

Default Domain Restriction Possibilities

Katherine Ritchie
University of California, Irvine

Henry Schiller
*The University of Sheffield*¹

Abstract. We start with an observation about implicit quantifier domain restriction: certain implicit restrictions (e.g., restricting objects by location and time) appear to be more natural and widely available than others (e.g., restricting objects by color, aesthetic, or historical properties). Our aim is to explain why this is. That is, we aim to explain why some implicit domain restriction possibilities are available by default. We argue that, regardless of their other explanatory virtues, extant pragmatic and metasemantic frameworks leave this question unanswered. We then motivate a partially nativist account of domain restriction that involves a minimal view of joint planning around broad shared goals about navigating and influencing our environments augmented with cognitive heuristics that facilitate these. Finally, we sketch how the view can be extended to account for the ways *non-default* restriction possibilities become available when conversationalists have shared idiosyncratic goals.

1. Domain Restriction and Book Cases

Definite noun phrases and quantifiers appear to have different domains of evaluation in different contexts. A simple demonstration of this can be given by considering a scenario like the one illustrated in the below image (*Fig. 1*), and whether (1) is true of it:




Figure 1. One book on a table

illustrated in the below image (*Fig. 1*), and whether (1) is true of it:

(1) The book is on the table.

Different judgements about (1) would be warranted if the depicted scenario was of an empty table, or of a table with a book under it, or of a table with one book on it and another book under it. The literature is full of answers to

¹ Both authors contributed equally: names are listed in alphabetical order.

questions about the meaning of sentences like (1) in context (as well as the context-invariant contribution made by sentences like it). For instance, since there are many books in the world and definite singular phrases are thought to require uniqueness,² one could hold that (1) expresses something false (or with no truth value) but via pragmatic reasoning we recover the intended meaning (something like: that the book in the scene depicted is on the table in that picture; Bach 1994, Cappelen & Lepore 2002). Alternatively, one could adopt the more common view, that (1) expresses something true given a semantic restriction, for instance given a presupposed relevant domain of objects (von Stechow 1994) or other additional content that restricts the domain (Sellars 1954; Sainsbury 1979; Neale 1990, 2000; Recanati 1986; Stanley & Szabo 2000). Different accounts cash out the ‘restriction’ in different ways. For instance, in terms of situations (sets of objects and relations; Barwise & Perry 1983, Elbourne 2013, Kratzer 2023), elliptical information (Collins 2018), or information that is presupposed about the speaker’s intentions (Stalnaker 2014). Much has also been written about how utterances of sentences interact with discourse environments to fix the domain of quantification, e.g., via a kind of lexical (von Stechow 1994, Roberts 1996/2012) or pragmatic (Abreu Zavaleta 2021) presupposition, discourse structure (Schwarz 2019), through the recovery of speaker intentions (Buchanan 2010; King 2014, 2022), through features of the conversational situation that are objectively relevant to conversational goals (Gauker 2003), or via prominence rankings determined by mechanisms of discourse coherence (Stojnić 2021).

Our aim here is not to directly engage in any of these debates. Rather, we focus on an underexplored question about what determines the implicit restriction possibilities in context.³ More specifically, we focus on how, by default, there seem to be certain restriction possibilities (e.g., those having to do with location) that are more natural than others (e.g., those having to do with historical properties, aesthetics, material, or color).

To illustrate, return to the example above. Suppose you walk into a room and see what is represented in *Fig. 1* and see no other books or tables anywhere else in the room, then (1) seems to be true and felicitous in virtue of conveying something like (1’).

² Within philosophy of language, the idea that singular definite descriptions require uniqueness can be traced back to Frege (1892) and Russell (1905).

³ There are a number of different ways that propositional contents get implicated in linguistic communication. Our paper is about the determination of the literal semantic meaning of a sentence uttered. In principle, this can be distinguished from both the message (content) the speaker intends to communicate, and the actual uptake of the utterance by the addressee (though often these are intertwined in metasemantic theses – see e.g., Section 2.3).

(1') The book in this room is on the table.

By contrast, consider someone uttering (1) in a situation like that represented in *Fig. 2*. In this situation (1) seems infelicitous and, one might think, for a quite obvious reason. There are *three* books on the table, 'the book' requires uniqueness, so an utterance of it is not felicitous. Yet,




Figure 2. Three books on a table

each book might differ from the others in myriad ways.

They might contain different information and be of different subkinds (e.g., one a cookbook, another a novel, the third a textbook). They might differ in being funny, depressing, or innovative. They might be written by different authors and published by different presses in different years. And, to judge them by their covers, each is a different color. If (1) could be interpreted as implicitly restricting the domain of books by topic, genre, author, publication year, aesthetic preferences, cover color, or

other features as in (2a-c), it would be felicitous and, on some views of domain restriction, true as well when uttered about a scenario like in *Fig. 2*.

- (2)
- a. The funniest book ever written is on the table.
 - b. The book on structuralism I read in graduate school is on the table.
 - c. The blue book in this room is on the table.

There appears, however, to be a stark contrast between uttering (1) and conveying (1') and uttering (1) and conveying any of (2a-c). *Prima facie*, it seems that a speaker cannot convey (2a-c) with an utterance of (1), even in a context in which there is a unique book satisfying the additional descriptive material. Further, even if the speaker has an aim to talk about only the blue things in the room (or books they read in graduate school or...), it still seems that they cannot assert (1) and thereby convey one of (2a-c).

Consider one more case. Suppose a speaker utters (3) in a room with only three books arranged as in *Fig. 2*.

(3) Every book is on the table.

Their utterance is clearly felicitous (and plausibly true) and seems to express something like (3').

(3') Every book in this room is on the table.

Now imagine the speaker is in a room in which there are three books on a table and several stacks of books on the floor. Even if the books on the table are all the hardcover, depressing, ... books in the room it seems a speaker cannot express (4a-c) with their utterance of (3).

- (4)
- a. Every hardcover book in this room is on the table
 - b. Every depressing book in this room is on the table.
 - c. Every coffee-stained book in this room is on the table.

By considering these contrasts one might conclude that an important pattern is revealed about the sorts of implicit domain restrictions that are possible. That is, one might think that domains can be implicitly restricted by *locational features* (e.g., the fridge I see in front of us, the room we are in) and features related to what is *presently* perceivable to the speaker at the time of utterance, but not by other properties, like historical, aesthetic, size, subkind, or color properties. And, in fact, in the domain restriction literature, the examples used almost always involve location features (e.g., many of the central examples discussed in Buchanan 2010, von Stechow 1994, Kratzer 2023).

While this generalization captures the patterns considered above, it is too broad. To see why let's consider another case. Suppose that A and B are working in a department store and have been asked to create a display featuring blue merchandise. Now consider the pattern in the dialogue in (5).

- (5)
- A: Where are the blue things?
B: Well the shirts are near the register, and the book is on the table.

Intuitively, B's utterance can be felicitous and true even in situations with three books, like that in *Fig. 2*. Moreover, the second conjunct in B's utterance seems to convey something like (2c). Similar cases could be worked up for other properties. This shows that implicit domain

restrictions are not limited to locational features or to what is perceptible in the proximal environment. Nevertheless, the contrast we began with is real.

These cases lead us to two questions. First, why is it that certain domain restrictions seem to be defaults? That is, why are implicit domain restrictions to features like location and what is presently visible easily available in a wide range of contexts? We will call the restrictions available in a wide range of contexts *default domain restriction possibilities*.⁴ People without shared conversational history can walk into a room with a single book on a table and one can utter (1) thereby conveying (1'). Our initial contrast showed this was not the case for features like subkind, historical properties, color, or many other properties. Given that in some discourse environments implicit restrictions to other features are possible (e.g., the conversation in (5)) we need to answer a second question: How can we account for the fact that non-default restriction possibilities are available by, e.g., altering the discourse environment through asking a question?

We begin by considering whether our first question is already answered by extant pragmatic or metasemantic frameworks (§2). We argue that it is not, and that each involves a way of reframing the question that is left open. We offer our own account that involves a minimal view of joint planning and forming shared goals augmented with cognitive heuristics that facilitate these (§3). We argue that this view—a partially nativist account of implicit domain restriction—answers why it is that certain domain restriction possibilities are available by default. The default domain restriction possibilities are those facilitated by cognitive heuristics. We then turn (§4) to a sketch of an account of the ways *non-default* restriction possibilities become available, arguing that this involves relevance to joint planning with idiosyncratic aims or goals, constrained by features of human psychology. We then offer suggestions for empirical tests of (elements of) our account (§5). Finally, we conclude and suggest future theoretical directions (§6).

⁴ Our use of the term 'default' is not intended in any strong theoretical sense – what we have in mind is simply a generalization about the interpretation you get without explicit contextual setup (e.g., by way of a shared conversational history). It could be argued that there is no such thing as a 'default' context free of contextual setup; then our inquiry could be recast into one about the nature of the background contextual information that we tend to reflexively impart on a conversation.

1.1 What is a Default Domain?

Before going further, it will be useful to clarify what we mean by default domain restriction possibilities to better draw out what we take the examples above to show. We started by using contrasting pairs of examples to generate intuitions about default conversational domains, making the observation that certain restrictions on quantifier domains are easily available, while others are not. The idea is something like the following: by default, certain sets of objects are available for quantifying over with words like 'every', 'the', and 'some' without any additional conversational setup, and this is in virtue of these objects satisfying certain *locational* descriptions (especially those like 'is in this room' that associate with some pretheoretic conception of where we are), whereas other descriptions are too narrow ('is on the left side of this room', 'is a red thing in this room') or too broad ('is on the planet Earth') or simply miss the mark ('is a red thing').

This might establish the following question as a natural starting point: what makes it the case that certain *locational descriptions* are privileged by default over others in serving to restrict quantifier domains? How is it that all and only *the objects in X location* rather than all and only *the objects in Y location* (where Y may even be a proper part of X) serves as an appropriate default restriction on quantifier domains? (Or to infuse this with even more presuppositions: why are we biased towards talking about what we take to be 'our' location?)

But this question already makes substantive commitments about a variety of other ways we might restrict domains. For instance, even though we discount restrictions based on non-locational predicates of objects (e.g., 'is red') we make ontological presumptions about what objects there are. And the examples we've given have all presumed that we're speaking only of some present temporal location. So another question arises now: whether we should restrict our investigation to only the question of how certain locational descriptions are privileged by default, or whether we should expand our investigation to other default-determining dimensions / sources of domain restriction.

As already noted, an utterance of (3), repeated here,

(3) Every book is on the table.

gives us a choice of interpretations between different locational predicates. But it also gives us a choice between temporal restrictions:

(3'') Every book is on the table right now.

(3''') Every book is on the table tomorrow.⁵

Just as there is a clear default bias towards certain spatial locations, there is a clear default bias towards certain temporal locations (e.g., the present, or 'now').

There is something appealing about limiting our investigation to only spatio-temporal dimensions of default restrictions. On an intuitive picture of folk metaphysics (one related to the theoretical picture we inherit from Kant) spatial and temporal features of reality provide us with something like the presupposed background against which we conceptualize objects. It would be natural, then, to think of these as the basic ingredients determining default quantifier domains, and this will be our presumption throughout. Our starting question, then, is *Why does there appear to be a default restriction to (what we conceive of as) the present location and time?*⁶

2. Explanations for Default Domain Restriction Possibilities?

We now have two questions to answer: Why does there appear to be a default restriction to (what we conceive of as) the present location and time? And, how can we account for the fact that non-default restriction possibilities are available given altered discourse environments? We start with the first question. In this section, we consider three possible accounts and argue that each fails to provide an explanation for why certain domain restriction possibilities are defaults.

