

JOINT ACTION AND COMMUNICATION, WITH A FOCUS ON DEMONSTRATIVES

by

Rory Harder

A thesis submitted in conformity with the requirements
for the degree of Doctor of Philosophy
Graduate Department of Philosophy
University of Toronto

© Copyright 2022 by Rory Harder

Abstract

Joint Action and Communication, With a Focus on Demonstratives

Rory Harder

Doctor of Philosophy

Graduate Department of Philosophy

University of Toronto

2022

This dissertation develops a view of human communication according to which it is the joint activity of making sense of one another, as fellow agents, in virtue of our actions. I consider the relationship between communication and joint attention in order to show how my view provides a superior successor to orthodox Griceanism. I also motivate my view by using it to explain how there can be successful communication, involving a demonstrative, despite the interlocutors not settling on one or another object as referent. Finally, I address a challenge, which arises for any Gricean view of communication, of explaining what is special about linguistic communication. My answer is based on a formal semantics for demonstratives and related expressions, and draws upon the tradition of dynamic semantics.

Acknowledgements

I would first like to thank Imogen Dickie, whose wisdom and rigour have been the guiding lights of my intellectual development—not only throughout writing this dissertation, but going back to my introduction to the philosophy of language in my undergrad. Thanks also to Nate Charlow for his indispensable role as co-supervisor—especially in helping with the formal semantics.

Many friends and colleagues at the University of Toronto helped immensely, and I would like to thank in particular Michael Szlachta, Lu-Vada Dunford, Rob Mason, Robbie Matyasi, Julia Lei, and Zach Weinstein. My time in the New York area was crucial for developing the ideas in this dissertation, and I would like to thank all those I met there for valuable discussions. Finally, I would like to thank those who helped with various forms of feedback: Gabe Dupre, Evan Westra, Brendan de Kenessey, John Campbell, Johannes Roessler, Hong Yu Wong, Simon Charlow, Chris Barker, Craige Roberts, Paul Pietroski, Dan Harris, David Boylan, Matt Mandelkern, Ernie Lepore, and Una Stojnić.

The lists of names above are in no particular order. All errors are my own. The work was supported by funding from SSHRC and OGS.

Contents

Introduction	1
1 Joint Attention and Communication	7
1.1 Direct Coordination and Perceptualism	10
1.2 The Communicative Suggestion and Griceanism	17
1.3 Communication and Social Motivation	25
1.3.1 The Necessity Worry for Griceanism	27
1.3.2 Direct Coordination and the Concept Acquisition Problem	33
1.4 Looking Ahead	38
2 Demonstratives and Joint Attention	40
2.1 Coordinated Attack and Common Knowledge	42
2.2 Transparency and Perceptual Demonstratives	48
2.3 Frege and Felicitous Underspecification	53
2.3.1 Challenging the Neo-Fregean Insight	54
2.3.2 Pathways of Thinking	58
2.4 Looking Ahead	62
3 Underspecification in Demonstrative Communication	65
3.1 Felicitous Underspecification and its Limits	68
3.2 Understanding and the Commonplace Effect	73
3.3 The Goal of Conversation	77
3.4 Looking Ahead	82
4 Salience and Deixis in the Anaphoricity of Demonstratives	85
4.1 Familiarity and Uniqueness	88
4.1.1 Presupposing Uniqueness	88

4.1.2	Introducing Familiarity	93
4.1.3	The Hybrid Approach	96
4.2	Saliency and Deixis	101
4.2.1	Core Proposal	101
4.2.2	Application	107
4.2.3	Pragmatics of the Referential System	111
4.3	Intensional Extensions	116
4.3.1	Deictic Features	118
4.3.2	Rigidity and Existentialism	124
4.4	Looking Ahead	130
.1	Formal System	133
.1.1	Syntax	133
.1.2	Semantics	134
	Conclusion	137
	Bibliography	146

Introduction

As noted by Aristotle (1984: 1253:a29–30), we are social animals. We do things together—go for walks, vote for leaders—in a distinctively human way. This dissertation is a modest attempt to shed some light on this aspect of human nature. I address two special and interrelated forms of our joint engagement—linguistic communication and joint attention—that have been studied by philosophers, psychologists, and linguists. I begin in chapter 1 with joint attention. As the dissertation develops, the focus turns to linguistic communication, and even more specifically on the meaning and use of demonstratives: words such as “this” and “that.” These words have been taken to play a special role in linguistic communication, in part because of an alleged connection to joint attention.

In accounting for the rationality of things we do—explaining our actions as sensible activities—it is tempting to treat us as ideally rational agents by appealing to complexes of sophisticated mental states that are neither accessible to our own introspection nor even empirically supported. In resistance of this temptation, there is an overarching challenge of capturing robust normativity without overintellectualization. I face the challenge head on by taking seriously both relevant philosophical and empirical work. On the side of philosophy, I take a background stance of non-naturalism and internalism about justification. According to the former, there are genuinely normative phenomena to be explained about thought and action; according to the latter, these phenomena are bound up with the subject’s point of view. Yet the normative features of thought and action are not independent of the empirical facts, for the account of how the things we do make sense requires appeal to how in fact we do them.

I look to Frege’s work for a clue as to what the relevant normative facts are, since that is where analytic philosophy began. Frege famously distinguished between sense and reference: a name such as “Aristotle” does not only point to Aristotle—its referent—but has also attached to it a way of thinking of its referent—its sense (Frege 1960a,b). One motivation for sense and its objectivity, which he fervently stressed, was sense’s role in communication. He

saw that for successful communication with a singular term it is not sufficient that the interlocutors merely think of the referent in any which respective way; rather, they must think of it in appropriately related ways. He concluded that in successful communication a sense must be shared: there may be ineradicable low-level differences between the perspectives of the interlocutors—imaginative differences, for instance, between which mental images each happens to associate with the word—yet in understanding each must, and somehow can, grasp the same objective sense.¹

Since we moderns are skeptical of sense and its objectivity, we may take Frege's insight about communication to simply raise the following question: *how can there be rationally irrelevant differences between two perspectives on the same object?* Suppose that we are inspecting an unhealthy house plant on the table in front of us. It is out in the open that we are both concerned with this particular plant; we communicate about it fluidly in various ways. But, given our physical nature, we must occupy different places relative to the plant. So there are necessarily differences between our perspectives on it. How is it possible for these differences not to matter for the sake of our rational engagement? Suppose, in contrast, that there are subtle and cleverly placed mirrors throughout the room, so we are looking in opposite directions yet happen to be seeing the same plant. In such a case, the differences in perspective are rationally relevant. So, in general, what is the basis of the distinction between rationally relevant and irrelevant differences between perspectives? This question is Frege's puzzle, in the broadest possible sense.

In chapter 1, I focus on the intersubjective version of the puzzle, illustrated by the cases just given, in which the perspectives are had by different individuals in the context of joint engagement.² In elaboration, note that, in the case of cleverly placed mirrors, it is possible for us to achieve rational engagement. We may bring our conceptual thought to bear: reflecting upon each other's experiences, we may come to know that we are in fact looking at the same plant. In such a case there is what we might call "indirect coordination." But when the differences between our perspectives are rationally irrelevant—as in the first, mirrorless case—the

¹The emphasis on shared sense in communication is from Frege 1960b; in the second footnote of Frege 1960a, he suggests another view of communication, which has come to be associated with Russell, that communication only requires that the interlocutors fix on the same object, in any which way. Despite the subtleties of Frege interpretation, the view that more is required is manifestly Fregean; see Evans (1982, 1985) and McDowell (1998a).

²For a seminal discussion of the intrapersonal case—for instance, an individual tracking an object over time—see Campbell (1987). He argues, following Strawson (1971b) and Evans (1982, 1985), that there must be direct intrapersonal coordination (the terminology of "coordination," introduced later in this paragraph, is from Fine (2007).) I discuss Campbell's intrapersonal argument in chapter 2, but my work in chapter 1 joint attention begins with an interpersonal analogue of that argument of Campbell's.

coordination is direct. In direct coordination, conceptual reflection is not required in order for there to be rational engagement; the individuals are rationally related to each other in an immediate way. But what could such direct coordination be? Must it ever hold? Perhaps the only way in which we rationally interact is mediated by conceptual reflection upon each other's minds.

I locate the basis of direct interpersonal coordination in joint attention: when two (or more) individuals attend together to some object. Over the past several decades joint attention has become an object of shared focus among psychologists and philosophers because of its Janus-faced nature. It has a rich epistemic component—"Everything is in the open, nothing is hidden" (Peacocke 2005: 298)—despite being a predominantly perceptual interaction. More empirically, infants begin engaging in it around their first birthday and it serves as a crucial milieu for the development of social cognition. Hence, it must be simple enough to allow more complex theory of mind abilities to develop on its basis, yet rich enough to foster their growth. Given the early onset of joint attention, a striking empirical result is that it is not until 3–4 years age that children demonstrate the ability to simultaneously compare two conflicting perspectives on the same object. Hence, if our primitive engagement in joint attention is rational, it involves direct coordination, since it does not involve reflection upon differing perspectives and establishing that they are of the same object.

Moreover, a recent trend in both the philosophical and psychological literature on joint attention involves developing the suggestion that joint attention holds partly in virtue of communication: it is because we share our thoughts or feelings about an object that our individual attention becomes joint. There might seem to be a tension between the communicative conception and the possibility of direct coordination, since communication might seem to be a relatively sophisticated type of interaction. But I show that the communicative suggestion in fact helps provide an account of direct coordination in basic cases of joint attention, since communication may be understood in a broad enough way so as to include the simple communication that is present there, as a matter of empirical fact. Crucially, such communication is driven by our deep-seated social motivation to make sense of one another, as fellow agents, in virtue of our actions—to achieve what I call "mutual intelligibility." Joint attention fulfills that deep-seated social need, so it has the positive normative status of practical justification. The immediate justification of direct coordination may then be grounded in that practical justification.³

³For more on the notion of practical justification, see Dickie (2015: §3) and references therein. In chapter 1, I provide a structurally similar account of direct interpersonal coordination to hers of the intrapersonal case.

Thus one upshot of chapter 1 consists in showing how the communicative suggestion—that joint attention holds partly in virtue of communication—helps explain the direct coordination present in basic cases of joint attention. But I also argue there that the communicative suggestion leads the way to an improved version of the orthodox Gricean account of communication. The Gricean orthodoxy holds that successful communication consists in the achievement of common knowledge, between the speaker and addressee, of the speaker's attempt to influence the addressee's mental life.⁴ But one crucial motivation for the communicative suggestion is that common knowledge between individuals of each other's individual attention, even if their attentional states are mutually responsive, does not suffice for joint attention: common knowledge between individuals can be attained in an impersonal, disengaged way. Thus, if communication between individuals is a form of joint activity that can bring about the jointness of joint attention, then the Gricean orthodoxy, in its appeal to common knowledge, does not capture the nature of communication. I propose instead a refined version of Griceanism, which holds that successful communication consists of the addressee recognizing the speaker's attempt to influence the addressee's mental life in virtue of the joint activity of achieving mutual intelligibility. Interlocutors' engagement in that joint activity need not involve their deep-seated social motivation—though, as I argue, such engagement does require the social motivation, in a subtle way.

