different areas along a mental pathway. We can see from our example of dogs barking that episodic memory is more dependent on context than semantic memory. Different elements of the context (a different neighbor, different night, etc.) distinguish one episodic memory from another. Regardless of context, however, the knowledge that dogs bark remains the same. In fact, we can understand this knowledge as the result of a generalization across many episodic memories (Piñango 2019). In the case of dogs, this would mean that we know dogs bark because we have experienced dogs barking in enough different contexts that we have generalized this information as a fact about the world. Piñango (2019) writes that "situations enter semantic memory through episodic memory as *contextualized episodes*, which, over time, become increasingly underspecified and therefore generalizable and more accessible to language (i.e. more "semantic") (p. 627).

And how can we represent the storage of situation-episodes in memory? Piñango (2019) does this using a model called the multidimensional space, or MdS.

2.2 Multidimensional space (MdS)

So far, the working model of the MdS has two dimensions: control asymmetry and connectedness (Piñango 2019). Control asymmetry has to do with how much control each entity (in our cases, each entity within a single two-entity construction) has over each other. What is relevant here is not the magnitude of control involved, but the *asymmetry* of that magnitude. That is, a construction involving two entities with equal, large amounts of power over each other (like the reader and her heart) has the same control asymmetry as a construction involving two entities with equal, small amounts of power over each other (like a rock and a leaf on the ground).

The other dimension of the working MdS model is connectedness. This can be thought of as an inverse of inextricability, or hard it is to separate the two entities. We can see how the two dimensions interact with each other in Figure (2) below:

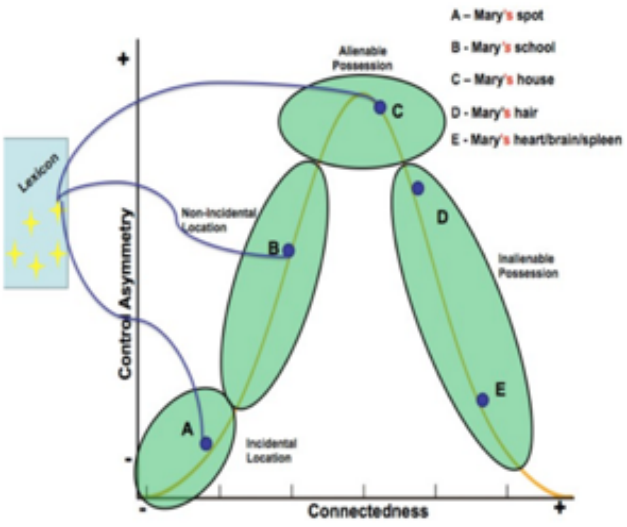

Figure (2). The two dimensions of the MdS (Piñango 2017 (1)).

To understand how *have* would be distributed in the above figure, we can replace the A, B, C, D, and E nodes with:

A — Mary has a spot (incidental location)

B — Mary has a school (non-incidental location)

C — Mary has a house (alienable possession)

D — Mary has hair (inalienable possession)

E — Mary has a heart/brain/spleen (inalienable possession)

In this way, we can see that the *have* construction can be used in the locative as well as the possessive domains, and that the two do not even be separate domains at all, but appear to exist on a continuous gradient.

2.3 Force in conceptual structure

One of the main utilities of the MdS is that it links domains (locative and possessive, at least) that might otherwise be thought of as separate. However, the model as described above may not enough to explain certain prepositions like *on*, which have meanings that arise from the force domain. (Throughout this essay, we will use the term *force* to mean perception of physical force, rather than any kind of linguistic force, like illocutionary force.)

Gärdenfors (2015) shows the force components present in prepositions such as the English *on*. Specifically, he claims that *on* has a force requirement rather than a visuospatial one. A visuospatial explanation of *on* would be that the entities in the relation "ENTITY1 is on ENTITY2" are vertically aligned such that ENTITY2 is resting on the upper surface of ENTITY1. In fact, there is space for this interpretation in Jackendoff's (1983) account of *on*:

> The most salient place-function expressed by "on" requires its reference object to have an upper surface. Another sense of "on" occurs in "the fly on the ceiling," in which the place-function involves the *outer* (i.e., visible) surface of the reference object. These two senses seem to be typicality conditions in a preference rule system in the lexical entry for "on." (p. 162)

Jackendoff (1983) does bring up situations where the upper-surface requirement does not check out, and proposes an outer-surface requirement in those cases. But this explanation does not make a watertight case for the necessity of the upper surface in the "most salient place-function expressed by "on."" When a fly is on the ceiling, it is true that the fly is on the "*outer* (i.e., visible) surface" of that section of drywall or plaster. However, my intuition is that you could also describe a fly as "on the end of a banana" or another object where there is no one outer/visible surface (a banana has many sides), yet the surface the fly is in contact with does not have to be the upper one. It can indeed be on the upper surface, but it could also be perched sideways or upside-down. Gärdenfors suggests an explanation that might be more precise: that the requirement for *on* is actually from the force domain. It just happens to be that most things exert a force downward due to the effect of gravity, so that the entity exerting the force on the other is usually the entity located above the other in a visuospatial domain (meaning that, in most cases, Jackendoff (1983)'s upper-surface requirement would hold). In the case of an entity that exerts a force upward, however, we would expect the visuospatial and force domain accounts to predict different things. This is indeed what we see, illustrated in the example of a balloon exerting a force upward on a lamp:

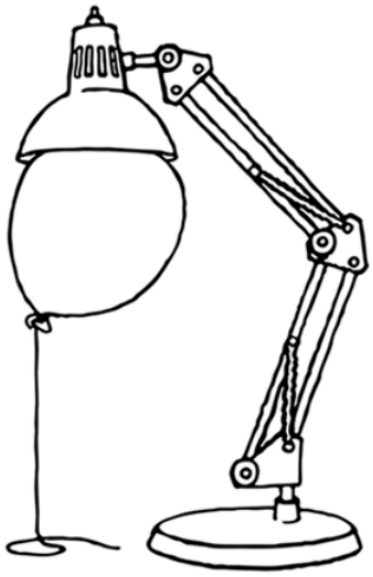

Figure 11. Is the lamp *on* the balloon?