⁵ It could be objected that (3'') is lexically mandated in a way that (3''') is not, by the use of the present-tense. But this is not so: if you ask someone to describe what the room will look like tomorrow, for instance, they can say things like (3) fairly unproblematically in response. It is also worth noting that in languages like Chinese, where there is no tense marking at all, such lexical mandates will not in general be possible. It might also be objected that while there is no lexical mandate here, there is a simple pragmatic inference from an unqualified use of 'is' to a present-tense context. But this is exactly our point: from where does this present-tense bias originate?

⁶ Could there be more dimensions of bias besides space and time that are important to determining domain restriction defaults? Perhaps. For example, we take it that there might be an ontological default about what counts as an object when considering constructions like "Everything is dirty". Is the left side of a chair or one of its legs a thing that is quantified over? It seems unlikely to us they are and we suggest that a cognitive story, on which this relates to cognitive heuristics that drive how we conceptualize or think of kinds of things, will be relevant.

We want to be clear that this does not imply that any of these accounts fails to explain other phenomena. In fact, we take it that none of them are aimed at answering our first question. And, as it turns out, each leaves it open.

2.1 A Rational Pragmatic Explanation?

As a first pass, one might think that the asymmetry between implicit restrictions based on location or the time of utterance versus restrictions based on history, color, size, or other properties can be explained by general pragmatic rules. For instance, one might suggest that the difference could be the result of adherence to the Maxim of Quantity, Relevance, or the Cooperative Principle (Grice 1975). Or that it could be explained using the tools of Rational Speech Act theory which posits that speakers and hearers engage in recursive reasoning relying on prior probabilities and reasoning about the likelihood of successful communication given considerations of informativity and cost (Frank & Goodman 2012; Degen 2023). We are by no means presenting an alternative to Gricean theories of communication in general. What we say is compatible with the idea that the Cooperative Principle explains much about pragmatic reasoning. Nevertheless, we argue that these accounts do not answer the question at hand, and therefore require supplementation.⁷

To see why, consider again the intuition that when (3) is uttered by a speaker in a room with an addressee, ‘every book’ is taken by default to quantify over all the books in the room, rather than, say, all the green books or all the funny books in the room. One might naturally try to explain this in terms of (a mutual presupposition of) the fact that the speaker is being cooperative and adhering to something like the maxim of Quantity: the principle that says that one should make one’s contribution to a conversation as informative as needed for the purposes of conversation—neither offering too much nor too little. Quantity explains why,

(6) I lost \$100 last night.

⁷ Others have attempted to explain pragmatic inference in terms of psychological, rather than rational, mechanisms of relevance (Sperber & Wilson 1986, Lepore & Stone 2014). Though our own account is broadly sympathetic to this approach, as it stands these views require supplementation just as the ‘rationalist’ accounts we discuss here do.

typically conveys that the speaker lost exactly \$100, even though it is strictly speaking compatible with the speaker having lost significantly more.⁸ One way of understanding this is in terms of a question under discussion (QUD) which reflects the immediate discourse goals of the interlocutors (Roberts 1996/2012). If the QUD is *How much money did the speaker lose last night?*, then a speaker only conforms to Quantity insofar as they provide information that fully satisfies that question.

Connecting this to (3) the idea would be something like: talking about all the objects in the room is strictly speaking more informative than talking about all and only the green objects in the room. But we do not always try to say what is most informative, but instead, as Quantity states, cooperation involves finding the right balance of informativity. The level of informativity one ought to meet is determined by what is needed for the purpose at hand. So, we need to know what the conversational goal or purpose is, in order for Quantity to be applied.

One could supplement this explanation with appeal to another, even more powerful Gricean maxim — that of Relevance — in an attempt to flesh out a general pragmatic explanation: Make your contribution relevant to the purposes of the conversation.⁹ Applied to objectual domains, this would require that one talk about the objects that are (most) relevant. So, if we are talking about things in the room, the relevant books are those in this room. If our conversation is about green objects, then, by the same maxims and the cooperative principle, one should talk about the green objects as these are the objects that are most relevant.

But even if this is correct it is no explanation of *why* certain domain restrictions are defaults. Relevance might be able to tell us why one restriction rather than another is in effect, but we want to know what makes something relevant by default. For instance, why is it that implicitly restricting to objects in a room is so much easier or more often permissible or more natural than restricting to green objects in a room or to beautiful objects? This is what we are seeking an answer to when asking why it is that certain domain restrictions seem to be defaults. Relevance, Quantity, and the Cooperative Principle do not answer this question.

⁸ There is an interesting open question whether this is the same explanation as the one needed for cases of 'loose talk' in which (6) might be used to convey that the speaker lost roughly \$100; see Hoek 2018.

⁹ Grice (1975) calls this 'relation' when introducing his maxims.

To re-construe our question in terms of a Gricean pragmatic explanation, we are asking why it is that conversations with an aim or purpose relating to how things are in the nearby environment are defaulted to over other purposes of conversation. If the default purpose of conversation were “how are things near us?” then the Cooperative Principle and Relevance and Quantity could do their work in explaining why domains are implicitly restricted as they are. But the Gricean framework does not require that this is the default purpose of conversation and so, in and of itself, it does not offer an explanation of why certain domain restriction possibilities are default. Thus it requires supplementation of some sort or other.

The same thing goes for other pragmatic frameworks like the Rational Speech Act framework (RSA). RSA provides an account of how speakers and hearers use Bayesian inference to recursively reason about what state of the world is likely to be like given that a speaker made some utterance while reasoning about how a listener is likely to interpret that utterance. A simple RSA model might assume a uniform prior probability distribution on which objects (or restrictions) are all equally likely. But, RSA is compatible with it being the case that our prior probabilities supply a higher likelihood to restrictions being made by, say, location than by history, how funny something is, or its color. For instance, Frank & Goodman (2012) and Qing & Franke (2015) consider priors related to perceptual salience (how attention-grabbing something is in one’s environment). Consider a simple conversational setup in which a speaker and a hearer are engaged in using signals to refer to some object in the array shown in *Fig. 3*.



Figure 3. From Frank & Goodman (2012)

In the conversational reference game a speaker can only use a single word “green”, “blue”, “circle”, or “square” as a signal and the listener is tasked with determining which object they meant to signal. In one study, participants played the role of a listener. A speaker issued an utterance in what they were told was a foreign language. The listener then had to choose which object to select as the referent, given the utterance. Speaker choice in this condition was used to determine perceptual salience. They found that participants took the object with a unique color to be most salient (i.e., the green square in *Fig. 3*), followed by the object with a unique

shape (blue circle), and finally followed by the object without a unique color or shape (blue square), which was taken to be the least salient (Qing & Franke 2015). These salience priors could be added to in an RSA model predicting speaker utterances and hearer interpretations.

Could a similar explanation be given for default domain restrictions? That is, could the answer be that certain restrictions (current time and place) rather than others (beautiful, red) are available by default is to be explained in terms of higher prior probabilities where this is to be accounted for in terms of perceptual salience? It is worth noting, first, an important difference in the sorts of reference games just considered and cases of domain restriction. In modeling interpretations about reference, a domain of objects is taken as a given. The objects that are there to quantify over are stipulated like in the simple diagram in *Fig 3*. In contrast, what the domain of objects is is what is at issue when considering domain restriction and RSA models have not been used to account for this. Nevertheless, an account in terms of prior probabilities of various restrictions might still seem promising.

As we'll argue below, we think that perceptual salience (as well as perceptual availability and salience related to manipulability) are important features to explain why certain domain restrictions are defaults. However, we do not take this to be *explained* by RSA theory. To see why, consider again how salience is incorporated into an RSA model. Qing and Franke (2015) used an empirical test they took to reveal perceptual salience and added this to RSA models in terms of modified priors (i.e., priors related to perceptual salience). The model does not explain why perceptual salience should relate to increased prior probability. Further, recall that color did not seem to be a default domain restriction. Without further conversational set-up, an utterance of (1) to pick out the red book on the table in *Fig. 2* was not felicitous. So more needs to be said about how perceptual salience (e.g., in terms of color) relates to other factors that might affect utterance interpretation.

It might be suggested that conversational participants determine a domain of objects in precisely the way that they determine the referential content of demonstratives, names, etc. That is, in addition to using pragmatic reasoning about speaker choice to determine which of a domain of objects a speaker has referred to, we engage in precisely the same reasoning to determine which domain of an (unspecified) set of domains they take to characterize the conversation to begin with. But even if domains are 'selections' in a sense relevantly similar to

reference, we are left with no explanation of the psychological features that explain *why* someone selects one domain for discussion rather than another.

Even if an RSA model could assign prior probabilities to restrictions so that default domain restrictions are more probable than non-default restrictions, this does not explain why we default to those restrictions rather than others. More needs to be said.

2.2 A Discourse Structure Explanation?

Perhaps a related account with richer resources connected to discourse structure can provide an explanation. Here's a first pass at how this might go. One might hold that there is a default conversational goal and then take it that that goal along with features about rational cooperative conversation explain why there is a default to figure out how things are "nearby" (e.g., in this room, in this city) before moving further afield (e.g., how things are with blue things or beautiful paintings).

In considering this view, we first ought to ask whether there is a default conversational goal. Grice, Stalnaker, and others assume there is (although sometimes making it clear that it is an idealization) and take the goal to be figuring out the way the world is via the exchange of information. For instance, Grice states:

The conversational maxims ... are specially connected (I hope) with the particular purposes that talk (and so, talk exchange) is adapted to serve and is primarily employed to serve. I have stated my maxims as if this purpose were a maximally effective exchange of information; this specification is, of course, too narrow, and the scheme needs to be generalized to allow for such general purposes as influencing or directing the actions of others. (Grice 1975: 28)

Stalnaker states that "The point of a discourse - at least one central kind of discourse - is the exchange of information" (Stalnaker 1998: 280). Put slightly differently he says "To engage in conversation is, essentially, to distinguish among alternative possible ways that things may be. The purpose of expressing propositions is to make such distinctions" (Stalnaker 1978: 85). Finally, Roberts develops her theory of discourse structure relying on the idea that the primary goal is to answer what she calls "the Big Question" which is "What is the way things are?" (1996/2012, 6).

While we, following many others, take focusing on information exchange as *the* aim of conversation to be too narrow and idealized, we'll suppose for a moment that conversation in general is aimed at this.¹⁰ Given this assumption, can a discourse structure account answer the question as to why, say, locational restrictions are preferred to historical, aesthetic, or color restrictions? We will argue that it does not. To anticipate our argument, the general point is that it will be relevant to a goal of figuring out what the world is like to address how things are with things in the room and how things are with blue things in the room and so on. Each of these will provide information narrowing down the way the world is. So a theory that relies on relevance to a conversational goal—in the sense of providing a partial answer or strategy for answering—will not be able to account for the differences in restriction possibilities we outlined above.

To build up this response, let's consider Roberts' view of discourse structured through questions under discussion in more detail. Roberts takes it that discourse involves goals that are carried out via conversational moves that are components of broader strategies. As we saw above, she takes the main goal of conversation to be to answer the question "What is the way things are?" (1996/2012, 6). As with achieving any ambitious goals, a natural strategy will be to meet less demanding goals along the way. Roberts spells this out in terms of answering subquestions which are logically related to the Big Question. The answers to subquestions are partial answers to superquestions. They narrow down the range of possible answers to the Big Question and so answering them helps to achieve our main conversational goal.