I move on, in chapters 2 and 3, to addressing a provocative consequence that the communicative suggestion has concerning the use of demonstratives. A commonplace view is that some uses of demonstratives require joint attention for their understanding. But the communicative suggestion flips that order of explanation: when a demonstrative is used to speak about a perceptually available referent, it is successful demonstrative communication that helps bring about joint attention. The central theoretical motivation for the commonplace view, which I deny, arises from the Neo-Fregean tradition. One part of this motivation that I critically discuss in chapter 2 is based on a version of Frege's puzzle that arises for demonstrative communication. Consider the following well-known case, from Loar (1976). John and Mary are at home with the television on in the background. They are talking about a man they regularly see on the train during their morning commute. That same man happens to be on the news program playing on the television. Mary notices the man on the television, and utters,

- (1) That man [pointing at the TV] is well-dressed.

⁴This orthodoxy stems from Schiffer's (1972) version of Grice's ([1957] 1989) account of speaker meaning.

Enthralled in their previous conversation topic, John does not notice Mary's pointing gesture or what is happening on the television. He assumes she is continuing to talk about the man they see on their commute, and interprets Mary's utterance accordingly. Intuitively, John has not understood Mary, despite the fact that they are thinking of the same individual. The Neo-Fregean argument for the claim that demonstrative understanding sometimes requires joint attention thus runs as follows: the lack of understanding in so-called "Loar cases," such as the one just given, shows that demonstrative understanding requires that the interlocutors have justification for taking themselves to be thinking of the same object; but in some cases this justification must consist of direct coordination—i.e. it sometimes must not be based on conceptual reflection—and it is precisely joint attention to the referent that provides that direct coordination.

In response to this Neo-Fregean motivation for the claim that joint attention is sometimes required for demonstrative understanding, I reject the claim that demonstrative understanding requires that the interlocutors have justification for taking themselves to be thinking of the same object. One way I reject that claim, in chapter 2, is by providing an alternative account of what is going wrong in Loar cases, where that alternative account is provided by an account of demonstrative understanding based on my refined Griceanism from chapter 1. Another way I reject the Neo-Fregean claim is by considering a range of underappreciated cases in which demonstrative understanding holds without the interlocutors even thinking of any object as referent. Such cases go back to Dummett (1973: 74), but the phenomenon has been recently revived by King (2014b, 2017), who labels it "felicitous underspecification." I argue that understanding is possible in cases of felicitous underspecification, without the interlocutors having justification for taking themselves to be thinking of the same object. In chapter 3, I provide positive account of how there is understanding present in cases of felicitous underspecification, which requires drawing upon Gricean ideas in a different way than was done in the development of my refined Griceanism in earlier chapters.

An important challenge to my accounts of demonstrative understanding in chapters 2 and 3, which arises for any Gricean account of communication, is that of explaining what is distinctive about communication with language. The final chapter 4 addresses this challenge by providing a formal semantics for demonstratives and related context-sensitive referential expressions. My formal semantics draws upon dynamic semantics,⁵ and explains some striking contrasts between demonstratives and definite descriptions that competing accounts cannot.

⁵For seminal developments of dynamic semantics, see Heim (1982, 1983), Groenendijk & Stokhof (1991), and Stokhof et al. (1996). I draw heavily upon the more recent system of Bittner (2009, 2011).

I now conclude this introduction by explaining the upshot that my appeal to dynamic semantics has, in the context of this dissertation, by explaining how it is that a dynamic treatment of the semantics of demonstratives lines up with a Gricean account of understanding. From a dynamic semantics perspective, all context-sensitive referential expressions, in all of their uses, are ultimately anaphoric, which is to say they pick up a discourse referent from the linguistic context. Talk of discourse referents informally recapitulates that there is semantic machinery coordinating existentially quantified information at the discourse level. The technical details are not for here. What matters for present purposes is that embracing the dynamic perspective means treating the following two uses of demonstratives on a par.

- (2) a. A man walked in. That man was sad.
 b. That man [pointing to a man] was sad.

The anaphoric demonstrative in the second sentence of (2-a) helps elaborate the information about the discourse referent introduced by the indefinite in the first sentence, and the content built up by the discourse in (2-a) is roughly the existential claim that there was a man who walked in and was sad. So something similar should be said about the deictic demonstrative in (2-b): the demonstration is akin to an indefinite to which the deictic demonstrative is anaphoric. Hence, the content of (2-b) is similarly existentialist: there is some man being demonstrated who was sad. This existentialist conclusion about deictic demonstratives might seem jarring, but not if one accepts the Gricean view that fully understanding an utterance involves making sense of the speaker. A sentence containing a deictic demonstrative, in virtue of its semantics, merely sets up a constraint on what the speaker might do with it.⁶ The semantically generated constraint is given by the existentialist content, but what the speaker might be doing (among other things) is referring to some particular individual.

⁶Early versions of similar constraint semantics are in Strawson (1950), Grice ([1969] 1989), and Bach (1987). For more recent developments, see Schiffer (2003), Neale (2005), Pietroski (2006, 2018), Heck (2014), and Harris (2020). None of the authors cited here draw upon dynamic semantics in elaborating their constraint semantics.

Chapter 1

Joint Attention and Communication

Joint attention occurs when two or more individuals attend together to some object. I focus on the two person case, as is standard. It is an early form of human joint engagement, and has been studied as such by a tradition of developmental and comparative psychology with roots in the pioneering 1970s work of Jerome Bruner and colleagues (Scaife & Bruner 1975, Bruner 1975a,b, 1977, 1983). From the psychological point of view, it is a topic of interest because it provides a crucial milieu for cognitive development, especially that of social cognition; it is an interaction simple enough for infants to engage in, yet rich enough to foster cognitive growth.

Even more recently, philosophers have turned their attention to the phenomenon (two seminal works are Campbell 2002 and Peacocke 2005). For philosophers as well, theoretical interest in joint attention lies in its being at once sophisticated and simple, but in slightly different ways: the sophisticated element is its epistemological profile; the simple, its predominantly perceptual nature. The epistemic component of joint attention has been called its “openness”: the activity of the participants is mutually manifest or transparent—*out in the open*—between them. So, where psychological interest has been in its role in development, the philosophical puzzle has been to account for its openness as at least largely perceptual. But these foci are clearly related, for the broad upshot of both the philosophical and psychological perspectives is that engagement in joint attention provides us with a basic grip on other minds, where “basic” is meant here roughly in the sense that sophisticated conceptual resources are not involved.

Moreover, an exciting and very recent trend in both philosophical and psychological work on the topic has consisted of developing a suggestion that joint attention holds partially in virtue of communication (Carpenter & Liebal 2011, Eilan 2015, Siposova & Carpenter 2019,

León 2021). Let me intuitively motivate this suggestion, as the cited authors do, with some contrived cases, which will also serve to pre-theoretically introduce the phenomenon of our interest. After doing so, I explain what makes the communicative suggestion so interesting, especially in relation to joint attention's basicness.

Suppose you are walking in the park and see a cute dog. The owner stands a short distance away. At first, you and the owner attend to the dog individually without noticing each other. Here are two ways this interaction might unfold.

- i. You cast a furtive glance to the owner; you want to check if she is looking away so that you can get a chance to pet the dog without the trouble of asking for permission. She in turn notices your presence and subtly shifts her eyes towards you. Furthermore, suppose that each of you notice that the other is spying on the other. But, to avoid any awkward interaction, you covertly monitor her attention to you, and she does the same. Now, you are both aware of each other's attention to each other and the dog. And potentially—through increasingly sneaky looks—you are each aware of each other's awareness of that awareness, and so on. So here there may be arbitrarily high-level mutual awareness of your attention to each other and the dog; nonetheless, you are not attending *together* to the dog.
- ii. Upon seeing the dog and being overcome by its adorableness, you look to the owner in expression of your fondness for it. You smile and wink in that perfectly ambiguous way. Your eyes meet and perhaps you are happy to find her express a similar feeling, though with a slight tinge of weariness—presumably over the dog's unbounded energy. Or perhaps you are surprised to find her totally fed up with the dog's naughty behaviour. In any case, you look back at the dog in consideration of her reaction, expecting the interaction to continue.

Joint attention, along with its characteristic openness, is present in ii, but not i. And the contrast between these cases intuitively motivates the suggestion that communication plays a crucial role in sustaining joint attention.

Given its intuitive motivation, the communicative suggestion challenges a widespread approach in philosophy and beyond that characterizes public information in terms of common knowledge. Common knowledge holds when two individuals know that such-and-such, they each know that they each know that such-and-such, and so on (forever).¹ In case i, there may

¹The notion of common knowledge goes back to Lewis (1969) and Schiffer (1972). For prominent uses of it beyond philosophy, see Fagin et al. (1995) and Clark (1996).

be common knowledge between you and the dog owner about your attention to the other and the dog. Yet, again, there is no joint attention. Hence, such common knowledge is not sufficient for the kind of publicity that holds in joint attention. As shown by ii, it is really communication that brings about the openness of the interaction. Moreover, there seems to be no good reason to rule out, ahead of inquiry, the possibility of communication doing so without requiring common knowledge. Thus, the communicative suggestion promises a fresh account—*independent of common knowledge*—of least some forms of publicity.²

The communicative suggestion also illuminates a way of making progress on the challenge of accounting for joint attention's basicness. For—and this is the guiding insight of the chapter—the subtle nature of joint attention might be captured by a suitably subtle notion of communication. But there is a significant amount of ground clearing that must be done in order to develop the communicative suggestion in such a way as to help explain joint attention's basicness, especially since the common knowledge tradition has its talons in the orthodox Gricean account of communication.

So the chapter goes as follows. In §1.1, I refine the considerations mentioned above about joint attention's basicness, and proceed to lay out Campbell's (2002) perceptualist account of joint attention, according to which joint attention is a purely perceptual affair consisting of experiences and associated perceptual mechanisms. The perceptualist account, as I explain, is naturally motivated by considerations of joint attention's basicness. But, in §1.2, I present Eilan's (2015) version of the communicative suggestion, which is specifically an objection to Campbell's perceptualism. I then begin developing the communicative suggestion by looking to the Gricean tradition. I propose a version of Gricean communication that does not rely upon anything like common knowledge. The final §1.3, however, begins with a couple of issues for my modified Griceanism, and then draws together the considerations of basicness and the role of communication from the previous two sections in order to meet those issues. Crucially, I do not argue that my version of Griceanism meets those issues; instead, I take them to reveal that there is a related, but simpler form of communication, driven by a deep-seated social need, present in early episodes of joint attention. In conclusion, I discuss an upshot that the present approach to joint attention has for joint action and publicity in general.