Fig. 3. Reproduced from Gärdenfors (2015).

According to Gärdenfors (2015), it would feel odd to say that the lamp is *on* the balloon. Further testing would need to be done in order to confirm that this intuition is true for most native English speakers, but for now let us explore what this intuition would mean for a visuospatial account versus a force account. Though the lamphead is located spatially above the balloon in this image, the lamphead is the entity that experiences the force from the balloon, and this force is what brings the two entities in contact with each other. Therefore, a visuospatial account would predict the lamp to fall into the ENTITY1 spot, whereas a force account would predict the lamp to be ENTITY2. Even the observation of a misfit between the lamp and ENTITY1 is enough to say that the visuospatial account is not sufficient to explain the requirements of *on*.

What I expect is that *have* has requirements that are similar to *on*; ENTITY2 in the construction "ENTITY1 has ENTITY2" must exert some force on ENTITY1, and that force must be the *reason* why the two are in contact, if they are. To quickly test our own intuitions on this, we can look at Gärdenfors's lamp example and replace the caption with

the phrase "The balloon has the lamp." Does this seem alright to say? Is "The lamp has the balloon" better? My expectation is that "The lamp has the balloon" is indeed better, and that the observation would be switched if the balloon were instead something exerting a force downward.

This expectation is motivated by an observation in the preliminary collection of contexts for the bare "have" sentences. In some cases, when allowed by the shape of the entities involved, subjects reported that one entity hung from the other. In this case, the vertical orientation of entities was switched from the other elicitations. The following are three examples of such cases, along with the prompt associated with each elicitation:

(15) *"The mirror has a cactus that is small."*

P: Person saying this is probably telling another individual, like, um a person saying this would probably like, like different types of plants, and so he would collect them, and so like, he would tell another person that…that…about like, about where they collected like all these different type of plants. And he's telling them like where he stored his partic—this partic—like where he stored like this particular plant, this being the **cactus on his mirror. Hanging from the mirror**.

(16) *"The honeydew has a bag of Gala apples."*

P: Um, uh…I'm picturing a **honeydew plant/tree that has a bag of Gala apples hanging from it.**

(17) *"The honeydew has a bag of Gala apples."*

P: Um, there's like a plastic grocery bag filled with gala apples that's like **hanging off of a honeydew**, somehow.

Another motivation for the consideration of the force domain as the relevant domain for a *have*-relation requirement is that it provides a possible reason for a phenomenon observed in the qualitative data analysis: subjects would often spontaneously include some way that the two entities given are attached to each other. They would introduce some medium (e.g. glue) or instrument (e.g. a chain) into the context that would perform this role.

Force explains attachment in a way that other domains (visuospatial, temporal, etc.) cannot. If the exertion of a force by one entity onto another facilitates or licenses the *have* relation, then we would expect any kind of adhesion force to also perform this function. This would explain the observed phenomenon in the preliminary context elicitation of subjects' spontaneously including tape, glue, metal locks, and other types of binders. When they include these binders between the two entities given, they are introducing an attachment force that might not be there otherwise. Subjects also used cultural, historical, or emotional "binders" to introduce an attachment force as well. (One entity might always be seen along with the other, or they may perform a special function when used together, or the two may come sold in a set).

This is not to say that physical domains are necessarily a primitive upon which other conceptual domains are based. That is, although I consider emotional or cultural force to fill the same role in facilitating—or even licensing—the *have* relation as physical forces like gravity, I make no claim as to whether more abstract types of forces are metaphors built from physical force. In order to make this clear in analysis of the elicited contexts, I have labeled the directions of these types of forces simply as "unspecified".

2.4 But which force?

According to Gärdenfors (2015), a declarative sentence is a construal of a situation. We can express the same situation in many, many different ways, which would all focus on different things. For example, when I hit the table, I could say it just like that, "I hit the

table." I could also say, "The table was hit by me." Both construals would express the same event in different ways, focusing on different details of the scenario.

We know from Newton's third law that a unidirectional account of force would be an incomplete one. When I step on the ground, I could say that my foot is exerting a force on the ground. I could also say that the ground is exerting an equal and opposite force on my foot. Both are true. In (18), when I say one entity (let's say, a chair) is experiencing a force from another (cardboard box), I could also say that the second entity is experiencing a force from the first. That is, the chair is exerting what is called a "normal force" on the cardboard box keeping it from falling through the chair to the ground below.

(18) The chair has a box that is cardboard.

So how do we decide the way we characterize the direction of the force? Is the box exerting a force on the chair? Or is the chair exerting a force on the box?

One answer is that it does not matter. As we've said, every sentence is a construal of a situation. The way we construe a situation—that is, the things we make prominent—will change the acceptability of what kinds of sentences are said. Let us imagine we are deep into a conversation about normal forces. Think about all the things that are working against gravity to hold objects up, the ground is holding you, keeping you from falling through to the earth's core, the chair you're sitting on is holding you from falling to the ground. Or perhaps we are especially interested in the structural integrity of different chairs. Particularly flimsy ones might not be able to hold the same weights as others, and it may be reassuring to know that a certain chair can hold a certain box. If we have established that we are focusing on normal forces, then I suspect that it may become somewhat more acceptable to say that the cardboard box has the chair.