Let's look at an example. To answer the Big Question one needs to answer the question "What is the way things are in California?" and more specific subquestions like "What is the way things are in Los Angeles coffeeshops?", even more specifically, "What is the way things are at Civil Coffee on Figueroa today?", and more specifically still, "What is the way things are with blue coffee mugs at Civil Coffee on Figueroa today?".

¹⁰As noted above, some of these theorists explicitly note this, for instance, as seen in the quote from Grice. Directives, like requests and commands, make it clear that this is too narrow. One of the things we might be figuring out, for instance, is *what to do*, rather than how things are. This was a central contention of Austin (1962) and Wittgenstein (1953) and has been a prominent line of criticism of philosophers of language working in social and feminist traditions (see, e.g., Langton 1993, Kukla 2014; McGowan 2019, Keiser 2022).

Now, let's consider whether this picture can explain the difference in availability of default and non-default restriction possibilities we began with. That is, can it explain why some restriction possibilities have a default status? To partially answer the question about Civil Coffee, one might utter:

(7) The front door is locked.

While there is not one unique front door in the universe, given the question under discussion (7) is plausibly interpreted as equivalent to something like (8).

(8) The front door of Civil Coffee on Figueroa is locked.

Here the implicit restriction is to the location information is being queried about and so a restriction to the only front door at that location is available. Suppose that the tables at Civil are mostly square, but there is one round and one rectangular table. Part of answering the question "What is the way things are at Civil Coffee today?" involves answering how things are with the round things at Civil. Since there is just one round table, in order to give a partial answer to this question we need to know how things are with that table. So, a conversationalist utters:

(9) The table is wobbly and has three laptops on it.

Yet, (9) is not going to be easily interpreted as expressing something equivalent to (9').

(9') The round table at Civil Coffee on Figueroa is wobbly and has three laptops on it.

This is the case even though answering how things are with round things at Civil (or round tables at Civil) does provide a partial answer to how things are at Civil (and, more generally, how things are in the world). The structure of conversation and the entailments between answers to subquestions and superquestions like "What are the way things are?" fails to provide the resources to explain why it is that certain restrictions are available as defaults.

In the context of a discourse structure based account, we can understand the issue of default restriction possibilities as about constraints on the kinds of *questions* and *subquestions* we are liable to treat as by default operative in conversation. The Big Question—"What is the way things

are?”—might be a default question guiding conversation, but it is too big for much inquiry to get going without subquestions. So, as Roberts puts it, “A reasonable strategy will involve a plan to” answer it “by developing subgoals which are easier to achieve and are logically related to each other in such a way as to facilitate achieving the main goal” (Roberts 1996/2012, 6). Our guiding cases seem to show that some such questions are more likely to find their way to the top of the question ‘stack’ than others. So, in the context of a QUD framework, we can understand our initial question as a question about why we default to certain joint subgoals. That is, why do we default to the joint subgoal of figuring out what’s going on with the books in this room, rather than with cookbooks or the green books in this room or Discourse structure alone does not answer this question.

2.3 An Intentionalist Explanation?

Let’s look at one last explanation now in terms of a metasemantic account rather than an account of rational communication guided by pragmatic principles or discourse structure. One might argue that an account based on speaker intentions can explain why certain domain restriction possibilities are defaults.¹¹ To see how an intentionalist explanation might go, let’s consider a version of a Gricean view of meaning. On a Gricean view (Grice 1957, 1969), a (declarative) utterance *u* by a speaker *S* means *p* just in case *S* intended:

- (i) for their addressee(s), *A*, to come to believe (or believe that *S* believes) that *p*,
 - (ii) for *A* to recognize *S*’s intention to get *A* to believe (or believe that *S* believes) that *p*,
- and
- (iii) that (ii) be part of *A*’s reason for coming to believe (or believe that *S* believes) that *p*.¹²

Applied to domain restriction, the basic idea is that while a particular sentence type like our original (1), repeated below, could be used to express various things, a particular utterance of it

¹¹ As a variant, one could argue that it is audience expectations about speaker intentions that determine meaning or what is said (see Austin 1962, Kukla & Lance 2009). This variant would not change the upshot of the arguments we make in this section, so we won’t consider it further here.

¹² Many have developed other versions of intentionalist accounts of meaning to cover a broader range of cases (e.g., Bach & Harnish 1978 extend intentionalism to speech acts other than assertion, see also Harris 2014, 2020) and to solve problem cases for Grice’s own intentionalist view (Schiffer 1972, Recanati 1989, Buchanan 2018, Stalnaker 2014). Intentionalism about meaning is also appealed to by many in formal pragmatics, e.g., Stalnaker and Roberts, whose theories provide constraints on the sorts of rational intentions a speaking agent may have. It does not really matter, for our purposes, how exactly the intentionalist program gets spelled out: what we’re interested in are the default interpretations that seem to pervade human communication.

by a speaker expresses a more restricted proposition given the speaker's meaning intentions (Buchanan 2014, Neale 2000).

(1) The book is on the table.

A speaker might utter (1) intending to express something about the book before them on one occasion and a different utterance on a different occasion might express a different proposition given different meaning intentions.

This account has the resources to explain how it is that domains are restricted—they're restricted via speaker intentions—but we can appeal back to exactly the same criticism we gave of the accounts that appeal to norms governing discourse: it does not explain *why* certain domain restriction possibilities are defaults. On an intentionalist account it is speaker intentions (of a particular sort) that determine meaning. An intentionalist could state that speakers more often than not intend to express propositions about things in their proximal environment. We are open to this being part of the story, but it does not explain why it is that certain domain restriction possibilities are defaults. Rather, it attempts to answer our question by pushing it back a level and holding that speaker intentions tend to, as a default, be about things in the proximal environment. But, why is that? Our question could then be asked again, now framed in terms of why speakers default to certain sorts of meaning intentions (e.g., about the local environment) but not others (about, e.g., subkinds, historical features, or colors of objects in their environment).

In the next section, we offer an answer that turns on the ways cognitive constraints and heuristics interact with our planning capacities. In so doing, we provide an explanation that does not beg the question against default domain restriction possibilities, by assuming that there are default purposes of conversation, subquestions, or meaning intentions. This is not to say that we reject any of these accounts. Rather, we take it that any one of them must be supplemented with a cognitive account like the one we offer here in order to more fully explain the way domain restriction functions in communication.

3. Joint Plans and Cognitive Heuristics

3.1 Joint Plans

We have argued that several major frameworks in and of themselves, do not provide an explanation for why some domain restrictions are available as defaults. In this section we first consider the role of representation in joint planning, particularly as it relates to domain restriction. Our aim in doing so is to sketch a view that we take to capture many of the shared insights in the three types of views just considered while remaining as neutral as possible. Like these views, the joint planning view we offer cannot explain why there are certain default restriction possibilities. We then go on to argue that an explanation for why we default to certain domain restriction possibilities and not others can be given in terms of cognitive features we deploy in navigating and influencing our (social) environments and engaging in joint activities.

We follow other theorists, including those discussed above, in holding that in conversations we (at least very often) have shared aims, goals, or purposes (see also, Gauker 1997, 2003). These can be represented by conversationalists and involve a form of joint action and joint planning to meet them. How the account is spelled out in more detail, depends on two choice points. First, it depends on the view of collective intentionality one adopts (e.g., Bratman 1992, 2014; Searle, 2010; Gilbert 1989, 2006). Second, the details of the view will vary based on what pragmatic and metasemantic accounts one adopts. For instance, joint plans and actions might relate to following the cooperative principle, to structured inquiry in discourse, to speaker intentions and constraints on intentions, to some combination of these, or additional factors. Our aim is not to settle these questions or argue that one account is best. Rather, our goal is to provide a minimal starting point relating to joint planning, attention, and relevance so that we have a foundation on which to consider the cognitive heuristics we take to provide the crux of the explanation for why there are default domain restriction possibilities. Thus our account is amenable to each of the frameworks discussed in the previous section.

Applied to domain restrictions, we take the following initial characterization of how representation relates to joint planning to be promising: the domain of quantification tracks the objects that are relevant to our joint purpose. One way of understanding this is in terms of the interlocutors' current joint domain goal or goals (see Gauker 2003; Roberts 2012: Afterword).

For instance, take as a case the situation in which Leto says to Duncan,

(10) Everyone is here.

And assume further that in this case Leto's intended message is extensionally equivalent to something like,

(10') Leto, Duncan, Amir, Jessica, and Yueh are here.

When Leto utters (10), it conveys something extensionally equivalent to (10') because the individuals mentioned in (10') are all the people that are relevant to a joint purpose. This may be because – for instance – Leto and Duncan are discussing potential co-conspirators, and the question,

(11) Who is a potential co-conspirator?

has the individuals mentioned in (10') as its extension (Roberts, 1996/2012, 2012; Kratzer 2023; Schwarz 2019).

Relevance is an intentional notion in at least the following sense: to say that some object or individual *a* is relevant to some task is not to say that *a* would be (e.g.,) good or useful to that task. Rather, it is to say that *a* is in the extension of objects determined by a question or planning state.

Of course, something's being 'out of sight and out of mind' does not mean that it is not relevant: if Lucy is guest lecturing for Laurenz, and asks what his students are like, Laurenz might respond by saying that they're all well-prepared. This is meaningful to Lucy—and we imagine she would form a representation that determines the relevant students (the ones in Laurenz's class, specified via that property, perhaps) despite her never having met any of Laurenz's students.¹³

¹³ We do want to emphasize that discourse participants are not infallible when it comes to the domain of discourse. Imagine a simple scenario in which Lucy and Laurenz are cleaning up the living room, and Lucy says to Laurenz,

(i) Every book is on the table.

So we want to say that the domain of quantification is what is relevant to our joint purpose (or aim, or goal), where this is broader than just those objects we see or are in other ways acquainted with, but narrower than all of those objects it might be helpful or useful to know about. The domain can be intensionally specified by one or more properties.

To hone in on the notion of what is relevant to joint planning a bit more, it will be useful to distinguish it from a related notion of *salience* in the sense of being attention-grabbing (see also, Gauker 2003, Ch. 4). There's some reason to think that what is salient is separable from what is relevant to joint planning. Further, there is reason to think that the latter rather than the former is doing the work in domain restriction cases. First, note that attentional salience and conversational relevance come apart in all sorts of cases. If a goat walks into the room (to borrow one of Stalnaker's (2014) examples) while we're talking about something else, the goat might grab our attention and we might take the goat to be something we have reason to attend to, without thereby taking the goat to be relevant to our original conversational aim or purpose.¹⁴

In general, what we might think of as perceptual *salience* and *relevance* come apart.¹⁵ Things like deixis appear to track salience: when the goat walks in, I can say,

(13) He looks angry.

to immediately and unproblematically refer to the goat.

In contrast, definite NPs are more responsive not to changes in perceptual salience, but instead to changes in what is relevant to us given our (conversational) goals. Consider a case to draw out the salience/relevance distinction. We're visiting a museum with plans to see a particular

Now imagine that there is a book under the couch that neither Lucy nor Laurenz can see and that neither know about, but which is relevant to their goal of straightening up the room. This hardly seems to make a difference to what gets conveyed by (i), but Lucy has said something false, and moreover she and Laurenz would acknowledge this were they to learn of the mistake.

¹⁴ We take care here to draw a distinction between what is relevant simpliciter and what participants in a discourse represent as or take to be relevant. Relevance is determined by the question under discussion or conversational subject matter, and there are various ways in which we might fail to represent or attend to (or misattribute attention) whatever satisfies this criteria.