²For recent work challenging the common knowledge orthodoxy, some of which I return to as the dissertation progresses, see Jankovic (2014), Tenenbaum (2015), Lederman (2017, 2018), and Harris (2019).

1.1 Direct Coordination and Perceptualism

The central line of argument for a perceptualist account of joint attention is that any alternative would appeal to beliefs or knowledge of the participants, yet such an appeal would overintellectualize the activity and not respect the basic way in which joint attention provides access to other minds. This core motivation is emphasized in the philosophical approaches of Campbell (2002, 2005, 2011, 2018) and Peacocke (2005),³ and here it is expressed by psychologist Tomasello (2009).

As the children play, they monitor the adult and her attention, and the adult monitors the child and the child's attention. No one is certain how best to characterize this potentially infinite recursion of monitoring, but it seems to be part of the infants' experience—at least in some nascent form—from before the first birthday.
(69)

Tomasello is suggesting here that joint attention involves a perceptual analogue of common knowledge, given presumably that bona fide common knowledge is too cognitively demanding for infants. Now, following a recent discussion by León (2021), I refine this line of argument by appeal to some specific results from the study of the development of social cognition, or THEORY OF MIND: our broad collection of cognitive skills for considering the mental lives of others.

³The view that Peacocke (2005) develops is ultimately non-perceptualist, but in a subtle way. In this footnote, I briefly present and criticize it, since Campbell's perceptualist view is the focus in this section. First, Peacocke appeals to a type of mental states of AWARENESS, which includes experiences, but also other occurrent, personal-level, nonconceptual, and non-inferential states. Then he characterizes the notion of MUTUAL OPEN-ENDED PERCEPTUAL AVAILABILITY, which holds between a state of affairs *s* and two individuals just in case "each [experiences] that the other [experiences] that *s* obtains; and if either is occurrently aware that the other is aware that he is aware... that *s* obtains, then the state of affairs of his being so occurrently aware is available to the other's occurrent awareness" (302). His overall view is then that two individuals *I* and *I'* are jointly attending to *o* just in case there is a state of awareness *A* that consists of (a) *I*'s and *I'*'s attention to *o*, (b) *I*'s and *I'*'s awareness that their attention to *o* has mutual open-ended perceptual availability, and (c) *I*'s and *I'*'s awareness that this state of awareness *A* exists. There is a lot of subtlety to this view, which I do not do justice to here. But here is nonetheless some quick criticism. The appeal to mental states of awareness that are not experiences is presumably intended to account for how joint attention involves mental states with more sophisticated content than experiences, but nonetheless more fundamental than beliefs. So Peacocke is trying to avoid an intellectualist account of joint attention. But, given that he does not give a positive characterization of non-doxastic and non-experiential states of awareness, the fundamentality of joint attention is not fully explained. Furthermore, the state of awareness *A* that characterizes joint attention is not a mental state of either participant, but rather must be a mental state of their plurality, or a mental state irreducibly *shared* by them. It is plausible that his appeal to this group mental state is meant to account for the mutuality involved in the openness of joint attention; however, he does not explain how there can be mental states of things other than individuals, which is *prima facie* implausible.

An important notion throughout this chapter is that of a BASIC episode of joint attention. Here are some empirical claims that partially characterize such episodes.

Primacy Research in the development of theory of mind has shown that children begin engaging in joint attention as early as 9–12 months of age, and it is partly in virtue of engaging in joint attention that they develop competencies in more complex cognitive skills, including some involved in theory of mind.

See Carpenter et al. (1998) and Moll & Meltzoff (2011) for overviews of empirical support, as well more detail, on Primacy.

One important aspect of theory of mind is so-called “perspective-taking.” This concerns primarily our abilities to engage with others’ visual perspectives (Flavell 1977, 1992). It is plausible that joint attention always involves some sort of perspective-taking, and it might be tempting to think that they way it does is by the joint attenders reflecting upon theirs and the other’s perspective and recognizing that each perspective is of the same object. The following empirical result, however, shows that basic episodes of joint attention do not involve that form of perspective-taking.

Confronting Perspectives Research in the development of theory of mind has shown that not until 4–5 years of age are children able to solve certain tasks that require them to CONFRONT PERSPECTIVES: simultaneously compare differing perspectives on a single object.

In support of this claim, a central study is that of Moll et al. (2013).⁴ They presented children with a blue object, and a yellow colour filter that, when looked through, makes the object look green. Children were given experience with the object and filter to discover its effect. Then, in a situation where there is an adult looking through the colour filter at the object, and where the child sees the object without a filter, the child is asked at once both how the object appears to them and how it appears to the adult. The child responds by either answering verbally or pointing to a colour patch. The experimenters found that 3 year olds respond that it appears blue to both them and the adult, but 4–5 year olds answer correctly that the object appears blue to them but green to the adult.

⁴For a recent overview of empirical support for both Primacy and Confronting Perspectives, see Tomasello (2018, 2019). As Tomasello outlines, other studies with results complementary to Moll et al.’s (2013) are Flavell et al. 1981, Doherty & Perner 1998 and Rakoczy et al. 2015.

It follows from Primacy and Confronting Perspectives that children engage in joint attention before they are able to confront perspectives. This empirical conclusion suggests that it is partly in virtue of their engagement in basic episodes of joint attention that they develop such sophisticated perspective-taking abilities. So if joint attention always involves some form of perspective-taking, then it is reasonable to take a step further and claim that basic episodes of joint attention provide a primitive way of *sharing* a perspective with another that precedes and aids in the development of the ability to explicitly consider another's perspective and compare it with one's own. This thought—that in basic episodes of joint attention there is a primitive sharing of perspectives—has informed recent psychological work on the topic (Moll & Meltzoff 2011, Tomasello 2018, 2019). Here, however, I wish to draw out a philosophical upshot concerning justification, which (as we will see) motivates the perceptualist account of joint attention.

The distinction between sharing and confronting perspectives brings to mind Frege's puzzle and his resulting distinction between names that share a sense and ones that differ in sense yet are nonetheless co-referential (Frege 1960a). Upon understanding the sentence "Cicero is Cicero," it strikes one as trivially true because there is no sensible question of whether the two occurrences of "Cicero" stand for the same individual. According to Frege, this is so because the name "Cicero" has a single sense (which all of its occurrences share). But one may rationally wonder whether "Cicero is Tully" is true, since "Cicero" and "Tully" differ in sense, despite the fact that they pick out the same object.

Instead of appeal to sense, a simpler and more modern way of addressing the phenomenon that Frege hit upon is in terms of coordination. Schematically, Frege's puzzle was originally about how sentences of the form $\ulcorner a = a \urcorner$ and $\ulcorner a = b \urcorner$ could differ in cognitive value, where $\ulcorner a = b \urcorner$ is true. This puzzle can be generalized to be about how $\ulcorner \phi[a] \urcorner$ and $\ulcorner \phi[b] \urcorner$ can differ in cognitive value when $\ulcorner a = b \urcorner$ is true. For instance, someone might understand both "Cicero was a roman orator" and "Tully was a roman orator," yet rationally accept one and not the other. As already glossed, in order to explain such differences in cognitive value, Frege posited the notion of sense. But what is sense? Talk of coordination allows one to capture rational relations between representations without reifying senses (Fine 2007). Instead of two expressions sharing a sense, one can say they are directly coordinated. Two co-referential singular terms are **DIRECTLY COORDINATED** when understanding both guarantees justification for taking them to be the co-referential; two co-referential singular terms are **INDIRECTLY COORDINATED** when understanding both does not guarantee justification for taking them to be co-referential.

Setting aside relations between linguistic expressions, more relevant to present purposes are phenomena amenable to coordination at the level of thought: the mental states of a single subject. An old insight is that there is direct coordination between beliefs formed when attending to an object over time (Evans 1982, 1985, Campbell 1987). Suppose that you are watching your cat slowly walk around the room: you first notice how it raises its shoulders, as if ready to pounce on an imaginary target; a moment later you see it sniffing one of your plants, and hope it does not poison itself again. In a standard episode of solo attention such as this, where different properties of the object are noticed over time, there is immediate justification for grouping together all of the resulting beliefs as beliefs about the same object; within such a procession of experiences, one does not have to reflect and determine on that basis that an object seen at one moment is the same as one seen a moment later. You do not have to think, for instance, *this thing approaching the plant is the same thing as was prowling a moment ago, since I was watching it the entire time as it followed a natural trajectory for an animate object, and I may assume that there is no evil demon that somehow swapped in a similar looking cat along its path*. Some have suggested that the immediate justification one has holds simply in virtue of the phenomenal character of the experiences themselves (Campbell 2002, Smithies 2011, 2019). The PHENOMENAL CHARACTER of a mental state is what it is like for the subject to undergo it. Furthermore, perceptual phenomenology is TRANSPARENT: in the case of an experience, what it is like to have it is how things appear to the subject of the experience (Strawson 1988). The suggestion here is thus that in sequences of experiences, modulated by attention, it is simply how things appear that there is an external world of enduring objects.

But the notions of coordination most relevant to present purposes are not intrapersonal, but interpersonal: they hold between one subject's and another's experience, in the context of a joint activity, when the subject has justification for taking hers and the other's experience to be of the same object. DIRECT INTERPERSONAL COORDINATION in the case of joint attention involves a participant having immediate justification for taking hers and another's experience to be of the same object. As above, justification is IMMEDIATE when it does not rest upon conceptual reflection. So there is INDIRECT INTERPERSONAL COORDINATION when the justification rests upon conceptual reflection. These interpersonal notions are the most relevant to what follows, yet the discussion of the intrapersonal notions in the paragraph above was not merely an aside: Campbell's perceptualist account of joint attention, as we will see, draws a tight analogy between the two (and I return to critically discussing this analogy in the final section, when my own view is on the table).

With all these preliminaries in place, we may establish a constraint on basic episodes of

joint attention.

Premise 1 Any episode of joint attention either involves direct or indirect interpersonal coordination between experiences.

Premise 2 If an episode of joint attention involves indirect coordination, then it involves deployment of the ability to confront perspectives.

Premise 3 In basic episodes, joint attention does not involve deployment of the ability to confront perspectives.

Conclusion In basic episodes, joint attention involves direct interpersonal coordination between experiences. [From P₁–P₃]

I justify the premises in turn. Premise 1 reflects an assumption that our engagement in joint attention involves a normative component. This assumption seems plausible if one considers what occurs in ii from the introduction, where there clearly is joint attention. However exactly the interaction goes, what you and the dog owner are doing—looking back and forth between each other and the dog, commenting on it in some way—makes it apparent that you are focusing together on the same dog. So I assume in general that in any episode of joint attention there is justification available to the subjects for taking their perceptual perspectives on the relevant object to be co-referential. And if this justification is not direct, then it must be indirect. Premise 2 holds because, in order for an individual to have justification for taking theirs and another's experiences to be co-referential, which is based on conceptual reflection upon two experiences, the individual must deploy the ability to simultaneously compare two perceptual perspectives. Premise 3 is supported by the empirical claims in *Primacy and Confronting Perspectives*: children are able to engage in joint attention before they have the ability to confront perspectives, and the development of the ability to confront perspectives unfolds partly in virtue of engagement in joint attention. So in these most basic cases of joint attention the ability to confront perspectives is not involved, and I label this conclusion the **DIRECT COORDINATION CONSTRAINT** on basic episodes of joint attention.