(19) Don't worry, the cardboard box won't fall to the ground. It has a sturdy chair.

We must also remember that, in any given situation, many forces are at play. When we consider the magnetic force between a refrigerator magnet and the refrigerator, we are simply pointing out one of many forces in the situation. Besides the magnetic pull of the magnet to the refrigerator (or the refrigerator to the magnet), there is also the weight of the magnet, the weight of the refrigerator, and, if the magnet is starting to slip, some friction between its surface and the refrigerator door. There may be varying degrees to which the perceptions of these different forces are accessible to a communicator in the environment.

So another answer is that, though there is no single correct force direction in an actual situation, we can look to the construal of the situation to understand which force direction is *relevant* for the relation between the entities involved. For Gärdenfors's proposal of force requirements in the meaning of *on*, as discussed in section (2.3), the relevant force would be the one that brings the two entities in contact with each other. And whether it is because humans have evolved to communicate with each other, or because of a biological basis, like the integration of gravity perception in the visual system (Maffei et al. 2015) or the link between non-visual head and body orientation cues and gravity (La Scaleia et al. 2019), we expect most people's construals will align.

2.5 Early childhood

The importance of force perception is evident even in our first conceptions of what possession entails. Infants' first introduction to the concept of ownership (and, arguably, their first experience of containment) plays out in the process of learning how to grasp and ingest. Here, we will use Rochat (2014)'s account of the origins of possession. High control asymmetry is necessary for success; the child can only grasp once they gain the motor skills necessary to exert control over an object. There is also some necessary spatial proximity between the infant and the object in order for the infant to grasp it, and that spatial relation exists because of a force. Once the infant successfully grasps the object, they can surround

it—either in their fist or their mouth—and, eventually, attempt to synthesize it with their own body through ingestion.

We can see that this milestone in the infant's development has elements of what might be called location or possession, but the thing that causes the grasping situation to take place is the force experienced by the infant. Until the infant gains the motor skills necessary to hold onto an object, to experience this force, the grasping action is incomplete.

2.6 Lexico-conceptual structure

Jackendoff (1983) treats verbs of possession as belonging to a family of nonspatial semantic fields, where possession "plays the role that location does in the spatial field, as the central element of a group of [STATE] and [EVENT] concepts." He represents these verbs, along with other events and states, using something called a lexico-conceptual structure, or LCS, which represents meaning structures. According to Jackendoff (1983)'s analysis of "have," the sentence "Beth has the doll" would reflect the state of the doll being at a place (Beth). Drawn out in a tree form, this would look like:

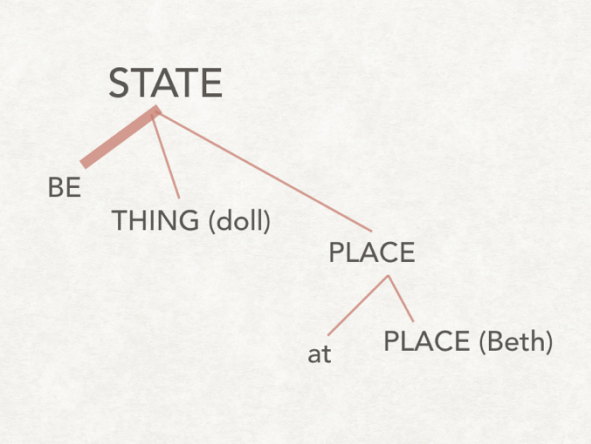

Figure (4): My depiction of the LCS of *have*, as proposed by Jackendoff (1983)

Zhang (2018) builds on this, treating possession as a bi-eventive causal structure, as follows:

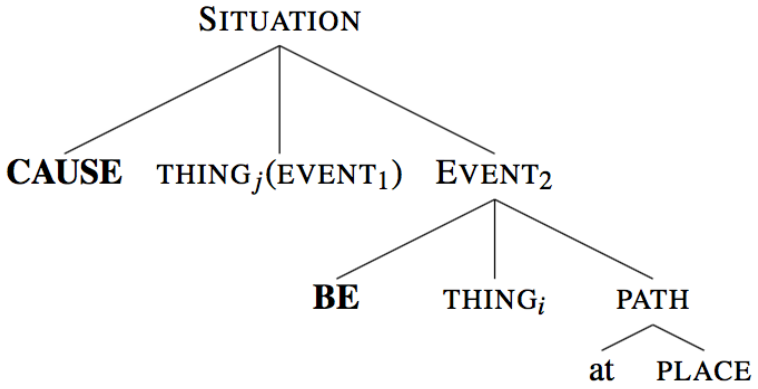

Figure 4: LCS of *have*, as proposed by Zhang 2018

Chapter 3: A return to the data

3.1 The current analysis

My analysis builds upon Zhang's structure, placing a specific restriction on the first argument of the causative. I propose that the first argument—the [THING$_j$/EVENT] that causes the second argument—must be a force pointing from [THING$_i$] to [PLACE]. This makes my analysis of the *have* relation similar to Gärdenfors (2015)'s force-dynamic analysis of the preposition *on*. Gärdenfors (2015) proposes that "...the meaning of "*x* is on *y*" is that *the force vector from x makes x come in contact with y, and a counterforce from y balances the force vector of x*." Likewise, I propose that meaning of "*y* has *x*" is that some force vector from *x* causes *x* to come in contact with *y*. In introducing this constraint in the lexico-conceptual structure for have, I agree with Jackendoff (1983) that humans use existing machinery to understand abstract concepts, and ask the question: in the case of the conceptual structure of *have*, what if this underlying machinery has roots in the force domain?

In order to examine this new lexico-conceptual structure, we can break it down into two parts and a rule connecting them:

(20) *y* has *x*.

(21) There is a (perceived) force pointing from *y* to *x*.