¹⁵ Roberts (2022) makes a similar point in relation to reference determination. Some notions of salience are closer to what we take to be relevance. We won't quibble with terminology here, but think there are clear ways that perceptual salience and relevance to a conversation can come apart.

painting. It's not the most famous painting in the museum, and in fact when we arrive at the gallery we see that there is a different painting surrounded by a crowd with a spotlight shining on it. That painting is most salient in important perceptual and social ways. An utterance of (14), involving a deictic use of a demonstrative, plausibly picks out the spotlighted painting.

(14) That one is overrated.

Here, the complex demonstrative will naturally pick out the most salient and attended to painting, even though it is not that which is most relevant given our plan in visiting the museum to see a lesser known painting. However, we can use "the painting" in (15a and b) to pick out the painting that is relevant to our plan, even though it is not, in some general sense, the most salient object in this class.

(15)

(a) Wow, the painting is even more amazing in person.

(b) Do you know where the painting is?

These cases show ways perceptual salience and relevance can diverge.¹⁶ Given our discussion, we suggest that people take restrictions on domains of objects to be available when they restrict to those things relevant to joint planning and joint goals.

This view alone, like related Gricean, RSA, and QUD frameworks, cannot answer the question why certain domain restrictions are defaults. History, color, location, aesthetics, or any other feature might be relevant to joint plans or goals. In order to answer our main question, we now turn to supplementing this account by considering cognitive heuristics and connections between linguistic, social, and spatial cognition.

¹⁶ There are additional complications related to competition between resolutions of definite NPs. For instance, while we take it (15a) can be used to pick out the painting we plan to see, it also requires having already seen the painting, given the predicate. Suppose instead that we are looking at the popular overrated painting. One of us then says:

(i) Ok, I'm ready. Let's go see the painting!

While we take it that the definite NP in (i) could be interpreted as being about the painting we have a plan to see, this interpretation is a bit less natural, given competition with the presently salient resolution of the NP. In contrast, if upon arrival at the museum one of us utters (i), the definite NP would be easily resolved as picking out the painting we have plans to see. Our claim is that there is a contrast between demonstratives and definite NPs and that these are resolved more in terms of salience and relevance, respectively. We do not, however, take that to be a complete story. Thanks to Craige Roberts for helpful comments and suggestions along this line.

3.2 Cognitive Heuristics

Our dispositions towards our environment — even our representational dispositions — derive from the demands made on us as agents by unpredictable environmental factors like unpredictable prey, as well as combative and cooperative conspecifics (see Tomasello 2022 for an accessible overview). The idea that what and how we're disposed to represent is in some way guided by general psychological features and principles is not new. It has been extensively discussed when it comes to how we represent things like properties of categories of objects and agents (e.g., in terms of prototypicality effects, Rosch 1975), communicative signals (Sperber & Wilson 1986), and even more abstract things like probability (Tversky & Kahneman 1974; Kahneman & Tversky 1979) and cost-benefit analyses (Angner 2020). In many domains, these can lead to situations where there is a tension between what feels natural and normal, on the one hand, and what is epistemically rational (e.g., given base rates or that category membership in categories like ODD NUMBER or BIRD is discrete) on the other.¹⁷

Now we want to ask what cognitive features govern how we represent and communicate about objects. We argue that which implicit domain restrictions are defaults can be explained in terms of features that shape how we navigate and influence our (social) environments and engage in (joint) actions. Humans are creatures that aim to understand and explain in order to facilitate prediction, manipulation, and control of our social and natural environment (Lombrozo & Carey 2006; Strevens, 2007; Woodward 2003). We engage in many of these projects together—from understanding the nature of fluid dynamics or the human mind, to building bridges, raising children and making meals. Communication is at its heart a joint endeavor, one which we take to be in the service of the more general projects creatures like us have, namely, understanding, predicting, and manipulating.

We propose three cognitive heuristics which relate to general ways we represent and categorize. Each fits into a view on which humans are understood as creatures seeking to carry out these broad aims relating to prediction, manipulation, and control in communicative, and

¹⁷ What feels natural and normal may nevertheless come under descriptions like 'resource rational' and 'ecologically rational'. Often we engage in reasoning — guided by heuristical reasoning tools — that allows us to approximate the behavior of a Bayes-rational agent; sometimes even resulting in behavior that is more accurate to Bayesian reasoning than reasoning we could have produced if we tried (Gigerenzer & Gaissmaier 2011). We are happy to accept that the heuristics we sketch enable us to engage in conversations in a way that does, indeed, maximally exploit our environment, or allow us to approximate the behavior of Bayes-rational agents. This is a separate claim, however, from that of RSA models on which speakers and hearers are claimed to genuinely engage in Bayesian inference.

therefore joint, interactions. In each instance, we'll explain the feature and why it is plausible, provide some empirical data supporting it, and then apply it to domain restriction. We propose that these features determine default domain restriction possibilities. We then provide further justification for these features by considering an additional cognitive feature—objectivity—and suggest further support for locational restrictions due to connections between language and spatial cognition.

The Perceptual Availability Heuristic

Earlier in this section we drew a distinction between something's being perceptually salient and something's being relevant. Just because some object *o* is available for reference in the sense of being something we could attend to doesn't mean it's relevant to a task or goal. But perceptual salience and relevance are not unrelated. Let's start, though, with a notion that is broader than salience—*perceptual availability*. Some things are perceptually available to a speaker *S* just in case they can be perceived in the environment in which *S* finds themselves with minimal distortions of their body (e.g., turning one's head, but not walking 50 feet).

For instance, let's say you're in a room with a bunch of green books visible on some bookshelves. However, you know (and in fact it is common knowledge) that there are numerous books in the room that are hidden from view. Utterances like (16) and (17) nevertheless strike us as felicitous, or at the very least as understandable: they are utterances about the books that are available to attend to without manipulating other things in the environment by, e.g., opening closet doors or drawing back curtains.

(16) They're all green.

(17) Every book is green.

This isn't surprising. For one thing, objects on which we can cast our attention are often relevant, especially in tasks involving a lot of open-ended exploration and manipulation of the physical environment.

Objects that are in the current location at the time of conversation are the only objects that are potentially perceptually available, without special tools like cameras. Given our physical perceptual capabilities, meeting these conditions—being in the here and now—must be met in

order to be perceptually available. The perceptual availability heuristic provides the first building block for a cognitive account of default implicit domain restriction possibilities.

Two Saliency Heuristics

Going further, on our account other properties come into play to filter the perceptually available objects, thereby providing a (potentially) smaller domain of objects. The key feature is saliency. Some objects that are available to attend to will stand out from the rest—these objects are very often natural ones to restrict a domain to.¹⁸ If you're faced with a task in an environment, it is natural to default to solutions to that task that involve attention-grabbing objects. This relates to the first of the saliency heuristics—the *perceptual saliency heuristic*. Saliency is manifest in a different way in a second saliency heuristic—the *manipulability heuristic*. To say that an object is manipulable is to say something about its perceived availability for control (i.e., I can control a hammer), manipulation (i.e., I can move a chair around), and interaction (i.e., I can press the buttons on an interface). Manipulability is a way of being salient. More manipulable objects are more salient.

Many properties related to being perceptually attention-grabbing and manipulable relate to locational features. For instance objects that are nearer to the speaker rather than those that are available to attend to, but are in the distance, are more salient in both ways and make for plausible default domain restrictions. This will particularly be the case when conversations occur in more open environments with less discrete boundaries. For instance, consider the difference between uttering (18) in an enclosed room versus in a large open meadow.

(18) Every plant is blooming.

Uttered in a room, plants very near to the speaker and plants across the room are both plausibly within the default restriction. If some plants near the door are not blooming but those within arms reach of the speaker are blooming, we take it (18) will still not be taken to be acceptable.

¹⁸ This isn't always true, of course; consider this example from Roberts (2010: 25):

“[Context: You and I are sitting in a café discussing how to understand Sperber & Wilson's (1985) definition of Relevance, and I say:]

(i) I see it now!

[Even though I'm holding a coffee mug by the handle right under your nose and shaking it for emphasis, you don't take *it* to refer to the concrete mug.]”

However, if some plants visible through a window are not blooming, these do not seem to affect the acceptability of (18). Why is this? We take it the answer relies on our representations of places and how places constrain what we can see and do. Rooms have discrete boundaries. In the standard case these are opaque, fixed, and limit actors movements. They affect what is perceptually and physically available for us to perceive and interact with. This relates to perceptual availability at a time, but it also lends support to salience heuristics as filters on what objects are available to quantify over. Perceptually available objects that are not within the boundaries of the room are less perceptually salient—they are further away and perhaps less attention grabbing. They are also less manipulable. Both salience heuristics filter the domain of perceptually available objects, thereby further restricting the default domain.

Contrast this case with an utterance of (18) made in an open meadow. If all the plants within the parts of the meadow near the conversationalists are blooming but some at its edges are not, we take it (18) is likely to be acceptable. Perceptual and manipulable salience, as well as the ways locations are represented as being bounded within spatial cognition, a topic we consider further later in this section, filter the perceptually available objects.

Other properties are relevant to salience too. These might relate to properties of objects including their color or size. The sorts of things that grab our attention are likely to be the sorts of things that we've adapted to attend to, because they have historically been relevant (Gibson 1977). Certain shapes and colors – for instance the shape of a snake, or the color red – plausibly stand out because noticing these sorts of things has historically been indispensable to our survival. And an object having a unique feature (like in the example involving Fig. 3 above) can also make it more perceptually salient.

This has a profound social component as well. For one thing, when we attend to something, we often assume that others attend to that thing as well. This tendency results in things like egocentric biases and the spotlight effect, but it is also an indispensable tool in communication. For instance, assuming that some thing(s) are mutually salient is part of what explains how we communicate with deictic referring expressions (Clarke et al. 2013; Rubio-Fernandez 2019). This isn't just a one-sided attentional bias either, research finds that speakers' production of referential expressions involves considering their audience's sequential visual search patterns and can, therefore, include additional adjectives for highly salient features, like color (Jara-Ettinger & Rubio-Fernandez 2022). For instance, speakers might use "the blue square" to refer

to the only square in an array of red circles and triangles. We also instinctively use certain bodily cues, like gaze direction, to figure out which objects others' are attending to in order to determine reference (Baron-Cohen et al 1995; Hanna & Brennan 2007; Garoufi et al. 2016).

Empirical work on language and perceptual availability and salience has focused on the resolution of referential terms, but we take it that these are also cognitive heuristics used to implicitly restrict the domain. That is to say, we use evidence of other people's attentional states and what we assume is shared and mutually salient to establish a domain of discourse given our aims to predict, understand, and manipulate together. Most linguists and philosophers of language do not take 'every book' to be referential and we do not intend to argue that it is here, rather our point is that what everyone agrees is important for establishing reference is also important for establishing a domain of (possible) referents.

So far we've focused more on perceptual salience than manipulability. Considering how we naturally interpret certain instances of domain restriction provides further intuitive evidence for the manipulability heuristic as a salience heuristic. Imagine a group of people staying in a secluded country house. They walk out into the backyard on a cold winter night after a snowfall and someone says,

(19) Everything's covered in snow!

and it is: all the lawn furniture, the potted plants, and the garden gnomes. But not the clouds and not the moon. Maybe not the trees visible in the near distance. Do these facts make (19) false? Maybe for the pedant. But there's an intuitive reading on which it expresses something acceptable and potentially true.