The direct coordination constraint refines talk of joint attention's basicness, and thus also refines the central motivation for perceptualism about joint attention. For, assuming that experiences are personal-level but nonconceptual mental states,⁵ we can see our way to an

⁵The assumption that experiences are nonconceptual is controversial. For early versions of the view, see Dretske (1981) and Evans (1982). More thorough motivation for and clarifications of that view are given by Peacocke (1992) and Heck (2007). A prominent defense of the view that experiences are conceptual is by McDowell (1994), but see Heck (2000) for a response.

account on which a proper configuration of experiences grounds the relevant epistemic facts of direct coordination.⁶ But, before presenting Campbell's perceptualism, let me pause to explain the background distinction between conceptual and nonconceptual mental states, as well as the related distinction between the personal- and subpersonal-level.

Beliefs and intentions are the paradigm conceptual mental states. For instance, I must have the concept of a frittata in order to believe that frittatas are tasty, or to intend to make a frittata for dinner. Experiences, I assume, are nonconceptual mental states. For instance, if you set a frittata down in front of me, I do not need the concept of one in order to see it. Furthermore, pre-experiential perceptual processing involves nonconceptual states: edge-detection involves calculations employing multi-variable calculus (Marr 1982), which my visual system performs with ease, despite that I struggled in first-year calculus. In general, following Peacocke (1992), a *CONCEPTUAL* mental state is one that, in order to be in it, the subject must possess the concepts required for the specification of its content. A *NON-CONCEPTUAL* mental state is one that is not conceptual.

Plausibly, any conceptual mental state is attributable to the person. For instance, it is I believing this frittata in front of me is tasty (so this clearly conceptual state is also personal-level); or, considering the contrapositive, it is not I calculating where the edges of it are on the basis of the retinal arrays of light intensities, since I do not know how to do those calculations (so this clearly subpersonal state is also nonconceptual). Yet there are nonconceptual mental states that are attributable to the person, the paradigm of which are experiences. For instance, it is I seeing the frittata. In general, a *PERSONAL-LEVEL* mental is one that is attributable to the person who is in it; whereas, a *SUBPERSONAL* one is instead attributable to one of their cognitive or physiological subsystems (Dennett 1969). Note that a crucial point later on in this chapter is that there are nonconceptual personal-level mental states other than experiences.

Returning now to the main thread of the discussion, Campbell's perceptualist proposal is that joint attention is a three-place experiential relation between two individuals, who are present to each other as co-attenders, and an object. The relation is experiential in that, when two individuals and an object stand in it, the individuals are in experiential states. Furthermore, the object and the other's presence as co-attender are parts of each individual's experience. These experiential states are sustained by subpersonal perceptual mechanisms that

⁶One immediate worry is that the direct coordination constraint rules out a perceptualist account of joint attention, since the constraint entails that in a basic episode participants must have the belief that theirs and the other's experiences are of the same object. But that entailment does not hold, since the relevant justification is propositional, and not doxastic. The justificandum of direct coordination is the content of a belief, but having justification for a belief-content does not require that one in fact form the corresponding belief.

track the object and the other's gaze; however, while the holding of the relation is sustained by subpersonal perceptual mechanisms—which may be deployed whether or not there is an object present or another individual present as co-attender—it is a relation that only holds if an object and co-attending individuals are in fact present.

This account faces an immediate problem (Peacocke 2005, Eilan 2015, Battich & Geurts forthcoming). The other must enter one's experience as a *co-attender*, and not merely as, say, a parallel attender. But Campbell does not explain what it means for another to be present in one's experience in the required way.

A natural way of elaborating Campbell's view in response to this initial problem, however, was anticipated above by the discussion of intrapersonal coordination. On Campbell's account of an individual's lone experience of an object, the role of subpersonal perceptual information-processing is in maintaining the subject's view of the object. A common way of understanding the role of this visual information-processing is in line with representationalism (see e.g. Harman 1990), on which the processing is glossed as the construction of a representation, which then exhausts the phenomenal character of the subject's experience. In contrast to the representationalist metaphor, Campbell suggests that the processing's role is simply that of making objects and their properties visible to the subject. Imagine looking through a pane of glass. The clearer the glass, the clearer the view; and there is no representation on the glass of what is seen through it. So the alternative metaphor is that the role of subpersonal perceptual processing is simply in maintaining the pane's transparency.⁷ Thus, by analogy with the lone case, it may be that in joint attention the subpersonal mechanisms of gaze-following make the other attender's activity transparent. Hence, the natural elaboration is as follows: when there is an object and another individual present, the deployment of the subpersonal perceptual mechanisms of gaze-following are sufficient for each individual's experience to contain the other as co-attender.

So, on the natural elaboration, there is a tight connection between solo and joint attention. In the solo case, subpersonal perceptual mechanisms, such as those bound up with edge-detection, play their role in maintaining a subject's view of an object in the world. In the joint case, additional mechanisms swing into play under the hood to track another's attentional state. These additional mechanisms may be in a computational sense more sophisticated than those in the solo case, yet they are crucially of the same kind as pre-experiential subpersonal

⁷In the terms introduced by Chomsky (1995), this way of drawing the contrast between representationalism and direct realism presupposes an "ersatz" view of some mental representation, since it appeals to metaphorical glosses on the role of subpersonal processing. Compare Egan (2014, forthcoming). Though, following Campbell, I am careful to restrict ersatzism to subpersonal mental states.

processing. It is precisely for this reason that Campbell's account is perceptualist. But, it is also for this reason that the account is open to another, more prescient, objection. I begin the next section by presenting that objection, which leads the way to the communicative conception of joint attention.

1.2 The Communicative Suggestion and Griceanism

Eilan (2015) argues that even Campbell's elaborated proposal fails to distinguish another's presence as a co-attender, as in truly joint attention, from presence as a mutual covert attender. She supplies a pair of cases, one of which is a counterexample to the elaborated proposal.⁸

Suppose you are at a mandatory meeting in which a university administrator is droning on about some new policy. Now contrast the following two continuations. First, you look up and make eye contact with your colleague across the table. You then look back together at the presenter, then look back at each other, exchanging looks of boredom. Second,

you become aware, out of the corner of your eye, that a colleague is watching you. As you become aware of his observation of you, you start employing something psychologists call 'covert attention' with respect to him. And as you begin to do this, he, in turn, becomes aware that you are attending to him, and thus you continue for a while, each of you dividing your attention between the speaker and each other. (Eilan 2015: 6–7)

It is plausible that in both the first case, where there is joint attention, and in the second, where there is not, the same subpersonal mechanisms of gaze-following swing into play. As emphasized, these mechanisms are of the same kind as those undergirding lone attention. Hence, it seems plausible that they could be deployed in both types of case. But then the natural way of extending Campbell's account does not provide a sufficient account of joint attention. The account predicts that another may be present as a co-attender without there being joint attention.⁹

⁸The cases i and ii from the introduction are similarly structured, and are partially based on her discussion. And, as cited in the introduction, Carpenter and colleagues were the first to provide contrasts of this kind, though they do not specifically target Campbell's perceptualism.

⁹Seemann (2010) elaborates Campbell's perceptualist proposal in order to meet the initial objection of underexplanatoriness, and the elaboration might also seem to be able to meet Eilan's more advanced objection. Seemann, drawing on the psychological work of Hobson (2002, 2005), holds that the sharing of feelings about the object is part of joint attention (some other discussions that draw on this material are in Roessler (2005)

Intuitively, what is missing in Eilan's board meeting case 2, and present in 1, is communication. So let me now present a general framework that captures the communicative conception of joint attention hereby motivated. After discussing the framework, I begin the task—which is the concern of the rest of the chapter—of filling out the framework.

So here is my attempt at systematizing the general conception of joint attention that is motivated by Eilan's objection to Campbell's perceptualism.

The Communicative Conception of Joint Attention Individuals I and I' are jointly attending to an object o if and only if they are undergoing respective experiences E_I and $E_{I'}$ that are sustained by:

1. attention to the object o ,
2. monitoring of the other's attention, and
3. communication between I and I' about o (*the communicative condition*).

The first two conditions follow Campbell in capturing the perceptual nature of joint attention. When individuals are jointly attending to some object, their experiences are sensitive to the object and each other's experiences. Both of these conditions can hold in virtue of subpersonal perceptual processing: our visual systems can track an object and another's gaze in a fast and automatic way. In an individual's lone attention, subpersonal perceptual processes play a causal information-processing role, and in virtue of doing so sustain a view of the external world for the subject via experience's phenomenal character. Analogously, as captured in conditions 1 and 2, in an episode of joint attention there are perceptual mechanisms of attention and gaze-following that help generate the phenomenal character of each participant's experience. But, crucially, the final condition represents that communication also plays a role in sustaining the experiences.

and Eilan (2005)). Crucially, however, the sharing is automatic and perceptual: one perceives another's bodily expression of a feeling towards the object, which is then automatically replicated in oneself. So the sharing of feeling to which Seemann appeals is not a form of communication, since communication is an action. As elaborated much later in this chapter, an action is something we do—not something that happens to us—and hence involves personal-level motivational states; whereas, Seemann (and Hobson) seem to be discussing a subpersonal and automatic process. Perhaps appealing to the mechanisms of the automatic echoing of affect makes it plausible that a different kind of subpersonal perceptual processing is being deployed between cases like board meeting 1 and 2. However, it seems plausible that the automatic sharing of affect could occur in cases like board meeting 2, where there is merely symmetrical covert attention: you may, for instance, feel the automatic pull of the other's strange excitement about the presentation, despite the fact that you have only noticed it in a covert manner.

When stated at this general level, the communicative conception does not necessitate the deep challenges to the common knowledge tradition mentioned in the introduction. Again, those challenges were that communication can bring about a kind of publicity independent of common knowledge and that capturing how communication may do so requires revising the traditional Gricean view of communication. In order to substantiate those upshots, the communicative component of the analysis must be elaborated. But let me first address some initial reservations one might have about the general framework.

According to some, certain cases of demonstrative communication hold partly in virtue of joint attention (Campbell 2002, Dickie & Rattan 2010, Seemann 2019).¹⁰ These cases are part of ordinary conversation and such that the demonstrative—e.g. “this” or “that”—is used to pick out an object perceptually available to the interlocutors. You point, for instance, at a parked Ferrari and declare “That car is my favourite”; intuitively, your friend—though perhaps annoyed by your obsession with sports cars—understands your uttered demonstrative partly in virtue of you both jointly attending to the car. But the communicative conception flips this order of explanation. Suppose you are standing beside that same friend and looking out a window. A loud crash reverberates from the street below. Both of you may look and see the passing truck responsible, and even notice that the other is doing so. Intuitively, however, there is no joint attention until there is communication between the two of you: perhaps a vocal expression of surprise, or simple raised eyebrows.