(22) *y* is in contact with *x*.

(23) The meaning of (20) is that (21) and (22) are true, and (21) is what *causes* (22).

Where Zhang (2018) placed the first argument of the causative—the [THING$_i$]—I propose to place (21), as follows:

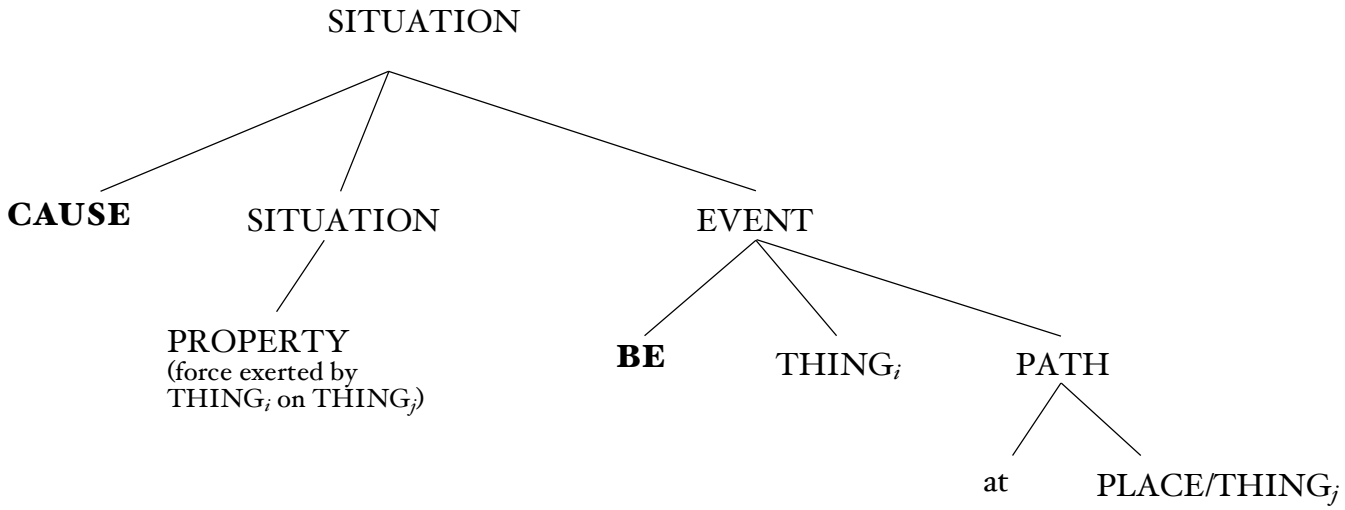

Figure (6): My updated LCS for *have*

What the addition of this force-dynamic requirement means is that, just as Gärdenfors (2015)'s analysis of *on* makes the preposition "invariant under spatial *rotation* as long as the

force relations stay the same," so too does my analysis of *have* allow the *have* relation to hold even when the spatial orientation of the entities changes. I have also slightly modified Zhang (2018)'s representation of PLACE to emphasize that this sub-branch of the LCS is doing the work of (22), showing that THING_i and THING_j are in contact.

Going back to the example sentence "Beth has the doll," let's first consider the alienable possession case. (I will later argue that the inalienable possession case is an abstraction of this structure). What I am proposing is that the thing that causes the doll to be in contact with Beth must be some force pointing from the doll to Beth.

Perhaps the most expected force that could relate these two entities is gravitational force; we can imagine that the doll is in Beth's lap, or in her arms, such that she is bearing the weight of the doll. Or maybe the doll is dangling from her hand—here, the spatial orientation is changed (Beth is no longer supporting from below, but above), but the weight of the doll is still something exerted by the doll and experienced by Beth.

We can imagine Beth failing to "have" the doll. The doll could fall from her fingertips or be knocked out of her arms, and then we could say:

(24) Beth had the doll, but she doesn't anymore.
(25) Beth had the doll, but she dropped it.
(26) Beth had the doll, but now it's on the floor.

Or even, if we wanted to have a little fun:

(27) Beth had the doll, but now the floor has it.

Gravity is a relevant force in these examples, and it is just one of the many possible forces that could be provided by the context. Given the right context, we can also run through the same exercise with some kind of attachment force at play. Maybe Beth's hands

are sticky with honey, and she has touched the face of a doll. When she moves her hand to the left, the stickiness of the honey makes the doll move leftward as well. When Beth moves her hand to the right, so does the doll. We could say, in this case, that Beth has the doll. Let's say that, after a bit of dragging across the carpet, the adhesion loses its grip, and the doll breaks free of Beth's sticky hand. Now, we can say the following:

(28) Beth had the doll, but she doesn't anymore.

Instead of honey, we can attach the doll to Beth's hand with something else, like tape or cling wrap. These last two examples of attachment (tape, cling wrap) were not invented by me. They arose spontaneously from the dataset from section (1.2). Now, we find ourselves with a new thing to look for in the data: direction of force.

In order to assess the forces at play in each elicited context in the dataset, responses were coded for force direction and vertical orientation of the entities involved. By vertical orientation, I just mean: Was one entity above the other? If so, which? The choice of vertical orientation rather than horizontal orientation (or any other kind of orientation) was simply to capture the largest variability in the dataset. I could imagine a future analysis looking at different types of orientation as well. Gravitational force was coded as a downward force, adhesion and other forms of attachment, as well as friction, were often coded as sideways forces, and tension was coded as an upward force. This may at first seem a bit too technical and unintuitive to be cognitively real, so the next section (3.2) will alleviate some of those worries.