We think manipulability imposes further restrictions on perceptual availability, again through limiting the perceptually available objects via features related to salience. In fact, in cases where we seem to be depending on perceptual availability to construct a domain, we tend to default to those visually available objects that would be manipulable, with which we could interact, were we to close in on them and be near them. As with perceptual availability, manipulable objects have a kind of prima facie relevance to certain tasks. In very simple, situated tasks (like chasing something) questions about what we can change in our environment are important.

Notice that given our physical limitations, an object being located nearby is required for it being manipulable (again setting aside special technology or tools like drones). Being present is also required given that we cannot manipulate objects in the past or future.

Claims about the relevance of manipulability are connected to claims made in 'ecological' and 4E approaches to psychology and cognition (Hurley 1998, O'Regan & Noe 2001). But we don't need to go even as far as saying that agents 'perceive' affordances (i.e., are confronted with perceptions that 'instruct' them how to interact with their environments). There is some evidence that what we're in a position to notice or attend to within our environments can be affected by our ability to interact with those environments.¹⁹ For example, more easily controlled tools are understood and judged as nearer (Wakslak & Kim 2015). We think that changes in manipulability may also correspond to changes in what objects we take to be potentially relevant to our shared goals. For instance, suppose we enter a room with our bare hands and one of us says,

(20) We're going to destroy everything.

Intuitively, what has been communicated is something about smaller objects and maybe some pieces of furniture — things that can be destroyed with one's bare hands. But now imagine that we enter the same room armed with humongous sledgehammers, and one of us utters (20). It seems that what has been communicated has changed. The sheetrock walls now seem to be included in the domain of *everything* even though they were not before. Given the general aims of predicting, understanding, and manipulating, what is represented as potentially manipulable, is a psychologically natural heuristic affecting implicit domain restriction possibilities.

The perceptual availability and salience heuristics facilitate coordination to meet the very broad aims we take humans to have and which communication as a joint action facilitates. Without further information about conversationalists' specific beliefs or shared idiosyncratic discourse goals, the information we go into a communicative exchange with is quite limited; we can assume very broadly that we are aiming to do something related to understanding or changing the world. Given this dearth of information, shared aims will be best met by coordinating on objects that are perceptually available, perceptually salient, and that are taken to potentially be

¹⁹ This is much weaker than the (very) controversial claim that *perception* can be influenced by top-down effects like cognition and ability (Firestone & Scholl 2015).

within our control. We argue that these cognitive features determine default domain restriction possibilities and explain why locational and temporal features, like presently being in the room one is in, are apt for implicit domain restriction. Let's spell this out a bit more concretely.

Objects in one's current spatial and temporal location are those that are likely to be perceptually available. The property of being nearby makes an object more likely to be perceptually salient than having been made in Prague, being brown, or being constructed out of concrete. That is not to say some of these other properties could not increase salience (many colors do, of course), but being in one's spatio-temporal vicinity is one property that scores highly as making an object both available and salient. Being in the vicinity of the conversationalists is, in ordinary scenarios, also a requirement for being manipulable. Other features can affect manipulability (size, weight), but if an object is not nearby (without access to technological innovations developed only in our very recent history) it simply will not be manipulable.

Perceptual availability and salience in both perceptual and manipulability forms individually and collectively support that locational and temporal features related to being present in the location of the conversation are default domain restriction possibilities. This is so not just because they are, say, locational features, but *because* they score highly on each of these cognitive features. For availability it is only present nearby objects that score highly at all. For salience properties, availability is required, so these features too require the present and nearby restrictions that availability yielded.

To provide further support for the account on which these heuristics provide default domain restrictions, we next consider two further cognitive supplements.²⁰

Adding Support: Objectivity and Spatial Cognition

When it comes to cooperating with others, we tend to sort the world into categories that can be understood as (relatively) objective and not (very) perspectival or subjective. The sort of

²⁰ We do not want to rule out the possibility of additional heuristics, and certainly there are others that are conceptually available. The heuristics we propose are tailored to the 'design' of human cognitive / perceptual systems. But consider a perceptual system more like a dog's, which relies on an olfactory system that is able to hone in on objects that were *previously* in a location. Perhaps an advanced communicative system that relies on this kind of perceptual organization would deploy different sorts of heuristics that pick out entirely different regions of spacetime as defaults than those we have argued ours do.

objectivity we are focused on here is objective in a psychological sense; it is about what people *take* to be objective. A person might take features related to location and material constitution as well as social features like race, gender, and nationality to be objective. Other features, like tastiness, beauty, and apparent size are taken to be less objective. One marker of this is the fact that disagreement over whether such properties are instantiated is typically considered to be faultless (Kölbel 2004; Scontras et al. 2017). In faultless disagreement tasks, people are asked whether a speaker, A, who says *p* and a speaker, B, who says *not-p* could both be right. Stronger judgments that both could be right are interpreted as judgments that the disagreement is faultless and that the feature (e.g., beauty) is presumed to be subjective rather than objective.

Objectivity obviously matters for interpersonal agency. By focusing on more objective features, coordination and successful joint planning and action is easier. This isn't to say that we don't care about coordination on 'subjective' standards, but merely that without information about more specific goals, it is natural to rely on more objective standards.

Evidence for the value of objectivity in communication comes from work on adjective ordering and reference resolution. Scontras et al. (2017) find that there is a strong bias towards adjective ordering that places more objective adjectives (e.g., 'blue', 'wooden') closer to the modified noun than more subjective ones (e.g., 'beautiful', 'big').²¹ To most people, "the big blue box" or "the beautiful wooden table" sound more natural than "the blue big box" or "the wooden beautiful table". The same pattern is cross-linguistically attested across a wide range of unrelated languages (*ibid.*). The explanation for this that they propose? That "less subjective content is more useful at communicating about the world" (2017: 64). Scontras et al. (2019) followed up on this work arguing that reference resolution is more effective with less subjective adjectives as they are better means to limit referential search (cf. Simonič 2018, Franke et al. 2019).

We take it that objectivity helps to explain why default domain restriction possibilities are determined by cognitive heuristics related to what is perceptually available, perceptually salient, and manipulable. Suppose the only information we have about our current state of interaction involves a very basic understanding that our aims involve, in some way or another, predicting, explaining, or affecting change. This is the sort of scenario we envisaged in offering our original contrast case involving one or three books on a table. Objectivity captures the psychological

²¹ Subjectivity was assessed via participant subjectivity ratings and on acceptability in faultless disagreement tasks.

naturalness of defaulting to objective rather than subjective features when there are no more specific goals or discourse aims already in place. More objective restrictions, like being in the room we are now in, will serve as better implicit restrictions than restrictions like being beautiful or being large. This is because more objective restrictions will be more likely to be shared among conversational participants and to be taken to be shared among participants than ones that rely on perspective, taste, or evaluative judgment. Joint endeavors to predict, explain, or manipulate our environment together will be better facilitated by defaulting to features assumed to be objective. Further, given that in many conversations only minimal assumptions about shared knowledge can be made (i.e., about what one expects their interlocutor to also know or believe) features related to location and presence are plausibly top amongst what is assumed to be objective and shared.

While other features might be equally objective (e.g., the material something is made of or its color), we take it locations of objects in space and time will be taken to be highly objective and more likely to be shared. This is different than, say, that it is an objective matter that yuccas are more cold hardy than the California lilac. While this is the case, and disagreement about it would not be faultless (one or the other person is wrong if they say opposite things), it is not plausible to assume that this is shared knowledge or relevant to conversationalists goals without further evidence. In contrast, even with minimal evidence it is reasonable to assume features related to what is presently in the environment are shared and objective and relevant to conversational goals. This makes these better default implicit restrictions, helping to explain why the cognitive heuristics we offered above are those that determine default domain restrictions.

Finally, we want to close by considering ways linguistic and spatial cognition might further support default domain restriction possibilities. There is evidence that language and communication are deeply tied to spatial cognition. It has been argued that the syntax of human language might have originally evolved for navigation (Bartlett & Kazakov 2005). Turning to semantics, Levinson (2023: 6) notes that a fifth of most common English words are spatial. Prepositions like “in” and “on” are clear cases of words that are (at least on one central use) spatial. Other function words, like “and” might also be systematically connected to spatial representations (e.g., Jackendoff 1983; O’Keefe 2003; Glenberg 2010; Langacker 2008; Guerra et al. 2013).

There is also evidence that components of language processing and learning are implemented in the hippocampus, the region of the brain activated when spatially orienting, navigating, and gesturing to give directions (O'Keefe & Nadel 1978; O'Keefe 2003). The right hippocampus has also been found to grow when people memorize routes, as in the famous London taxi driver study (Maguire et al. 2000). Similarly, the left hippocampus has been found to increase in volume when learning a new language (Mårtensson et al. 2012). While the hippocampus is also implicated in memory, some hypothesize that it is central to episodic memory—including memories of experiences—but perhaps not semantic memory—the ability to retain and recall factual information (Vargha-Khadem et al. 1997; Tulving & Markowitch 1998). Others suggest that the hippocampus is central to episodic memory retrieval as well as the retrieval of semantic memories particularly when those involve spatial information (Hoscheidt et al. 2010). Episodic memory has been argued to be structured by spatial relations both in the ways memories are recalled and in the ways memories are abstracted from in generalization (Aronowitz & Nadel 2023). So, there is some evidence to think the hippocampus—as it functions in navigation, language learning, and memory—might be intimately connected to spatial cognition. Levinson (2023) postulates that gesture could be the lynchpin connecting language and spatial cognition. He argues for an evolutionary account on which early hominid communication relied primarily on gesture and spatial organization implemented in the hippocampus. When communication shifted to be vocalized, he hypothesizes, it retained implementation in the same brain region and a spatial “underlying conceptual framework” (2023: 7).

While much more research is needed, given the potential syntactic, semantic, neural, and evolutionary connections between language and communication, on the one hand, and spatial cognition, on the other, there might be a defeasible linguistic bias to default to spatial and locational features over others (e.g., color, material, or aesthetic properties). This explanation differs from the explanation based on perceptual availability, perceptual salience, and manipulability. We argued that each of these cognitive heuristics facilitate the human aims of prediction, explanation, and manipulation and do so in communication as a part of joint planning and collective goal-making. We then argued that as it turns out spatial and temporal restrictions—to the here and now—meet each feature to a significant degree, thus explaining why they are default restriction possibilities. In contrast, a locational bias like that we've just been considering is, we take it, less related to shared human aims or ways to facilitate joint planning and cooperation and more connected to the language faculty as it has evolved in humans.

Nevertheless, the two sorts of cognitive explanations are compatible and, we take it, both are promising and worthy of further exploration.²²

4. From Default to Non-Default Restrictions

We now turn to our second question: How can we account for the ways non-default restriction possibilities become available? This question can be answered in different ways depending on the more general framework of communication and metasemantics one adopts. For instance, an intentionalist will appeal to speaker intentions given representations of an addressee's mental states. Our answer, which is most closely related to a discourse structure account, relies on relevance to joint planning, aims, and goals. The explanation we give here is also quite minimal, as this question has not been neglected in the way we take our first to have been. Nevertheless, it will be useful for completeness to lay out how we see the second question being answered.

We take it that available non-default restrictions are those that are relevant to a joint idiosyncratic purpose (as determined by intentional planning states). These more particular purposes can come about in two ways. First, they can be introduced through explicit discourse moves. Discourse moves like questions, assertions, and directives change conversational aims or discourse structure thereby changing what restrictions are relevant.²³ Second, a particular purpose can be put in place prior to the initiation of a conversation. Let's start by considering several ways utterance force affects restriction possibilities.