I return in the next chapter to a more thorough discussion of the relative priority of joint attention and demonstrative communication. The present point is simply that the communicative conception’s verdict has intuitive support. In addition, the intuitive considerations at the end of the previous paragraph show that the communicative conception is not in tension with the distinction between “top-down” and “bottom-up” joint attention (Carpenter & Liebal 2011: 170–171). In a top-down case, one participant draws the other’s attention to the object, and such cases are naturally covered by the communicative conception of joint attention. In a bottom-up cases, however, an object makes itself salient by, for instance, making a loud noise. Such cases might seem to pose a problem for the communicative conception; however, as illustrated above, despite the fact that an object might draw attention to itself,

¹⁰Relatedly, there have been a number of recent developmental and evolutionary considerations given in support of there being a tight connection between demonstratives and joint attention (Diessel 1999, 2006, O’Madagain & Tomasello 2019, Diessel & Coventry 2020, Rubio-Fernandez forthcoming). So, following this work, it may be thought that the evolutionary function of demonstratives is to initiate joint attention, and that the success of that function is crucial for early cognitive development. Yet those facts are compatible with the current proposal that, whenever a used demonstrative leads to joint attention, the joint attention holds partly in virtue of the demonstrative understanding.

there is intuitively no joint attention to that object until the participants communicate about it in some manner.

I turn now to unpacking the communicative condition, which places an important constraint on theories of communication. Recall the contrasts given already as motivation for it: the one between i and ii in the introduction and the one in Eilan's board meeting. Each of these contrasts includes a case where the perceptual components 1 and 2 of the communicative conception are satisfied, yet there is no joint attention. Hence, if it is the addition of communication that brings about joint attention, then successful communication must itself be sufficient to help bring about the jointness of an activity. A simple way in which communication might be sufficient in this regard is if communication is itself a joint activity. I assume that this simple way in which communication brings about jointness is correct. Hence, the constraint that the communicative conception of joint attention places on theories of communication is that communication must be represented as a joint activity. After all, the claim that communication is something we do together has intuitive support independently of the present considerations about joint attention; the communicative conception of joint attention simply brings into focus the joint nature of communication.

Let me begin the search for an account of communication that respects its status as a joint activity by considering the general nature of joint activity. According to the influential view of joint action of Bratman (1993, 2014), common knowledge of interlocked intentions is sufficient. On this view, for instance, we are walking together if the following holds: we both intend that we walk and that our resulting sub-intentions—to take one step, then to take a second—should influence one another, and all of this is common knowledge. But here is a case from Tenenbaum (2015) that forcefully illustrates how Bratman's account falters even for the prosaic case of walking together.

Let us suppose Larry has an inordinate amount of money coming to him. But the bank will give Larry the money only if he shows up at the bank with Mary so that a bank teller can confirm that Mary is alive. Unfortunately, as Larry knows, Mary would prefer him to be penniless, so she's not willing to go to the bank with Larry simply to help him. However, as luck would have it, Larry finds out that Mary is in the perfectly symmetrical predicament. As one would expect, Larry feels the same spitefulness toward Mary. Their mutual hatred runs deep but not as deep as their self-love; each would rather have the money rather than let the other suffer in poverty. Unfortunately, more powerful than their mutual hatred or their inflated self-love is their sense of dignity. Neither would stoop to ask the

other a favour or propose a truce or an agreement. They all know all this and are thus incapable of retrieving their respective fortunes. One day Larry is walking and he sees Mary; he immediately realizes that if he walks towards the bank, she'll follow him there in the hopes she'll get the money. Mary is hit by the same thought. They immediately notice that both would have had the same thought, and that they both realize that the other would have had noticed that both had the same thought, etc. They walk towards the bank, each carefully monitoring that the other is going. (3385)

Here each of Larry and Mary intends that *we walk to the bank*, and that their steps influence one another's. Furthermore, there is common knowledge of this whole affair. Yet there is intuitively no joint action: they are not walking to the bank *together*.¹¹

Tenenbaum suggests that what is missing is that Larry and Mary are not intending to do something that is essentially joint; in particular, they are not intending to walk to the bank *together*. One might have an at least prima facie worry that this suggestion is circular—as Tenenbaum does, following Bratman (2014)—for the holding of a joint activity is to be explained in part in virtue of an intention with a content whose specification involves appeal to a concept of that joint activity. But there is nothing in principle incoherent or underexplanatory about this form of circularity: as Fine (2012) observes in a related discussion, perhaps all there is to being cool is to be taken to be so. So I wish to combine Tenenbaum's suggestion with some central aspects of the Gricean approach to communication, in order to give an account of communication that secures communication's status as a joint activity. I return to the circularity worry after refining a Gricean view of communication and contrasting it with the standard version.

For Grice ([1957] 1989), the speaker's part of an episode of communication includes an intention to bring about a certain change in the addressee's mental life. But this intention—called the *INFORMATIVE* intention—is not the whole story: a parent might physically discipline their child and thereby make the child believe that they acted wrongly. But that morally dubious use of force is not communication. In addition, the speaker must have what is called a *COMMUNICATIVE* intention: as a first gloss, the speaker must intend that her informative intention be recognized by the addressee and for that recognition to serve as at least a part

¹¹A related point has been emphasized by Gilbert (1990, 2013), though she focuses on the claim that common knowledge of interlocked individual intentions is not sufficient for the socio-normative commitments intuitively present in joint action. Following Tenenbaum, I do not take the upshot of the case given here to be primarily about joint commitments, but I return to that point in the conclusion.

of the reason for the addressee's bringing about of the mental change. But this first gloss is not the whole story: a neuroscientist might have a machine that can manipulate brain waves via electrical shocks in order to bring about certain mental states. The neuroscientist might have an informative intention and use this device to bring about others' recognition of the intention. But in such a case there is intuitively no communication, even if the neuroscientist intends that the unfortunate "addressee" bring about the intended mental change with the induced recognition as a part of the reason for doing so.¹²

One might (correctly) suspect that the notion of recognition should rule out the neuroscientist case: the speaker must intend that the addressee *recognize* their intention. But an account is needed of what that recognition involves. I propose to use the notion of acting together, taken from Tenenbaum, as the basis for the required account: in an episode of joint action, there is a joint goal, which is an essentially joint activity represented in the intentions of the participants. I claim that this joint goal is that of mutual intelligibility: making sense of one another as fellow agents. So the speaker's communicative intention involves the aim that she and the addressee work together in making sense of one another's action, and that successful engagement in this activity generates the addressee's recognition of the informative intention.

Here is the proposal stated in systematic fashion.

Refined Gricean Communication There is communication_{ig} (in the joint-activity-based Gricean sense) between speaker *S* and addressee *A* just in case *S* makes an utterance where:

1. *S* intends that a particular reaction *r* is brought about in *A* (*the informative intention*),
2. *S* intends to engage with *A* **in the joint activity of making sense of one another's actions**, where successful engagement in this activity generates *A*'s recognition of the informative intention (*the primary communicative intention*)
3. *S* intends that this recognition serve as at least a part of *A*'s reason for producing that effect (*the secondary communicative intention*), and
4. the primary communicative intention is fulfilled (*the uptake condition*).

¹²This type of case, which appeals to a neuroscientist that can bring about mental states with shocks to the brain, is from Schiffer (1972).

Crucially, as indicated by the subscript of *jk* on “communication” in the analysis above, I have only provided necessary and sufficient conditions for a specific type of communication: the joint-activity-based Gricean one. I am not claiming that there are no other forms of communication. Nonetheless, I do think that the constraint that communication is a joint activity, which was brought to light by the communicative conception of joint attention, shows that the standard version of Gricean communication is incorrect. For that reason, which I am about to elaborate, I take my proposal to be a superior successor to standard Griceanism.

In order to compare my proposal with standard Griceanism, let me first say a few words about the concept of the joint activity of making sense of one another, which I have claimed is crucial for specifying the content of the speaker’s primary communicative intention. The paradigm instance of making sense of another’s actions is the attribution of conceptual mental states such as beliefs and intentions. And that paradigm is exactly the intended result of the speaker’s primary communicative intention, which is that the addressee attribute an intention to her. But simpler forms of making sense of one another should be possible: one might use whatever cognitive resources one has in order to see how another is “like me” (Tomasello 1999: 70). And these simpler forms are relevant to how that overall result—the attribution of the speaker’s intention—is brought about (in addition to the potential attribution of other intentions or beliefs in bringing about that result). The addressee should, according to the speaker’s primary communicative intention, work together with the speaker in order to discover the informative intention, in such a way that indicates to the speaker that the addressee is a fellow agent. I return to this particular joint activity in the next section, as I will argue that the concept of it is gained from a deep-seated social motivation.

On the standard Gricean picture, the bold part of the primary communicative intention above is instead filled out by appeal to common knowledge: the speaker intends that her and the addressee come to have common knowledge of the informative intention. So, on the standard view, communication is constituted by common knowledge of the speaker’s attempt to influence the mental life of the addressee.¹³ In the present context, however, one should be suspicious about whether the appeal to common knowledge helps provide sufficient conditions for communication. Recall the cases given in support of the communicative view of

¹³This standard view is from Schiffer (1972). An important alternative version of it is given by Sperber & Wilson (1986), who appeal instead to mutual manifestness, which is a weaker notion than common knowledge. But I will not elaborate upon their alternative, since the difference between the two views does not matter for present purposes: the criticism below of the common knowledge orthodoxy I am about to present is that it does not provide sufficient conditions for communication, so appeal to a notion weaker than common knowledge will not address that criticism. This same point applies for the even weaker version of Griceanism in Neale 1992.

joint attention: Eilan's example of the board meeting example and case i from the introduction. These cases illustrate that common knowledge between the participants of their attention to some object is consistent with a lack of genuine engagement between them. Thus, it should be possible for there to be common knowledge of a speaker's informative intention, yet no genuine joint activity. After all, as observed by Eilan (2020), Grice's original account of communication represents its phenomenon as a relatively impersonal affair: the speaker acts with certain intentions, which the hearer then recognizes. And, since common knowledge does not suffice for joint activity, adding common knowledge of the speaker's intention should not make the overall account sufficient. It seems instead that the standard Gricean view simply pushes back a step the problem for common knowledge raised by the cases motivating the communicative suggestion about joint attention.