3.2 Diagnostics

Before we get into the force analysis of this dataset, we should lay down some ground rules. In discussing forces, we must keep in mind that what we are interested in is not force itself, but the perception of force. We can only discuss force in the way it exists in the

mind, and one way to get at how it exists in the mind is to look at how we talk about it. When people talk about ourselves or some object experiencing a force, we talk about burdens, tensions, and exhaustion. Just like people can groan under the weight of expectations, a table can groan under the weight of a heavy microwave.

To test whether we perceive some entity as experiencing a force, we can use an *endurance failure* test, or a *crack/deformity* test. The test is as follows:

(29) If a scale modification in *y* might cause *x* to "crack" or otherwise deform, *y* is most likely exerting a force on *x*.

(30) A candle is on the table. If the candle were enormous, it might cause the table to crack.
(31) A magnet is on the fridge. If it were a superpowered magnet (like the magnets used to lift junk cars), it might cause the fridge to crack or otherwise deform.
(32) There is something worrying on my mind. If it were an extremely obsessive worry, it might cause damage to my mind.

(33) A candle is next to a table, and both are sitting on the floor. If the candle were enormous, it might cause the floor to crack, but not the table.

We conclude from (33) that the candle is exerting a force on the floor, but not the table. We also see in (32) that the concept of "force" can exist beyond just physical laws.

Another way to get to an intuitive understanding of force perception would be to strip the semantic information from the entity that is potentially exerting a force, and ask whether we would be careful with what we would replace it with. For example,

(34) "What kinds of things would you place on your computer? What kinds of things

would you not?"

For anyone who cares about their computer, there is an important distinction between these two lists. In contrast, when we ask:

(35) "What kinds of things would you place next to your computer? What kinds of things would you not?"

that distinction mostly disappears. The gravitational force that brought entities together in the "on" relation is absent in the "next to" relation. And the situations where such a distinction does exist are not random; they are invoking some force, just not a gravitational one. For example, we might not want to put a strong magnet by the computer for fear of wiping the memory, and it's possible we would not want to put an agitated dog there, in case it decides to take a bite. The relevant forces here are magnetic force and the force exerted by the dog's jaw.

This test also works for more abstract perceptions of force:

(36) "What kinds of things would you place on your mind? What kinds of things would you not?"
(37) "What kinds of people would you place in your friend group? What kinds of people would you not?"
(38) "What kinds of things would you place in your schedule? What kinds of things would you not?"

If we agree that (36-38) characterize important distinctions between the set of things you would place on your mind/in your friend group/in your schedule and the set of things you

would not, we can determine from those intuitions that there is some perception of force at play here. In this way, thoughts can exert a perceived force on the mind, friends' presence can exert a perceived force on a social circle, and appointments can exert a perceived force on a schedule.

3.3 The interpretations

Now, let us walk through the interpretation associated with each possible result:

If we found no effect of vertical orientation or contact on a downward force direction, then we would have little reason to believe that the conceptual structure of *have* includes a force component or requirement. In this case, what we understood about have's meaning might be fully explained using the visuospatial domain only. The semantic information of have would determine the vertical orientation of the entities because its conceptual structure would have requirements in the visuospatial domain, but this would have nothing to do with the perception of force.

If we found an effect of vertical orientation on downward force direction but no effect of contact, we would have to reassess our understanding of the LCS of *have*. This result would suggest that the second argument of the causative in Figure (6) is less necessary than previously thought.

If we found an interaction between vertical orientation and contact on the presence of a downward force in the context, this would mean our result supports a force requirement in the LCS of *have* that acts as the first argument of a causative (where the second argument represents contact between the two entities), as described in section 3.1. The LCS I have laid out here predicts such an interaction.

3.4 Results and discussion

In order to perform statistical analysis (I used a generalized linear mixed-effects model) on this dataset and generate plots, I used the the lme4 and ggplot2 packages in R. I

found a significant interaction between vertical orientation and contact on downward force direction in the elicited contexts ($p < 0.05$).

Contact and no-contact:

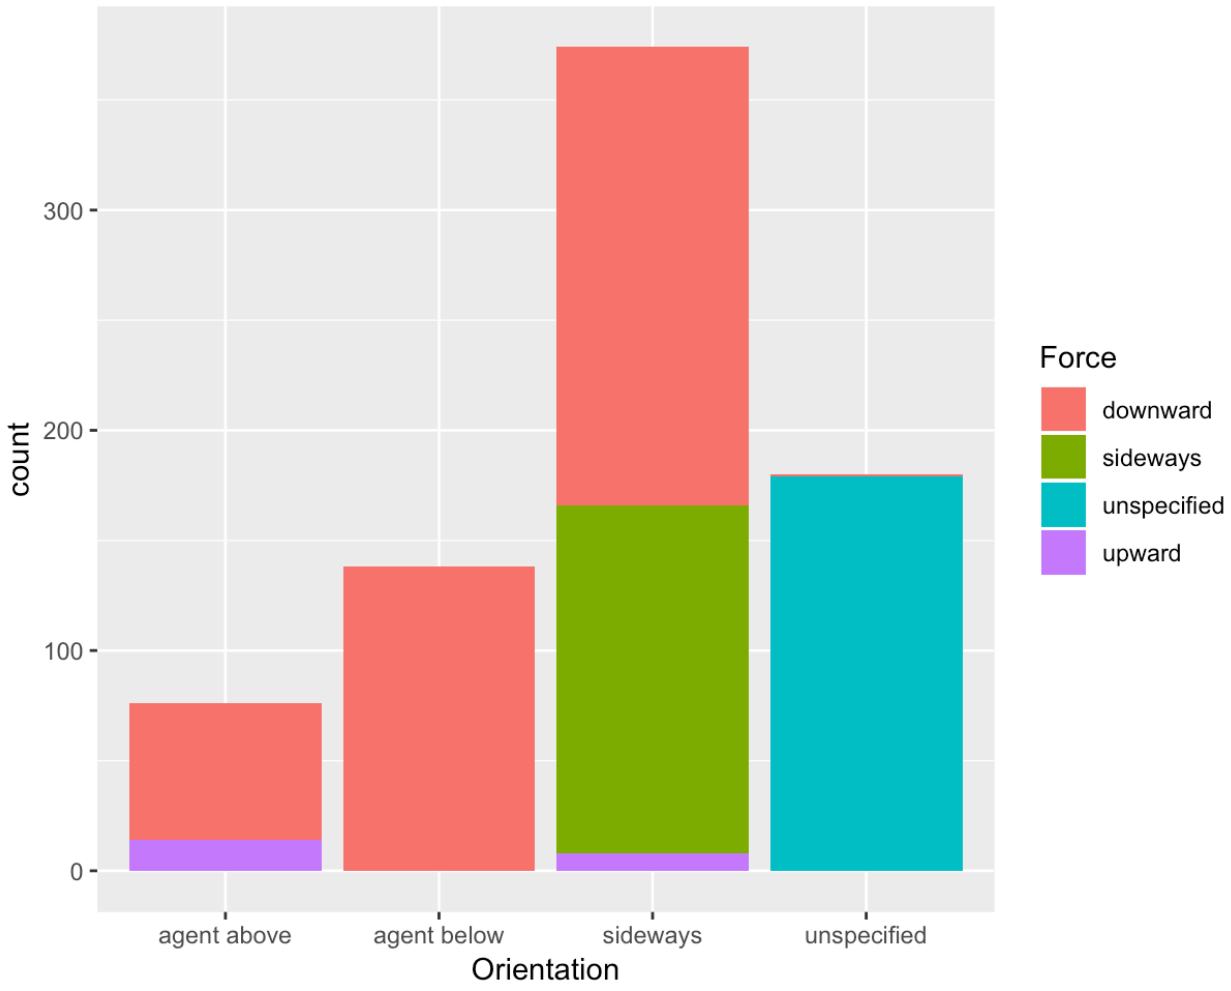

Figure (7): No clear alignment of force direction and orientation in the full dataset

We can see in Figure (7) that, aside from the case of unspecified orientation, downward force direction dominates across the board. In agent-above, agent-below, and side by side orientations, the force direction was majority downward. We can see the presence of sideways-force direction in the sideways orientation, but we cannot say that the orientations are clearly aligned with the force directions.

So what is happening when there is no contact between the two entities? Some of the items do not lend themselves well to contact, like the maple tree and the red car. In order to make contact between these entities, subjects would have to describe an extremely low-hanging tree branch or a crash scenario, which several subjects did indeed do. What we also see, however, is that subjects often placed the entities side by side in an incidental location relation. This might be explained by the structure of the task itself, as explained in section 3.5.

Contact only:

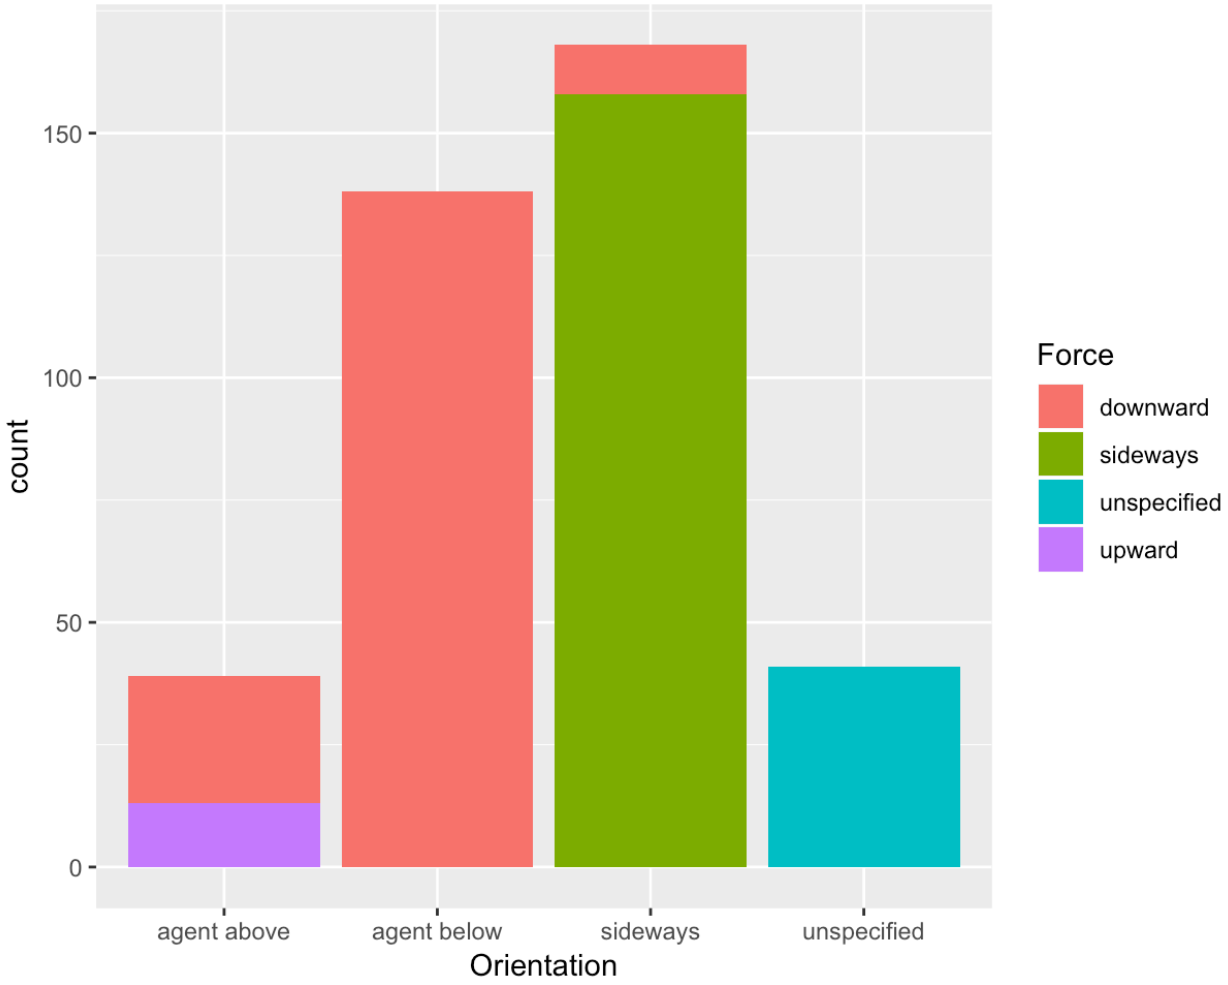

Figure (8): When there is contact, the agent-below and sideways orientations show clear alignment with force direction.

When there is contact, and the spatial orientation is such that the agent is spatially below the theme, the force direction was always downward (138/138, or 100%). When there is contact, and the context places the two entities side by side, there is a higher chance that the force direction will be sideways rather than downward (158/168, or 94%). This is represented by the green section in the plots above. In fact, subjects seem to go to great lengths to make sure the experimenter understands that there exists such a sideways force, and that it is a strong one.

In conclusion, the responses from the dataset indicate a force-based requirement in the conceptual structure of *have*, as described by my updated LCS for *have*. When there is contact between the entities involved, vertical orientation aligns with force direction such that the entity that *has* the other entity is experiencing a force from the other entity.

The following are some examples from the variety of attachments subjects included in the elicited contexts, along with the prompt the subject is responding to for each example, and a force-dynamic analysis of each. Usually, analysis of individual responses, especially of unusual ones, would not be very informative when it comes to describing a dataset. However, I present these here in order to show the patterns they exemplify.

"The maple tree has a car that is red."

1. P: Um, I imagine a maple tree with a car stuck inside it.

E: And how do you mean, stuck inside?

P: Um there is a tree that has grown around the car. Literally, the car is inside. I guess, like, perhaps half of it's sticking out. The front half is lodged inside the wood of the tree.

There is contact between the two entities here, and force direction was coded as horizontal because the force that holds the car in contact with the tree is friction between the sides of

the car and the wood of the tree. It is true that gravitational force also contributes to this orientation, but it is important to note the way the subject chose to construe the situation. They did not state that the car was resting on an opening in the tree, which would place emphasis on the gravitational force the car would be exerting on the tree. Instead, the subject chose to describe the car as "stuck inside" the tree and "lodged inside the wood."

> *"The signpost has a scooter that is pink."*
>
> 2. P: um, uh, [sigh], there is a scooter attached to the signpost. E: and how would it be attached? P: I don't know, like welded on or taped on or something

Here, there is no visuospatial vertical hierarchy between the two entities. The subject chose two unusual ways of attaching the scooter to the signpost (welding and tape). It seems that the particular method of attachment was less important than the requirement that the two entities be attached in *some* way, whatever way that may be.

> *"The signpost has a scooter that is pink."*
>
> 3. P: Imagining uh, a pink scooter that is like, as locked as a scooter can be to a signpost.

The subject made sure to note that the entities were not only in contact, but in as close contact as possible.

> *"The honeydew has a bag of gala apples."*
>
> 4. P: Oh…is the bag of gala apples stapled to the honeydew?

Like in the scooter example, this subject chose an unusual way of attaching the bag of apples and the honeydew, indicating that the attachment itself is more important than the particular method of attachment.

> *"The mirror has a cactus that is small."*
>
> 5. P: Um, maybe, like, I, I imagine like a magnetic cactus that's, like, on the mirror.

The subject used magnetic force here in order to show the connection between the "cactus" and the mirror.

3.6 Task dependency

We also must discuss some level of task dependency in the results we saw from this experiment. There is reason to suspect that when subjects produce contexts placing the two entities next to each other spatially with no contact between them, they are not responding to the lexical item *have* itself. In these cases, subjects may just be employing a task-dependent heuristic.

In a book chapter about the subprime mortgage housing crisis, behavioral economist Dan Ariely (2009) writes:

> "When we can't figure out the right answer to the question facing us, we often figure out the answer to a slightly different question, and apply this answer to the original problem." (p.286)

Ariely uses this characteristic of human behavior to describe why so many home buyers in the U.S. leading up to the year 2008 took out higher loan amounts from the bank than were reasonable. Ariely's take on the situation is that "instead of figuring out the

answer to the correct question (how much should we borrow?), they focus on a different question altogether, one that is not the correct question, but one that they can easily answer: how much can we borrow?" (2009). A difficult task is replaced with a much easier task that will not provide the correct answer, but that feels close enough.

We know that task of context elicitations for the stimuli in the present dataset is difficult. The production of a context for these particular sentences requires some special level of mental work. This is because, as discussed in section (1.2), stimuli were chosen so that the most common use of lexical *have*—as something denoting a containment or ownership relation—was dispreferred. Even though the task was difficult, subjects were expected to successfully construe a context that would relate the given entities with *have* if they were at all capable of doing so. What was not accounted for in the design of the study was that subjects may, without even being aware of this, perform the same decision-making heuristic that led to the subprime mortgage housing crisis of 2008.