Recall (5) repeated here.

(5) A: Where are all the blue things?

B: Well the shirts are near the register, and the book is on the table.

²² It could be noted here that one variable that seems to play a role in determining what objects are available for domain restriction is just whatever participants in a discourse represent as being 'in their environment' or the 'place' that they occupy. These concepts are themselves in need of explanation, and our view (on which cognitive heuristics drive the representation of objects as relevant for quantification) basically treats them as an epiphenomenon. But an alternative approach (one that we do not have the space to consider in detail) might treat something like a 'place' concept as the genuine driver of default domain possibilities, and what we have labeled as 'heuristics' as explanatory posits in evolutionary psychology.

²³ Other discourse features within various discourse moves—like previous content, *if*-clauses, and focus—can also restrict domains (see, e.g., Roberts 1995; Beaver & Clark 2008).

When introducing this example, we noted that raising the question about the location of blue things allowed for “the book” to be restricted to (something like) “the blue book in the room”. It is not just questions that can introduce non-default restriction possibilities. For instance, consider a case of an assertion and a directive:

(21) A: The blue merchandise belongs in the display by the door.

B: I'll go grab the book and the shirts.

(22) A: Bring the blue merchandise to the display by the door.

B: OK! I'll go grab the book and the shirts.

B's utterances in (21) and (22) can easily be interpreted as restricting “the book” and “the shirts” to the *blue* book and the *blue* shirts, where those items might be located in another room or in a warehouse across town. Discourse moves make non-default domain restrictions available by introducing new more specific discourse goals.

Since our framework relies on relevance to joint aims, adopting particular aims, purposes, or goals even without doing so via a discourse move can affect what domain restrictions are available. It's hard to imagine how this might happen within a conversation. But, depending on how one individuates conversations, it is natural to think that some people might have a particular aim prior to beginning a new conversation. For instance, suppose that Arman and Bea are planning to visit a museum to see a recently acquired Cy Twombly painting. They agree to meet in the room where the painting is being displayed along with a number of other paintings. As Arman walks up to Bea he says:

(23) The painting is even bigger and more erratic than I had expected.

Even though there are many other paintings in the room and some others might be more perceptually salient in ways we discussed earlier (e.g., spotlight, surrounded by a crowd), we take it Bea will easily interpret Arman as using “the painting” to pick out the Twombly painting they intended to see together.

The cases considered here provide ways in which a discourse move or a shared goal, aim, or purpose can affect restriction possibilities in ways that go beyond defaults. We take it, however, that these are still affected by and interact with the principles discussed in the last section. Particular and more specific discourse aims still fall under the superordinate category of aims to predict, understand, and manipulate the environment. Directives focus on manipulation and planning ways to do so (Portner 2004, Kaufmann 2011). Assertions share information to facilitate understanding (Stalnaker 2014). Questions guide inquiry by structuring ways to understand or explain (Roberts 1996/2012). Speech acts like promising and marrying also involve manipulating the social world (Austin 1962), as do uncooperative and oppressive forms of social control like those involved in ranking individuals and groups in hierarchies (Langton 1993).

When a more specific aim is adopted, conversationalists have more shared information to go on when implicitly restricting the domain. This allows for the default domain restrictions delivered by cognitive heuristics to be overridden.

For instance, if two parents have an aim to collect all the library books that their children have checked out, they won't rely just on what is perceptually available or salient. They may need to look under couches, in the car, and in backpacks. In saying "I think we're done! I put all the books on the table" the speaker is not restricting via the cognitive defaults we considered before, but rather by the shared aim the parents have.

Consider, as another example, objectivity. More objective features are easier to coordinate on than less objective ones. This will be the case even for idiosyncratic goals. For instance, library books checked out by one's children is a more objective restriction than library books that I enjoyed reading. But less objective features can still make their way into restrictions. For instance, we might set ourselves the task of finding the *biggest* painting by each of the master painters housed in the museum. "The Rembrandt is on the fourth floor" I say; "The Bruegel is in the basement" you respond, "But I don't think I've found the Bosch". In each case what we mean is 'the *biggest* Rembrandt', 'the *biggest* Bruegel', and 'the *biggest* Bosch'.

One way to bring this into stark relief is by considering again the ordering effects that Scontras et al. (2017, 2019) discuss that we considered above. They found that "the big blue box" was much more common and sounded more natural than "the blue big box". However, this too can

be overridden in certain discourse contexts. For instance, suppose that we are looking for all the big boxes. I then say:

(24) The *blue* big box is in the kitchen and the *red* big box is still upstairs.

We take it in this context, (24) sounds reasonable, especially with emphasis on the color terms. It might even sound more felicitous than using a sentence in which the ‘big’ precedes the color adjectives.

Our view is that certain cognitive heuristics are relevant to facilitating the very general human aims of prediction, explanation, and manipulation when carried out together. These explain the implicit default domain restriction possibilities that we find. When more specific aims are adopted prior to or in the course of conversation, other restrictions will be relevant given these more specific idiosyncratic aims. This explains the wider availability of implicit restrictions within particular conversational contexts.

4.1 Displacement and Non-Default Restrictions

To close this section, we consider a worry that might arise for our account. An interesting feature of linguistic communication is its capacity for displacement: language is “able to free itself from the ‘bonds’ of the immediate spatio-temporal surroundings of its production” (Auer 1988; see also Hockett 1960; Hockett & Altmann 1968; von Stechow & Heim 2011; Kratzer 2013). This manifests in many interesting ways. Much of the focus in philosophy of language and linguistics has been on modal displacement—the ability to talk about non-actual possibilities. We focus on something less removed from actuality here: the ability in conversation—whether face-to-face, via text message or email, or on the phone—for the context of that conversation (specifically, its indices and domain of objects) to pick out another spatio-temporal location.

Consider the following sentence, uttered discourse-initially, in a context where two philosophy professors who are eating lunch at a cafe have common knowledge about a friend Carl who works in the UCLA linguistics department:²⁴

(25) Carl managed to have his sensitive phone call overheard by everyone again.

²⁴ Thanks to an anonymous reviewer for suggesting this example.

Similarly, imagine a situation where the authors of this paper are out shopping for supplies for a party hosted by our friend, Dan. One of us gets a phone call from Dan, who informs us about the state of his kitchen after his 5-year old daughter's attempt at baking a cake for the party, and says:

(26) Dan's daughter was trying to make a cake, and apparently everything is covered in flour!

In each of these cases, there is a natural reading of the restricted quantifiers: something like 'everyone in the UCLA linguistics department'²⁵ or 'everything in Dan's kitchen'. At first blush it might even seem like these instances of displacement are default restrictions, and that our account – focused as it is on cognitive heuristics that predict default restrictions to the present physical surroundings of the conversational participants – cannot explain how these restrictions arise. Shouldn't our account predict that 'everyone' is read as 'everyone in this cafe and 'everything' as 'everything in this grocery store'? We think that it would be mistaken to jump to this conclusion, given that when we start talking about people like Carl and Dan, or what has happened 'again', this triggers a non-default context: we are talking about what is going on with Carl and Dan.

What is puzzling, however, is how it is that even when we displace a conversational context to, e.g., what is going on where Carl was, or where Dan is, we are able, without explicit contextual setup, to restrict the domain of 'everyone' or 'everything' to precisely those objects that are relevant for the story being told.

We think that this ability reflects the use of the cognitive heuristics and joint aims in conversation we have proposed, but in addition to them it reflects a more general capacity for perspective taking.²⁶ So, when one of us reports to the other that 'Everything is covered in flour' we are able to understand this as referring to every object of a certain sort in Dan's kitchen because we apply the perceptual availability and manipulability heuristics *from Dan's perspective*.

²⁵ (25) might involve an interesting ambiguity. There is a *de re* reading on which 'everyone' refers to exactly that group of people who overheard Carl's embarrassing conversation last time (e.g., the other members of the UCLA linguistics department. On another reading, it might refer to everyone in a specific location; on yet another reading, it might refer to everyone *wherever Carl is*. Some of these ambiguities might constitute a kind of felicitous underspecificity (cf: King 2022).

²⁶ This kind of perspective taking is sometimes called 'viewpoint construction', especially in the stylistics literature.

On one view, we take the perspective of another by serial adjustment from one's own point of view to theirs (Epley et al. 2004). Perspective taking is thus a relatively complicated cognitive process, and it is noted to arise relatively late in development (usually preschool age children are able to engage in some form of perspective taking; Newcombe & Huttenlocher 1992). The linguistic capacity for displacement is also noted to come ontogenetically late (Hockett 1960; Auer 1988). That it might depend on the similarly 'advanced' capacity for perspective taking is, we think, plausible, but requires further investigation.

5. Empirical Predictions

To summarize, our proposal is that the domain of quantification tracks the objects that are relevant to our joint purpose. On our view, default implicit domain restrictions are determined by cognitive heuristics (and perhaps other cognitive resources) given the minimal shared aims of explaining, predicting and manipulating together. Non-default restrictions are those that are relevant to a joint idiosyncratic goal—set out in discourse structure or prior to the start of a conversation. Our account is empirically testable. This is already made clear to a certain extent given the empirical work, especially that focused on reference resolution, we discussed in §3. Future work could investigate how perceptual availability, perceptual salience, manipulability, and objectivity affect implicit domain restriction, as well as how shared idiosyncratic goals affect domain restriction. We suggest some possible empirical tests below.

Availability

The perceptual availability heuristic predicts that without further information about shared knowledge or idiosyncratic goals speakers tend to speak about and be interpreted as speaking about perceptually available things, rather than things that they are not in perceptual contact with. To test this prediction, we might set up an experiment like the following: participants see a speaker in one environment (a room, forest, field, street scene...). The person has a text bubble saying something of the form "Every F is G" or "The F is G". Study participants are also able to see several other scenes of other locations (e.g., other rooms, streets). These scenes also include Fs. Participants are then asked to click on the objects that are relevant to whether what the speaker said is acceptable or true. For instance, a participant might see four rooms, one of

which the speaker is located in. The speaker says “every book is on the table”. There are books in all four of the rooms. The participant will click on all the books they take to be relevant to what the speaker in the vignette says. Our prediction is that without further information about the speaker’s goals or a larger discourse context, participants will be more likely to click on only the books in the room in which the speaker is located. If participants learn that the speaker has a goal to collect every book in the house or find as many books as possible, we predict they will click on every book the participant can see, even those that are not perceptually available to the speaker.

Another version of this experiment might build on work related to false belief tasks that have been used to test for theory of mind. For instance, a participant might be presented with two characters who enter a novel environment together, and who the study participant learns have only just met and do not yet share any specific goals about what to do in or learn about the environment. The participant again has a larger view of the environment (i.e., they can see areas that the two characters have not yet entered). One of the characters utters “Every F is G”. The portion of the environment the characters have explored contains some Fs, but there are also additional Fs the characters have not yet seen. The participant would then, as in the first study, be asked to click on all the objects relevant to the acceptability or truth of the speaker’s claim. Our prediction is that adult participants would click only on the objects that are perceptually available to the characters. Once additional goals are specified (e.g., if the characters are described as needing to find every F in the forest in order to win the game), additional Fs are then relevant to this joint aim and we predict that a non-default restriction connected to this more specific goal becomes operative. We hypothesize that in this condition participants would click on every F they can see, including those that the characters in the game have yet to encounter.