As an alternative, I have proposed a refinement of Griceanism that is in line with various recent proposals that communication should be a genuinely joint endeavour (Jankovic 2014, Siposova & Carpenter 2019, Eilan 2020). Crucially, my proposal draws on Tenenbaum's suggestion that genuine joint action requires that the participants intend to engage in an activity that is essentially joint. In the case of communication, again, I propose that the intended activity is that of achieving mutual intelligibility—that of making sense of one another's action—something that must be done together. In order for my alternative to ultimately prove a genuine successor to the common knowledge orthodoxy, it must be shown that it can also handle the main motivation for appeal to common knowledge in the primary communicative intention. The main motivation is the problem of sneaky intentions, and see this footnote¹⁴ for a

¹⁴The problem of sneaky intentions was first raised by Strawson (1964). It was then crucially elaborated by Schiffer (1972), who provided cases of increasing complexity that suggest the addition of common knowledge to fill out the speaker's primary communicative intention. Grice ([1969] 1989), however, suggested that there is in fact something else going wrong in all of the cases put forward by Schiffer. Grice observes that in such cases the speaker intends that the addressee rationally transition from recognition of the informative intention to its fulfillment in a certain way, but at the same time intends that the addressee falsely believe that the speaker intends him, the addressee, to *not* make the transition in that precise way. So, in such a case, the speaker acts with an intention that *P*, but at the same time intends that the addressee think falsely that she, the speaker, has an intention that *not-P*. So Grice suggested that, instead of appealing to common knowledge, his account of communication may simply add the requirement that the speaker does not have intentions of that form. But an immediate problem with that suggestion is that it seems ad hoc. So let me suggest that my joint-action-based Griceanism promises to provide a principled way of ruling out the kind of scheme that Grice identified is present in cases illustrating the problem of sneaky intentions. In general, it is plausible that one cannot act with the intention of working together with another, with the aim of making each other's actions mutually intelligible, while at the same time intending that the other form a false belief about the intentions behind one's action. Finally, see Jankovic (2014) for a closely allied version of Griceanism, which stresses how communication is a form of joint action, and in which it is argued in detail that such a version of Griceanism can deal with the problem of sneaky intentions. Though it is a subtle question whether I can adopt her solution, since she stresses Bratman's account of joint action, yet I go along with Tenenbaum's alternative. Overall, I do not wish to get into the details of how the

brief discussion of how my version of Griceanism resolves that problem.

Before moving on, I wish to briefly respond to a worry that the fan of common knowledge might raise concerning the view of communication laid out in this section, since responding to it anticipates a consequence of the view that I will return to in the next chapter. To understand the worry, recall first how Tenenbaum's suggestion was motivated by a case in which common knowledge of interlocking intentions with non-essentially-joint content was not sufficient for joint action. Given that motivation, it might seem that the view of joint action from Tenenbaum holds that common knowledge still plays a role, as it is simply claiming that there must be common knowledge of intentions with essentially joint activities represented in their content. Thus, the worry is that, if that is the right version of Tenenbaum's view, then the account of communication I have provided in this section does not allow that joint attention is fully independent of any sort of common knowledge—contrary to what was promised in the introduction to this chapter—as communication on my view requires common knowledge.

The response to this worry is that the account of joint action from Tenenbaum does not require common knowledge of the relevant intentions. In the next chapter, I am going to return to this point in showing that the refined Gricean view of communication does not require common knowledge of any sort. Now, in the remainder of this chapter, I discuss two issues for my refined Gricean view of communication, the resolution of which guides the development of my overall account of the openness of joint attention and its special role in our social lives.

1.3 Communication and Social Motivation

Recall the circularity issue for Tenenbaum's suggestion: something, X , is claimed to hold partly in virtue of individuals being in mental states such that the specification of the content of those mental states involves appeal to a concept of X . My joint-activity-based Griceanism also partakes in that form of circularity. I already suggested, drawing on an observation from Fine (2012), that there is nothing in principle incoherent or underexplanatory about this form of circularity. Furthermore, it is not even clear why these authors—Fine, Bratman, and Tenenbaum—even take there to be a *prima facie* issue with this form of circularity. There

present account fully solves the problem of sneaky intentions; instead, what I say in response to the problem here is intended as an indication that the present view is on the right track as a replacement of the common knowledge one.

does not seem to be anything wrong with saying, for instance, that someone's action counts as one of intentionally shutting the car door partly in virtue of them intending to shut the car door. Analogously, the actions of two people may be said to amount to the joint activity of, say, walking together partly in virtue of each of them intending to engage in that joint activity.

There is, however, a crucial difference between the two examples just given. In the case of shutting the car door, there is a familiar story as to how one may come to possess the concepts required for specifying the intention's content. By perceptually interacting with the world, one can learn about car doors and what it is for something to shut. In contrast, the acquisition of a concept of a joint activity does not seem to fit the familiar mold of learning from perceptual experience, for if such a concept is one that is gained, then it is one that seems like it must be gained partly in virtue of engagement in joint action—not mere observation of external reality. But then there is a conflict with Tenenbaum's suggestion that engagement in joint action holds partly in virtue of the deployment of a concept, already possessed, of an essentially joint action: it cannot be both that joint activity holds partially in virtue of the deployment of a concept, where that concept itself must be gained partially in virtue of engagement in joint activity; otherwise, either engagement in joint activity is impossible or the concept is not one that is gained.

So any account of joint action that takes Tenenbaum's suggestion on board faces an issue concerning how to properly account for how we gain the concept of the joint activity that the account appeals to. Again, the issue arises because it is plausible that the concept of any joint activity is one that is gained partly in virtue of joint engagement. I label this issue the "concept acquisition" problem. Here is how this problem applies in particular to my account of joint attention and communication from the previous section. I have claimed that episodes of joint attention hold partially in virtue of a form of communication that involves deployment of the concept of the joint activity of mutual intelligibility. But it is plausible that the concept of that joint activity is gained partially in virtue of engagement in the joint activity of achieving mutual intelligibility—in basic episodes of joint attention, for instance. By the end of this section, a solution to the concept acquisition problem will be given on behalf of my account.

The second issue I address in this section concerns the necessity of any form of Gricean communication for sustaining joint attention. Recall the Primacy claim from §1.1: infants begin engaging in joint attention very early on. So one might worry about attributing Gricean intentions and their recognition to infants.¹⁵ So if the communicative conception of joint

¹⁵For more developed arguments against Gricean communication along these lines, see Roessler (2005) and Campbell (2018).

attention is correct, then the relevant form of communication in basic episodes of joint attention is not the joint-activity-based Gricean one, on pain of overintellectualization. It is additionally important, in the present context, to isolate the form of communication that is present in such cases, since, as was argued in 1.1, the main motivation for the perceptualist view of joint attention is based on considerations of basicness, and in particular the direct coordination that must be present in early episodes of joint attention. So it is worth investigating whether the communicative conception can provide an alternative account of direct coordination.

The rest of this chapter uses the necessity worry as a springboard: given that there is some intuitive pull to the claim about overintellectualization, addressing it leads the way to a communicative account of direct coordination—which brings together the seemingly independent considerations from §1.1 concerning joint attention’s basicness with those in §1.2 motivating the communicative suggestion—and also provides an answer to the concept acquisition problem.

1.3.1 The Necessity Worry for Griceanism

The first step in meeting the necessity issue is in generalizing the Gricean conception of communication by recognizing that genuine communication may be based on personal-level motivational states other than intentions. Recall from §1.1 that intentions are *conceptual* personal-level motivational states. Following Dickie (2015), I label *nonconceptual* personal-level motivational states “needs.” Paradigm examples of needs are hunger and thirst. Animals and infants may be so driven without possessing the concept of, say, nutrition. Personal-level motivational states in general, whether needs or intentions, have the common feature that they are intimately related to genuine action: an action—as opposed to a reflex or something that happens to us—is a movement guided by a personal-level motivational state (Frankfurt 1978). Given the existence of needs, it is another short step to recognizing the possibility of a form of communication that is genuine action, yet not based on intentions.

Accordingly, I claim that there is a form of communication driven by a social need. The content of this need is to make sense of each other’s actions: we are driven to mutually establish one another’s intelligibility. Specifying the content of this need appeals to the same concept as for the communicative intention of the joint-activity-based Gricean account of communication from the previous section. Recall how on that account the speaker must have the informative intention to bring about a change to the addressee’s mental life, but also the

communicative intention that the addressee recognize the informative intention *by virtue of the joint activity of making sense of another*, and for that recognition to be part of the addressee's reason for bringing about the intended effect. Though, again, needs are nonconceptual, so one can be driven by a need without possessing the concepts required for its specification. I now discuss empirical evidence that supports the existence of this need, as well as recent related proposals by philosophers and psychologists. This discussion will help elaborate my claim about the need's content.

The claim that some form of communication occurs before and within basic episodes of joint attention is empirically supported by what psychologists call the transition from "primary" to "secondary" intersubjectivity (Trevarthen 1979, Rochat & Striano 1999, Adamson & Russell 1999, Hobson 2002). Throughout the first year of life, infants engage in one-to-one emotional communication with others: primary intersubjectivity. There is a huge amount of evidence that as early as the first couple months infants are active and aware participants in emotionally laden social exchanges. See this footnote¹⁶ for some empirical detail. Here is a psychologist's description of a 2-month-old (Tronick et al. 1978).

Baby is looking off to side where mother will come in. He sits completely quiet, back in his baby seat, face serious, cheeks droopy, mouth half open, corners down, but there is an expectant look in his eyes as if he were waiting. His face and hands reach out in the same direction. As his mother comes in, saying, "Hello" in a high-pitched but gentle voice, he follows her with his head and eyes as she approaches him. His body builds up with tension, his face and eyes open up with a

¹⁶In particular, there is the still-face effect, which was first established by Tronick et al. (1978) in studies involving infants between 2 and 20 weeks old and their mothers. The behaviour of the infants was compared in two kinds of interaction with their mothers: normal and still-faced. In the normal interaction, the mother played with the infant. Normal playing involves, for instance, miming of the infant's facial expressions, as well as intervening to coax it away from unhappy expressions, or to contain over-excited ones. In the still-face interaction, the mother sits in front of her infant and stares at it with a neutral face. In these cases, the infant attempts to in many ways to engage the mother, but after this is unsuccessful it withdraws and expresses wariness and helplessness. The still-face effect is the significant difference in infants' behaviours between these two different kinds of interactions. According to the meta-analysis of Mesman et al. (2009), similar results have been found in over 80 studies since Tronick et al. 1978. These studies vary with respect to the age and gender of the children, as well as whether the adult engaging with the child is a parent or stranger. Several studies have been done with non-Western children, and found the effect as well. In addition to the still-face effect, an important source of empirical support for the connection between the social motivation and early episodes of joint attention is that there are negative correlations between Autism Spectrum Disorder (ASD) and engagement in joint attention, among other aspects of social cognition, which supports the present proposal given the recent view that ASD is characterized by a deficit in social motivation (Chevallier et al. 2012, Nyström et al. 2019). For in-depth overview of developmental as well as evolutionary considerations in support of the social motivation's relevance for joint attention, and infant sociality more generally, see Tomasello (2014, 2019).

real greeting which ends with a smile. His mouth opens wide and his whole body orients towards her. He subsides, mouths his tongue twice, his smile dies, and he looks down briefly, while she continues to talk in an increasingly eliciting voice. During this, his voice and face are still, but all parts of his body point toward her. After he looks down, she reaches for and begins to move his hips and legs in a gentle, containing movement. He looks up again, smiles widely, narrows his eyes, brings one hand up to his mouth, grunting, vocalizing, and begins to cycle his arms and legs out toward her. (5)

In addition to one-to-one affective communication—the primary intersubjectivity illustrated in the passage just given—throughout the first year of life infants separately engage in gaze-following and lone attention to objects. Then, around the first birthday, all of these activities come together and infants begin sharing with others their emotive reactions to objects: the dawning of secondary subjectivity (Striano & Bertin 2005). Here is a psychologist’s description of an early episode of secondary intersubjectivity (Stern 1985).