Instead of responding to the prompt given (produce a context in which this phrase might be uttered), subjects may be choosing to respond to something *similar*, and much easier: produce a context, any context, in which these two entities might be related. For example, in response to the item "The maple tree has a car that is red," subjects may simply try to come up with any relation that could hold between a maple tree and a red car. The simplest answer? Cars can be parked next to trees. Perhaps this is the most common situation in which we have witnessed a car and a tree in the same scene.

"The maple tree has a car that is red."

(39) Produce a context in which this phrase might be uttered.

(40) Produce a context in which a tree and a car might be related.

This type of response has little to do with *have* and is mostly based on the two entities given and the structure of the task. If *have* does contribute anything in this case, it

is simply the requirement that the two entities be related in some way. The tasks involved in answering the two related prompts (the actual one, and the heuristic) are broken down in the following way: If you are responding to (39), you are working with

1. an entity (a maple tree),
2. the lexical item *have*,
3. another entity (a car), and
4. an adjective to describe the second entity (red).

You are asked to find a way to make sense of all of these elements and put them together in a context. It is not the case that you are given *only* this information; you are given syntactic structure and phonetic information in addition to meanings, but this alone is not enough to license *have*. You will need to introduce some perceived force in the context, a force that causes the two entities to come in contact with each other. This part is difficult. A quick way to dramatically reduce the difficulty of the task, while keeping it as similar as possible to the original, is to just bypass item 2. Then you are left with

1. an entity (a maple tree)
2. another entity (a car), and
3. an adjective to describe the second entity (red).

A simple context that relates the entities is that of a car being parked by the maple tree. This is a much easier task, and subjects may decide to answer this instead when the original task gets too difficult.

Subjects' employment of such a heuristic would likely manifest in elicited contexts where the prominent force is downward, there is no contact between the two entities, and the entities' spatial orientation is side by side. These criteria characterize 198 of the 768

elicited contexts, or 26% of subjects' reponses, distributed across items and subjects. This does not necessarily mean that the full 26% of these responses are due to a heuristic. What I am claiming here is just that the presence of this heuristic may help explain the presence at all of such data.

The explanation I am proposing here is that, in some of the elicitations, the prompt (40) was so much easier to answer than (39) that the subject chose (consciously or less-than-consciously) to respond to (40) instead. This could be due to something like the subject's level of fatigue or a particularly large challenge for this subject in response to this item. As we remember from our initial description of the elicited contexts, this quarter of the data can explain the surprisingly high occurrence of contexts involving incidental location. If such a heuristic was indeed employed, then perhaps these contexts are not indicating anything special about subjects' understanding of *have*. Because our discussion of the results in section (3.4) was focused on contexts in which there was contact between the two entities, the existence of this task-dependent heuristic does not affect the results discussed there.

References

Ariely, D. (2009). Thoughts about the Subprime Mortgage Crisis and its Consequences. *Predictably Irrational, Revised and Expanded Edition: The Hidden Forces That Shape Our Decisions*. HarperCollins.

Caramazza, A., McCloskey, M., Green, B. (1981). Naïve beliefs in "sophisticated" subjects: misconceptions about trajectories of objects. *Cognition*. 9(2), p. 117-123.

Carey, S. (2009). *The Origin of Concepts*. Oxford University Press, Inc.

Deo, A. (2015). The semantic and pragmatic underpinnings of grammaticalization paths: The progressive to imperfect shift. *Semantics & Pragmatics*. 8(14).

Freeze, R. (1992). Existentials and Other Locatives. *Language*. 68(3). p. 553-595.

Gärdenfors, P. (2015). The Geometry of Preposition Meanings. *The Baltic International Yearbook of Cognition, Logic, and Communication. Volume 10: Perspectives on Spatial Communication.*

Gärdenfors, P., Jost, J., and Warglien, M. (2018). From Actions to Effects: Three Constraints on Event Mappings. *Frontiers in Psychology*.

Jackendoff, R. (1983). *Semantics and Cognition.* The MIT Press.

La Scaleia, B., Lacquaniti, F., and Zago, M. (2019). Body orientation contributes to modelling the effects of gravity for target interception in humans. *Journal of Physiology*. 597.7. p. 2021-2043.

Maffei, V., Indovina, I., Macaluso, E., Ivanenko, Y.P., Orban, G.A., Lacquaniti, F. (2015). "Visual Gravity Cues in the Interpretation of Biological Movements: Neural Correlates in Humans." *NeuroImage*, vol. 104. p. 221-230.

Myler, N. (2016). *Building and Interpreting Possession Sentences.* The MIT Press.

Piñango, M.M. (2017). *The conceptual substance of lexical meanings: implications for meaning composition*. New Haven, CT. SYNC at Yale University, Linguistics and Interdepartmental Neuroscience Program.

Piñango, M.M. (2019). Concept composition during language processing: Two case studies and a model. *The Routledge Handbook of Chinese Applied Linguistics.* eds. Chu-Ren Huang, Zhuo Jing-Schmidt, and Barbara Meisterernst. Routledge. p. 624-644.

Ritter, E. and Rosen, S.T. (1997). The function of *have*. *Lingua*. 101(3-4).

Rochat, P. (2014). *Origins of Possession: Owning and sharing in development*. Cambridge University Press.

Zhang, M., Piñango, M.M., Deo, A. (2018). Real-time roots of meaning change: Electrophysiology reveals the contextual-modulation processing basis of synchronic variation in the location-possession domain. *CogSci*. p. 2783-2788.

Zwarts, J. (2005). Prepositional Aspect and the Algebra of Paths. *Linguistics and Philosophy*.

Zwarts, J. and Gärdenfors, P. (2016). Locative and Directional Prepositions in Conceptual Spaces: The Role of Polar Convexity. *Journal of Logic, Language, and Information*. 25(1). p. 109-138