Perceptual salience

The perceptual salience heuristic predicts that we tend to be speaking about what is salient. That is, while relevance and salience come apart (as in example (13)) we default to treating salient items as relevant to discussion. Individual dispositions to find certain things salient vary (and can have striking normative consequences; see Munton 2023), but there are general patterns to our salience attributions roughly corresponding to what we are most likely to give attention to in a particular environment. This makes loud noises, bright / flashing lights, and

noticeable movement particularly easy targets for salience attributions. One way of testing the prediction that we tend to focus conversation on what is salient, then, would be to compare two virtual situations – one in which a number of shapes (hearts, clubs, diamonds, spades) are motionless, and another in which only some of those shapes (only the hearts) are moving around the space. Judgments of sentences like ‘They’re all hearts’ or ‘None of them are spades’ could then be assessed between the two scenarios, and for a number of other attention-grabbing features as well. What our account predicts is that in the scenario in which some of the hearts are moving around, these target sentences will read as more acceptable to participants than they will in a scenario in which all the shapes are motionless.

This proposed kind of experiment tests whether objects that are more salient are easier to quantify over. But other sorts of experimental questions related to our claims about salience might be probed as well, for instance whether the presence of salient objects creates defaults for restriction in a given environment.

Manipulability

The relationship between manipulability and domain restriction is, in one sense, quite easy to empirically establish, by attempting to verify whether there is any link between changes to what a subject is able to manipulate in their environment, on the one hand, and what sorts of quantificational statements about that environment they are disposed to accept.

For instance, manipulability could be empirically established by testing whether groups of individuals with the ability to manipulate some subset of objects in a (virtual) environment are more likely than others to assent to or assert various claims involving implicitly domain restricted quantifiers than are individuals presented with the same objects but a different manipulable subset. For instance, imagine an array of different shapes, including several diamonds, hearts, spades, and clubs. If the set of manipulable objects (perhaps these can be dragged around the screen, rotated, etc.) contains only one of the hearts, participants may be tested on whether claims made about ‘The heart’ seems felicitous as compared to a control group in which multiple hearts or no hearts are manipulable.

This can be tested in physical environments with tools as well. We might imagine an experiment in which an individual is sitting at a table in a room (with instructions not to get up), with a cup on

the table and another cup on a far bookshelf. Some participants are given a long stick that allows them to reach items across the room. We could attempt to establish whether individuals in the stickless group would respond with less confusion to a command like ‘Knock over the cup’ than would individuals with the ability to reach both cups with the stick.

Objectivity

In addition to the heuristics we’re relying on, we argued that objectivity helps to support our account. We suggested that default restrictions delivered by the cognitive heuristics to being here and now are objective and will be taken to be so given broad shared aims to understand, explain, predict, manipulate, and control the environment. Non-default restrictions are, we suggested, also more likely to be objective than subjective, although here they can relate to very specific idiosyncratic goals. To test for this, standard faultless disagreement tasks, like those discussed above could be used to test for location. For instance, participants in a study might be told that two speakers will issue utterances about things in a room or in an array of objects. Speaker A could then say something about an image being located in the image (e.g., “There is a dog”) or about the locations of objects in the image (e.g., “There is a dog next to a ball”) and Speaker B would then utter the negation of A’s statement (e.g., “There is no dog”). Participants would then be asked to judge whether both A and B could be right. We predict that participants would confidently judge that both speakers could not be correct. If confirmed, this would suggest that locations—like being in a room—are deemed to be objective.

We certainly do not take this to be an exhaustive list of ways our account is empirically testable. For instance, we have not discussed ways in which people represent the boundaries of places—especially outside of the built environment—within spatial cognition. We take this to be an area ripe for further exploration. Our discussion of empirical tests has focused on conversations involving adults. The account also opens up further avenues for developmental work examining how children implicitly restrict domains. This may be particularly interesting as we, like many others, take implicit domain restriction to relate to joint plans and goals—some that are general and hold broadly of creatures like us, some that are idiosyncratic and specific to conversational contexts—which might develop as children come to better understand conversational pragmatics and develop a more robust theory of mind.

6. Concluding Remarks and Future Directions

We began with a puzzle—certain implicit domain restriction possibilities seem to be available as defaults, while others are only available in certain conversational contexts. We argued that extant theories based on rational communicative practices, discourse structure, and speaker intentions do not explain this striking contrast. We argued that a cognitive explanation—relying on heuristics and features of our cognition—is needed to answer this initial question. This serves to motivate a partially nativist account of domain restriction. Of course, we do not take all domain restriction to be innate—we take relevance to specific joint plans to be central to many cases of domain restriction. What these plans are can vary widely in ways that affect domain restriction. Moreover, even the ways a restriction possibility is taken to be relevant to a goal might vary in important ways that eschew a nativist interpretation. We have argued, however, that features underlying human cognition play an important role in delimiting default domain restriction possibilities.

Our aim here was not to offer a semantics for quantifiers or definites that could account for the ways these can be implicitly restricted. Nevertheless, it is important to note that while we take it that domains can be restricted in ways that are default and non-default, this does not require a disjunctive first-order semantics. For instance, suppose one favored situation semantics as a way to account for domain restriction (Barwise and Etchemendy 1987; Kratzer 1989; Recanati 1996; Elbourne 2013; Abreu Zavaleta 2021). This sort of account would cohere nicely with our view that domains can be restricted in two sorts of ways.²⁷ One could adopt a single entry for universal quantifiers, say, but allow for situation variables to be fixed either through the cognitive heuristics we argued explain default restrictions or through discourse structure. Other semantic accounts are also consistent with the account we have offered here.

As we noted above, our account opens up empirical questions about our proposed cognitive heuristics, connections between spatial cognition and language, the relationship between

²⁷ Note that in situation semantics, situations are not taken to be identical to concrete spatio-temporal regions (Kratzer 1989). Rather, on the standard Kratzerian picture, situations involve more minimal information that includes what she calls “thin particulars”, entities with just some of their properties. These cannot be identical to spatio-temporal regions as these regions include “thick particulars”. An argument that spatio-temporal regions (and perhaps those with boundaries relying on particular features of the environment, like the walls of a room) are privileged as default domain restriction possibilities could be added as a supplement to situation semantics, but it does not simply fall out of a situation semantics account. We thank an anonymous referee for suggesting this connection.

reference resolution and domain restriction, and so on. It also suggests possible theoretical extensions. For instance, what is the relationship between default domain restriction possibilities and normative and social restrictions on domains? Consider, for instance, utterances of sentences like,

(27) Any bartender knows how to make a manhattan.

(28) Jocks sit at the back of the school bus.

(29) Everyone loves ice cream.

In these cases, the restrictions appear to be to normative domains: any *real* bartender; *normal* jocks; every *ordinary* person. These might be assumed to be objective or, in cases of conceptual engineering or metalinguistic negotiation, proposals for standards or definitions.

Finally consider an utterance of (30) made at a formal fundraising gala.

(30) Everyone had a great time.

Suppose that at the gala the invitees did indeed all enjoy themselves, but the waitstaff were treated so terribly that some cried and several walked out. We take it that if the addressee of (30) pointed out that the servers did not have a great time, the speaker might double down and say that they meant everyone "who matters" or "who was invited" had a great time. Such a response is clearly worrisome and relates to normative and moral considerations, but cognitive features like perceptual salience (who has the microphone, who is wearing a colorful gown rather than "blending in" in a uniform) might be relevant here as well. While normative cases like these have not been our focus, connecting this research to our more cognitive framework suggests another useful line of inquiry.

Thus we believe that the work done in this paper opens up important directions for further empirical and theoretical work at the interface between communication and cognition. But for now, our account explains our initial contrast case: why it is that certain implicit restrictions on quantifier domains are more natural and widely available than others.

Acknowledgements

We have benefited from conversations and feedback from many in the course of working on this paper. In particular we thank Sara Aronowitz, E.J. Green, Julian Jara-Ettinger, Edouard Machery, Craige Roberts, Greg Scontras, participants in a meeting of the Words Workshop, The Dublin Philosophy of Language Workshop, and the Postdoctoral Open Day at the University of Sheffield, and an editor and two anonymous referees for their comments and suggestions. We thank Luke Allen for drawing Figures 1 and 2.

Works Cited

Abreu Zavaleta, M. (2021). Communication and indifference. *Mind & Language*, 36(1), 81-107.

Angner, E. (2020). *A course in behavioral economics*. Bloomsbury Publishing.

Aronowitz, S., & Nadel, L. (2023). Space, and not Time, Provides the Basic Structure of Memory.

Auer, P. (1988). On deixis and displacement. *Folia Linguistica*.
<https://doi.org/10.1515/flin.1988.22.3-4.263>

Austin, J.L. (1962). *How to Do Things with Words*. Oxford University Press.

Bach, K. (1994). Conversational implicature. *Mind & Language*, 9(2), 124-162.

Bach, K. & Harnish, R.M. (1978). *Linguistic Communication and Speech Acts*.

Baron-Cohen, S., Campbell, R., Karmiloff-Smith, A., Grant, J., & Walker, J. (1995). Are children with autism blind to the mentalistic significance of the eyes?. *British Journal of Developmental Psychology*, 13(4), 379-398.

Bartlett, M., & Kazakov, D. (2005). The origins of syntax: from navigation to language. *Connection Science*, 17(3-4), 271-288.

Barwise, J. & Etchemendy, J. (1987). *The Liar*. Oxford University Press.

Barwise, J., & Perry, J. (1983). *Situation semantics*. MIT Press.

Beaver, D. I. & Clark, B.Z. (2008). *Sense and Sensitivity: How Focus Determines Meaning*. Wiley-Blackwell.

Bratman, M. E. (2014). *Shared Agency: A Planning Theory of Acting Together*. New York, NY: Oxford University Press.

Bratman, M. E. (1992). Shared Cooperative Activity. *Philosophical Review* 101 (2):327-341.

Buchanan, R. (2010). A puzzle about meaning and communication. *Noûs*, 340-371.

Buchanan, R. (2014). Reference, understanding, and communication. *Australasian Journal of Philosophy*, 92(1), 55-70.

Buchanan, R. (2018). Intention and the Basis of Meaning. *Ergo, an Open Access Journal of Philosophy*, 5.

Cappelen, H., & Lepore, E. (2002). Insensitive quantifiers. *Meaning and truth: Investigations in philosophical semantics*, 197-213.

Clarke, A. D., Elsner, M., & Rohde, H. (2013). Where's Wally: the influence of visual salience on referring expression generation. *Frontiers in psychology*, 4, 329.

Collins, C. (2018). Quantifier domain restriction as ellipsis. *Glossa: a journal of general linguistics*, 3(1).

Degen, J. (2023). The rational speech act framework. *Annual Review of Linguistics*, 9, 519-540.

Elbourne, P. (2013). *Definite descriptions*. Oxford: Oxford University Press.

Epley, N., Keysar, B., Van Boven, L., & Gilovich, T. (2004). Perspective taking as egocentric anchoring and adjustment. *Journal of personality and social psychology*, 87(3), 327.

von Fintel, K. (1994). *Restrictions on quantifier domains*. University of Massachusetts Amherst.

Firestone, C., & Scholl, B. J. (2016). Cognition does not affect perception: Evaluating the evidence for “top-down” effects. *Behavioral and brain sciences*, 39, e229.

Frank, M. C., & Goodman, N. D. (2012). Predicting pragmatic reasoning in language games. *Science*, 336(6084), 998-998.