A nine-month-old girl becomes very excited about a toy and reaches for it. As she grabs it, she lets out an exuberant “aaah!” and looks at her mother. Her mother looks back, scrunches up her shoulders and performs a terrific shimmy with her upper body, like a go-go-dancer. The shimmy lasts only about as long as her daughter’s “aaah!” but is equally excited, joyful, and intense. (140)

Note that described here is a basic episode of joint attention—the type appealed to at the beginning of §1.1. Finally, the active and aware engagement illustrated in the above passages is explained if the infants’ activity is genuine action driven by a motivational state that concerns others.

An action driven by the social motivation need not draw upon any intentions or other conceptual mental states. Compare action driven by hunger, another human need. In some cases, one forms beliefs and corresponding intentions about how best to satisfy one’s hunger. Imagine you are deliberating about which restaurant to visit. But sometimes one’s action is driven by hunger without any intervening conceptual mental states. Imagine your ravenous appetite has you pulled in, without reflection, by the smell of a familiar fast food chain. Similarly, I claim that the social motivation may guide one’s action in communicating without the formation of beliefs or intentions. In particular, one need not have the Gricean informative intention—a part of the Gricean account of communication from the previous section—to

influence the addressee's mental life in some way. In the simple form of communication I am here considering, the infant is driven by their need to make sense of and be made sense of by others in whatever way they can, which in early life involves the simple expression of emotion.

There are a number of recent proposals regarding a social motivation from philosophers and psychologists. Let me say a few critical words about these alternatives in order to sharpen my proposal. First, Eilan (2015) speaks of a "desire for relatedness" or "basic human urge to connect" (12), which may be unpacked slightly as a need for social bonds. Yet even when there is a strong bond between two individuals—a child and caregiver, for instance; or simply two adult friends—there is continued drive to engage in joint activities. Thus, Campbell's (2018) proposal that we have a need for "cooperative social interaction" seems superior (124), since instead of binary relatedness the need Campbell proposes has ternary content: we have a need to engage with others in joint activity.¹⁷ Campbell also proposes that the need might be specifically for joint attention, and a similar proposal is made by Dickie (2020), who claims that the need is to engage with others towards objects in the world in such a way as to mutually and rationally gain information about those objects. I take from Dickie and Campbell the insight that the social motivation has ternary content. But I hold that the ternary content does not appeal directly to things in the world about which we are trying to engage with others. Rather, the motivation is to engage in the joint activity of making sense of one another in light of our actions; so there is more than simply the binary content of connecting with another, but the additional aspect does not directly concern independent objects. By not appealing to independent objects, as Dickie and Campbell do, the present account explains how the social need drives early one-to-one episodes of communication, and thus also how basic episodes of joint attention are related to those one-to-one interactions.

Dickie's and Campbell's proposals are closely allied with Tomasello et al.'s (2005) influential "Shared Intentionality Hypothesis." This hypothesis is about what is unique about human cognition and its development, and it claims that it is our engagement in joint action. Accordingly, the hypothesis includes the claim that we have a "species-unique motivation to share emotions, experience, and activities with other persons" (Tomasello et al. 2005: 675). As the hypothesis is elaborated by Tomasello (2009, 2019), communication is a slightly sophisticated form of joint activity, which comes after and partly in virtue of participation in joint attention.

On my view, however, which takes seriously the communicative suggestion about joint attention, the most basic type of joint action is communication. Let me explain. Recall first the

¹⁷Thanks to John Campbell for discussion on this point.

proposal about joint action from the previous section: two individuals are doing X together just in case each are doing X because both intend to do X together. In the beginning of the present subsection, the existence of personal-level motivational states other than intentions was revealed, so it is natural to broaden that account of joint action in the following way: two individuals are doing X together just in case each are doing X because both are in a personal-level motivational state with the content to do X together. On this broadening, it is not a requirement that both participants intend to perform the relevant activity together. Instead, since the only motivational states other than intentions are needs—again, intentions are the conceptual personal-level motivational states, and needs are the nonconceptual ones—it may be that at least one of the participants is acting under a need to do X together. Now, I am proposing that the social motivation is the only human need with the relevantly joint content. Thus, achieving mutual intelligibility is the only joint activity that can be engaged in without an intention with joint content. Furthermore, as outlined above, it is one that we engage in very early on in expressing emotions, and we do so without reliance upon any conceptual mental states. Finally, I propose that joint engagement in the activity of achieving mutual intelligibility itself is a form of communication. Thus, the only kind of joint activity that can be engaged in without any intentions—and which we do in fact engage in right out of the gate—is a form of communication.

So what I am claiming about the nature of basic episodes of joint attention is that they are simply the joint activity of achieving mutual intelligibility, but where there is the expression of emotive states concerning some object, which helps sustain mutually responsive experiences of attention to that object. Now, one issue for my view, which seems to favour the Shared Intentionality Hypothesis and its philosophical allies, can be put crudely with the following question: Why, in joint attention, does the object matter at all? There are many ways of making sense of one another: as already stressed, there is a kind of mutual intelligibility attained in virtue of the direct communication of affect—not about any object. So if basic episodes of joint attention are driven by the need for mutual intelligibility, and infants throughout their first year of life are already engaging in direct communication of affect, then why is it that joint attention emerges?

The answer I propose draws on a general fact about need-driven action. The fact is that not all ways of satisfying the content of a need lead to a release of motivational pressure. Consider, for instance, the feeling of hunger and consequent satisfaction upon eating. In order for one's feeling of hunger to be reduced, there are specific ways in which nutrients must be

processed.¹⁸ If one were to place food directly into the stomach or inject calories, one would not feel less hungry. One might hold that specifying the content of motivational states does involve specifying particular routes to fulfilment. Fully responding to this rejoinder would take us too far afield from the present topic, since there are general issues about the content of motivational states that are relevant.¹⁹ Nonetheless, packing in, for instance, ways of getting nutrition into the content of the state of hunger seems immoderate. So I am hopeful that the intuition that the content of, say, hunger does not involve a way of gaining nutrients can be theoretically vindicated.

Applying this general insight to the present topic—the social motivation—I propose that, as an infant’s development proceeds, it becomes crucial for abating the motivational pressure of the social need that the activity be one of making sense of one another in relation to some perceptually available object. In these contexts, sharing attitudes about objects is crucial for releasing the relevant motivational pressure. One might further ask why the relevant motivational pressure has this feature. The answer to this further question presumably lies in the role joint attention plays in cognitive development. Within episodes of joint attention, there is cognitive scaffolding that parents and older peers provide, for which the presence of an object is more beneficial than what is available in simpler one-to-one engagement. So, from an evolutionary point of view, early social interaction is a means to the child getting into a position in which they have the independence to engage with the world. Nonetheless, from a developmental point of view, social interaction is the foreground, since it is a basic need for the joint activity making sense of others and being made sense of driving the child’s activity, and this need itself does not involve some further purpose.

The crucial point from the previous paragraph is that the causal role of motivational phenomenology explains the transition from primary to secondary intersubjectivity. In primary intersubjectivity, recall, there is direct one-to-one affective communication, which is one particular way of satiating the social need. Secondary intersubjectivity involves another particular way of satiating the social need, in which one communicates with another concerning some object. I claim that, as development proceeds throughout the first year, motivational phenomenology pushes one towards the latter way of establishing mutual intelligibility.

Now let me say more about motivational phenomenology, which is the feeling of pressure and release associated with needs. Philosophers have traditionally focused on perceptual phe-

¹⁸For overviews of relevant empirical work on motivation, see Toates (1986) and Berridge (2004).

¹⁹The issue, for instance, of whether the content of a motivational state is an accuracy condition—how things will be—or a fulfillment condition—how things should be (see Dickie (2015: §3) and references therein for further discussion of this distinction).

nomenology, but it should be uncontroversial that motivational phenomenology exists. There is something it is like to feel hungry and sated. Similar phenomenology is associated with the social need: there is a way it is like to need to connect with others, and for this motivational pressure to be released when one engages with another human. Consider, for instance, the satisfaction one feels from doing even mundane activities with others. Suppose you are forced to spend months in social isolation (please forgive the outré philosophical thought experiment). During such a period, one “hungers” for social interaction. And, after such isolation, consider what it is like to finally sit and talk with a friend, or even a stranger. This particular phenomenology is generated from the motivational pressure and release of the social need.

Other defenders of the communicative conception of joint attention have stressed the importance of a similar feeling of connection had when communicatively interacting with others (Eilan 2015, 2018, Siposova & Carpenter 2019). So here is a point, as with many others above, where I am drawing together and refining recent psychological and philosophical work surrounding joint attention, with special attention to that of Eilan (2015, 2018, 2020). But, in the next subsection I show how the motivational phenomenology I have highlighted helps make progress with regard to the direct coordination constraint, from §1.1, and the concept acquisition problem raised at the beginning of the current section.

1.3.2 Direct Coordination and the Concept Acquisition Problem

Recall the direct coordination constraint from §1.1: basic episodes of joint attention are interactions that make other minds transparent, which is to say that participants have immediate justification for taking theirs and the other’s experience to be of the same object. That justification is immediate in the sense that it is not based upon conceptual reflection on one’s and the other’s perspectives on the object. This direct coordination constraint refined the introductory talk about joint attention’s basicness, and is accordingly a desideratum for any account of joint attention.

I also explained in §1.1 how the direct coordination constraint naturally motivates Campbell’s perceptualist account of joint attention: transparency, according to perceptualism, is purely a matter of perceptual phenomenology, which is nonconceptual. But Eilan’s objection to perceptualism, given at the beginning of §1.2, casts doubt on the claim that the transparency of other minds in joint attention is purely a matter of perceptual phenomenology, since she provides a case in which all the conditions perceptualist account hold, yet in which there is intuitively no joint attention. Therefore, *if it is assumed that the only phenomenology relevant*

for direct coordination is *perceptual*, then there must be a non-phenomenological component to direct coordination. Presumably, the non-phenomenological component would be the conceptual content of certain mental states.