Franke, M., Scontras, G., Simonič, M. (2019). Subjectivity-based adjective ordering maximizes communicative success. In *Proceedings of the 41st Annual Meeting of the Cognitive Science Society*, pp. 344–50. Seattle, WA: Cogn. Sci. Soc.

Frege, G. (1892). Über Sinn und Bedeutung. In *Zeitschrift für Philosophie und philosophische Kritik*, 100: 25–50; translated as ‘On Sense and Reference’ by M. Black in Geach and Black (eds. and trans.), 1980, 56–78.

Garoufi, Konstantina; Staudte, Maria; Koller, Alexander & Crocker, Matthew W. (2016). Exploiting Listener Gaze to Improve Situated Communication in Dynamic Virtual Environments. *Cognitive Science* 40 (7):1671-1703.

Gauker, C. (1997). Universal instantiation: A study of the role of context in logic. *Erkenntnis* 46 (2):185-214.

Gauker, C. (2003). *Words without Meaning*. MIT Press.

Gibson, J. J. (1977). *The theory of affordances*. Hilldale, USA, 1(2), 67-82.

Gigerenzer, G., & Gaissmaier, W. (2011). Heuristic decision making. *Annual review of psychology*, 62, 451-482.

Gilbert, M. (1989). *On Social Facts*. Routledge.

Gilbert, M. (2006). *A Theory of Political Obligation: Membership, Commitment, and the Bonds of Society*. Oxford: Oxford University Press.

Glenberg, A. M. (2010). Embodiment as a unifying perspective for psychology. *Wiley interdisciplinary reviews: Cognitive science*, 1(4), 586-596.

Grice, H. P. (1957). Meaning. *The philosophical review*, 66(3), 377-388.

Grice, H. P. (1969). Utterer's meaning and intention. *The philosophical review*, 78(2), 147-177.

Grice, H. P. (1975). Logic and conversation. In *Speech acts* (pp. 41-58). Brill.

Guerra, E., Marghetis, T., & Knoeferle, P. I. (2013). Spatial meanings for function words? The link between conjunctions and spatial representations. In M. Knauff, M. Pauen, N. Sebanz, & I. Wachsmuth (Eds.), *Proceedings of the Annual Meeting of the Cognitive Science Society* (pp. 2458-2463). Austin, TX: Cognitive Science Society.

Hanna, J. E., & Brennan, S. E. (2007). Speakers' eye gaze disambiguates referring expressions early during face-to-face conversation. *Journal of Memory and Language*, 57(4), 596-615.

Harris, D. (2014). *Speech Act Theoretic Semantics*. Doctoral Dissertation. CUNY Graduate Center.

Harris, D. (2020). We talk to people, not contexts. *Philosophical Studies* 177 (9): 2713-2733.

Hockett, C. F. (1960). The origin of speech. *Scientific American*, 203(3), 88-97.

Hockett, C.F. & Altman, S. (1968). A note on design features. In T.A. Sebeok (Ed.). *Animal communication: Techniques of study and results of research*. Bloomington: Indiana University Press. (pp. 61-72).

Hoscheidt, S.M., Nadel, L., Payne, J. & Ryan, L. (2010). Hippocampal activation during retrieval of spatial context from episodic and semantic memory, *Behavioural Brain Research*

212 (2): 121-132. <https://doi.org/10.1016/j.bbr.2010.04.010>.

Hurley, S. L. (1998). *Consciousness in action*. Harvard University Press.

Jackendoff, R.S. (1983). *Semantics and cognition*. Cambridge, MA: MIT Press.

Jara-Ettinger, J., & Rubio-Fernandez, P. (2022). The social basis of referential communication: Speakers construct physical reference based on listeners' expected visual search. *Psychological Review*, 129(6), 1394–1413.

Kahneman, D., & Tversky, A. (1979). On the interpretation of intuitive probability: A reply to Jonathan Cohen.

Kaufmann, M. (2011). *Interpreting imperatives* (Vol. 88). Springer Science & Business Media.

Keiser, J. (2022). *Non-Ideal Foundations of Language*. Routledge.

King, J.C. (2014). Speaker intentions in context. *Noûs*, 48(2), 219-237.

King, J. C. (2022). *Felicitous Underspecification: Contextually Sensitive Expressions Lacking Unique Semantic Values in Context*. Oxford University Press.

Kölbel, M. (2004). Faultless disagreement. In *Proceedings of the Aristotelian society* (Vol. 104, No. 1, pp. 53-73). Oxford, UK: Oxford University Press.

Kratzer, A. (1989). An investigation of the lumps of thought. *Linguistics and philosophy*, 607-653.

Kratzer, A. (2013, January). Modality for the 21st century. In *19th International Congress of Linguists* (pp. 181-201).

Kratzer, A. (2023). Situations in natural language semantics. *The Stanford Encyclopedia of Philosophy*, Edward N. Zalta & Uri Nodelman (eds.), URL = <https://plato.stanford.edu/archives/fall2023/entries/situations-semantics/>.

Kukla, Q. (2014). Performative Force, Convention, and Discursive Injustice. *Hypatia*. 29(2): 440–457.

Kukla, R., & Lance, M. (2009). *'Yo!' and 'Lo!': The pragmatic topography of the space of reasons*. Harvard University Press.

Langacker, R. W. (2008). Cognitive grammar. *Cognition and Pragmatics*, 77.

Langton, R. (1993). Speech acts and unspeakable acts. *Philosophy and Public Affairs* 22 (4):293-330.

Lepore, E., & Stone, M. (2014). *Imagination and convention: Distinguishing grammar and inference in language*. OUP Oxford.

Levinson, S. C. (2023). Gesture, spatial cognition and the evolution of language. *Philosophical Transactions of the Royal Society B*, 378(1875), 20210481.

Lombrozo, T., & Carey, S. (2006). Functional explanation and the function of explanation. *Cognition*, 99(2), 167-204.

Maguire EA, Gadian DG, Johnsrude IS, Good CD, Ashburner J, Frackowiak RS, Frith CD. (2000). Navigation-related structural change in the hippocampi of taxi drivers. *Proc. Natl Acad. Sci. USA* 97, 4398–4403. (doi:10.1073/pnas.070039597)

Mårtensson J, Eriksson J, Bodammer NC, Lindgren M, Johansson M, Nyberg L, Lövdén M. (2012). Growth of language-related brain areas after foreign language learning. *Neuroimage* 63, 240–244.

McGowan, M.K. (2019). *Just Words: On Speech and Hidden Harm*. Oxford University Press.

Munton, J. (2023). Prejudice as the misattribution of salience☆. *Analytic Philosophy*, 64(1), 1-19.

Neale, S. (1990). *Descriptions*. Cambridge, MA: MIT Press/Bradford Books.

Neale, S. (2000). On Being Explicit: Comments on Stanley and Szabo, and on Bach. *Mind & Language*, 15(2-3): 284–294.

Newcombe, N., & Huttenlocher, J. (1992). Children's early ability to solve perspective-taking problems. *Developmental psychology*, 28(4), 635.

O'Keefe, J. & Nadel, L. (1978). *The Hippocampus as a Cognitive Map*. Oxford University Press.

O'Keefe, J. (2003). *Vector grammar, places, and the functional role of the spatial prepositions in English*. Oxford University Press.

O'regan, J. K., & Noë, A. (2001). A sensorimotor account of vision and visual consciousness. *Behavioral and brain sciences*, 24(5), 939-973.

Portner, P. (2004). The semantics of imperatives within a theory of clause-types. In Watanabe, K. and Young, R., Eds, *Proceedings of SALT 14*. CLC Publications.

Qing, C., & Franke, M. (2015). Variations on a Bayesian theme: Comparing Bayesian models of referential reasoning. *Bayesian natural language semantics and pragmatics*, 201-220.

Recanati, F. (1989). The Pragmatics of What is Said. *Mind and Language* 4 (4):295-329.

Recanati, F. (1996). Domains of Discourse. *Linguistics and Philosophy* 19: 445–475.

Roberts, C. (1995). Domain restriction in dynamic semantics. In *Quantification in natural languages* (pp. 661-700). Dordrecht: Springer Netherlands.

Roberts, C. (1996 / 2012). Information structure: Towards an integrated formal theory of pragmatics. *Semantics and pragmatics*, 5, 6-1.

Roberts, C. (2010) *Retrievability and incomplete descriptions*. Ms., The Ohio State University. <http://www.ling.ohio-state.edu/~croberts/Retrievability.pdf>

Roberts, C. (2012). Information Structure: Afterword

Roberts, C. (2022). Coherence, Salience and Anaphora: The role of the QUD.

Rosch, E. (1975). Cognitive representations of semantic categories. *Journal of experimental psychology: General*, 104(3), 192.

Rubio-Fernandez, P. (2019). Overinformative speakers are cooperative: Revisiting the Gricean maxim of quantity. *Cognitive science*, 43 , e12797.

Russell, B. (1905). On Denoting. *Mind* 14 (56):479-493.

Sainsbury, R.M. (1979). *Russell*. London: Routledge & Kegan Paul.

Sellars, W. (1954). Some reflections on language games. *Philosophy of Science*, 21(3), 204-228.

Schwarz, F. (2019). Definites, domain restriction, and discourse structure in online processing. *Grammatical Approaches to Language Processing: Essays in Honor of Lyn Frazier*, 187-208.

Schiffer, S. (1972) *Meaning*. Oxford University Press.

Scontras, G., Degen, J., & Goodman, N. D. (2017). Subjectivity predicts adjective ordering preferences. *Open Mind*, 1(1), 53-66.

Scontras G, Degen J, Goodman ND. (2019). On the grammatical source of adjective ordering preferences. *Semant. Pragmat.* 12(7):1–21

Searle, J. (2010). *Making the Social World: The Structure of Human Civilization*. Oxford: Oxford University Press.

Sellars, W. (1954). Presupposing. *Philosophical Review*, 63, 197–215.

Simonič, M. (2018). Functional explanation of adjective ordering preferences using probabilistic programming. Unpublished master's thesis. University of Tübingen.

Sperber, D., & Wilson, D. (1986). *Relevance: Communication and cognition* (Vol. 142). Cambridge, MA: Harvard University Press.

Stalnaker, R. (1978) Assertion. In *Syntax and Semantics 9: Pragmatics*, edited by P. Cole, New York: Academic Press. 315–332.

Stalnaker, R. (1998) On the representation of context. *Journal of Logic, Language and Information* 7 (1):3-19.

Stalnaker, R. (2014). *Context*. OUP Oxford.

Stanley, J., & Gendler Szabó, Z. (2000). On quantifier domain restriction. *Mind & Language*, 15(2-3), 219-261.

Stojnić, U. (2021). *Context and Coherence: The Logic and Grammar of Prominence*. Oxford University Press.

Tomasello, M. (2022). *The Evolution of Agency: Behavioral Organization from Lizards to Humans*. MIT Press.

Tversky, A. & Kahnemann, D. (1974). "Judgment under Uncertainty: Heuristics and Biases", *Science*, New Series, 185, 4157, 1124-1131.

Vargha-Khadem, F., Gadian D.G., Watkins, K.E., Connelly, A., Van Paesschen, W., Mishkin, M. (1997). Differential effects of early hippocampal pathology on episodic and semantic memory. *Science*. 277(5324):376-80. doi: 10.1126/science.277.5324.376.

Wakslak, C. J. & Kim, B. K. (2015) Controllable objects seem closer. *Journal of Experimental Psychology: General* 144:522–27.

Wittgenstein, L. (1953). *Philosophical Investigations*. Blackwell, Oxford Publishing.

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