Eilan (2015) seems to make the assumption just emphasized, and is thus driven to claim that in joint attention one's experience has conceptual content. She characterizes this content as follows: *we (you and I) are sharing this experience*.²⁰ Crucially, this conceptual content involves a primitive second-personal way of thinking of the other, which is the concept that plays a crucial role in bringing it about that the interaction is genuinely joint. On her view, we are doing something together because we are thinking of each other under "primitive you-awareness" (6). But such an appeal to conceptual content brings along a worry of overintellectualization related to the necessity worry for Griceanism above. More importantly for present purposes, such an appeal means that Eilan's view cannot resolve the concept acquisition problem. If the second-person concept were one that is gained, it would plausibly be so in virtue of engagement in joint activities such as joint attention. So the sense in which Eilan calls the second-person concept "primitive" comes into view: it is not a concept that we gain.

But, as emerged in the previous subsection, there is relevant motivational phenomenology from the social need. Thus, we may follow Eilan, against Campbell, in accepting that perceptual phenomenology is not sufficient for the transparency of joint attention, yet at the same time reject the crucial assumption, which both these authors seem to share, that the only phenomenology relevant for direct coordination is perceptual. Accordingly, as I show in the remainder of this subsection, we may provide a fully nonconceptual account of direct coordination as well as an answer to the concept acquisition problem.

Beginning with direct coordination, I propose that the justification had for taking the two experiences to be of the same object is given by the practical justification of one's activity in jointly attending.²¹ Let me explain the background notion of practical justification. All actions

²⁰See especially pages 15–16 of Eilan 2015. I explained above that experiences have nonconceptual content, but that claim is consistent with them having some conceptual content, for it may be that the way experiences contrast with, say, beliefs is that the latter do not have any nonconceptual content. Regardless, I am rejecting Eilan's view, and my view is consistent with the denial that experiences have any conceptual content.

²¹Here I follow the structure of Dickie's (2015) discussion of coordination in lone thought: she is concerned with *intrapersonal* coordination, in contrast to the present focus on *interpersonal* coordination. Dickie claims that the immediate justification one often has for taking one's beliefs to be of the same object is practical, given that the activity of forming such beliefs is driven by a need to think about ordinary objects. But note that the correctness of her account of the aboutness of individual thought is orthogonal to the present topic. As a brief aside, I suspect that, following suggestions in Campbell (2017, 2018), our individual rational thought about objects may be in some way parasitic upon interpersonal engagement. Further development of the proposal to come in this section might shed light on that suspicion. However, regardless of how that further investigation into the relationship between intrapersonal and interpersonal coordination turns out, where I differ from Dickie here

have a certain baseline normative or rational status, since they are things we do and not mere worldly happenings. But, in addition, some actions have the further positive normative status of practical justification. Theoretical justification may be more familiar to many philosophers. Traditionally, theoretical justification is had by beliefs. For instance, in normal conditions, if one sees or is told that things are thus-and-so, one's resulting belief that things are thus-and-so has the positive normative status of being justified. Contrast such a belief to one formed on a hunch. But it is also natural to talk of actions having similar normative status. Contrast the movements of a skilled archer in aiming at a bullseye with those of a novice. In general, we may say that an action has PRACTICAL JUSTIFICATION just in case it is appropriately related to its guiding motivational state. And, just as with theoretical justification, one might refine this appeal to an appropriate relation along either externalist or internalist lines.

One aspect of practical justification is reliability. Consider how the skilled archer's actions reliably bring about the guiding aim, in contrast with those of the novice. In basic episodes of joint attention, it is plausible that this reliability claim holds true: basic episodes of joint attention consist of emotive communication about the relevant object as well as mutually responsive experiences of attention to that object. In the context of such an interaction, it is plausible that a participant's activity reliably brings about the content of the social need: alongside gaze-following, the sharing of emotive states about the relevant object allow the participants to reliably make sense of what one another is doing. In addition, as mentioned in the previous subsection, the sharing emotive states, driven by the social motivation, is a type of communication that need not involve conceptual resources. And it is also plausible that the capacities of gaze-following neither need involve conceptual resources. Hence, the reliability component does not threaten the possibility of direct coordination.

A simple externalist account of practical justification claims that this reliability component exhausts practical justification. According to this view, an activity has RELIABILIST PRACTICAL JUSTIFICATION in virtue of it being a reliable way of satisfying the content of its guiding motivational state. The proposal that reliability exhausts practical justification is simple and hence appealing. Yet I have already given reason for doubting it. Compare again eating food in the regular way with, for instance, somehow directly placing food into one's stomach. Eating food in the regular way has a special status that other ways of reliably fulfilling the content of

is that I am concerned with the latter phenomenon and she the former. But, in following the structure of her account, the two ways in which the discussion to come is similar to hers are as follows. First, I also explain epistemic phenomena by likening them to more paradigmatic cases of rational action. Second, in doing so I consider two competing refinements of the general notion of practical justification, one broadly naturalist and externalist and the other broadly non-naturalist and internalist, and ultimately endorse the latter.

one's hunger do not; and that difference in special status is explained if not all ways of reliably satisfying the content of a motivational state have practical justification.

There is thus a component to practical justification in addition to the reliabilist one. Now, the insight behind an internalist view of justification is that reliability is not sufficient. With regard to theoretical justification, the subject must have a reason for the belief. A reason may in some cases be a little argument one holds inside their head. But that cannot be all cases, on pain of overintellectualization. In simple cases, the internalist component may be satisfied by the phenomenological character of the relevant mental states (Campbell 2002, Smithies 2011, 2019, Dickie 2015). From the practical perspective, what the subject is doing must make sense to her, for which phenomenology may play a role. Consider a simple prehension: you are thirsty so reach out for your bottle of water on the table in front of you. Your experience of the bottle plays a role in *explaining*—in more than a mere causal sense—the way your arm moves in grasping it; your experience of the bottle as being where it is does not simply cause your arm to move in the right direction, but also illuminates to you as the agent why it should do so. In addition, your thirst and its consequent reduction play roles in explaining that same activity: your thirst is part of *why* you act in that way—in more than a mere causal sense—and the action only makes sense in that light if it reduces that feeling. Or, in our running example, only the ways of reliably satisfying the content of one's hunger that are appropriately related to its motivational phenomenology of pressure and release have practical justification. Generalizing, an activity has PHENOMENOLOGICAL PRACTICAL JUSTIFICATION in virtue of it (i) being a reliable way of satisfying the content of its guiding motivational state, but also (ii) bearing an appropriate relation to the phenomenology of the guiding motivational state as well as the relevant experiences.²²

Accordingly, both perceptual and motivational phenomenology play a role in accounting for the practical justification had by the activity of a participant in a basic episode of joint attention. The total phenomenological character of such an episode involves perceptual phenomenology from the experience of attending to the object and the other, and motivational phenomenology from the social motivational state. Recall Campbell's perceptualism, on which the total phenomenological character of an episode of joint attention is exhausted by perceptual phenomenology. That account of joint attention is a simple extension of Campbell's approach to lone attention: in joint attention, perceptual phenomenology is like a win-

²²This account might need to be generalized further, so that the internalist component may also be satisfied by beliefs, since beliefs are also sometimes rationally related to action. For instance, you believe that collecting possessions is the key to happiness, so every day you go to the mall. But I set aside beliefs since they are not relevant to the simple activity relevant to the present topic.

dow that makes other minds transparent alongside an external world of objects and their properties. So on his view there is a parallel between the basic way that the external world and other minds are revealed to us. But one might be unhappy with the metaphor this strong analogy generates: in our relation to social reality it is not as if we are alone, each peering at one another through our solitary windows, not noticing the empty street below.²³ And, as Eilan argues, that Campbell's view provides such a straightforward analogy leads to its failure in capturing genuinely joint attention, since there are cases of mutual covert attention where the relevant perceptual phenomenology is present but not joint attention. In contrast, it is plausible that the present proposal can meet Eilan's challenge, since both perceptual and motivational phenomenology play a role in the practical justification that grounds direct coordination, and the relevant motivational phenomenology is not present in cases of mutual covert attention.

Having shown that the direct coordination present in basic episodes of joint attention is nonconceptual, we are now in a position to address concept acquisition problem as it arises for my account. Recall that resolving that problem involves explaining how the concept of the essentially joint activity of making sense of one another may be gained. In order to provide that explanation, we must identify a way in which concepts may be gained partly in virtue of engagement in joint activity, and not merely on the basis of perceptual experience.

There are two important features of experiences, already emphasized, which together show how experience provides the paradigm basis for gaining concepts. The first feature is nonconceptuality: having an experience with a certain content does not require one possess the concepts necessary for the specification of that content. This feature secures the possibility that experience may explain how one gains a concept, since having an experience does not itself require that one possess the concepts relevant for specifying its content. The second feature is rich phenomenology: what it is like to have an experience consists of a detailed view of external reality. This feature fills out a positive account of the way in which experience explains how we gain concepts on that basis: what it is like to have an experience is a rich presentation of how things are "out there."

The social need is relevantly similar to experiences in having those two features. Recall that the need is a nonconceptual motivational state. It also has a rich motivational phenomenology of pressure and release related to certain joint interactions with others. But there is one crucial qualification: this phenomenology is not transparent in the sense that perceptual phenomenology is, as it does not consist of a view of external reality. Instead, that

²³Thanks to Johannes Roessler for discussion on this point.

motivational phenomenology is based upon our own need to understand others and make ourselves understood. Thus, the way in which we gain the concept of the essentially joint activity of making sense of one another is by conceptualizing ourselves—our own “form of life” (Wittgenstein [1953] 2009: §19)—as we are pushed early on by our social need into basic episodes of joint attention.

1.4 Looking Ahead

Let me conclude by first summarizing the main steps of the discussion in this chapter. I began in §1.1 by establishing the direct coordination constraint on basic episodes of joint attention, which refined the central motivation for a perceptualist account. In §1.2, I systematized the communicative conception of joint attention, as recently motivated in contrast with perceptualism, and identified a version of Gricean communication that could play a sufficient role in filling out the communicative conception. In the final section §1.3, I began by responding to the necessity worry that Gricean communication is not present in basic episodes of joint attention, and in doing so uncovered a more basic type of communication, driven by a deep-seated social need. I then showed how appeal to the social need helped provide a non-perceptualist account of direct coordination, and also meet a related challenge to the version of Gricean communication identified in §1.2.

As already stressed, much of the chapter has consisted of drawing together and developing related and recent ideas from philosophers and psychologists: how joint attention is a basic form of social interaction, the communicative suggestion about joint attention, and the existence of a social motivation. So let me now consolidate the view of joint attention and communication that has emerged and how it makes progress on understanding joint attention and its special role in our social lives. The overall proposal is that there are two types of communication that may sustain joint attention: the sophisticated Gricean one, and the simpler one based on the social need. Sophisticated episodes are sustained by joint-activity-based Gricean communication, in which the concept of making sense of one another is a part of the conceptual content of the relevant intentions. Basic episodes of joint attention are sustained by communication based upon the social need, in which jointness is brought about by our intentional activity, yet not in virtue of conceptual content.

The relationship between these two types of communication is as follows. Right out of the gate, we are pushed by our social need into communicative exchanges, including joint attention. We may then gain a conceptualization of ourselves and thereby gain a concept of

the essentially joint activity that is deployed to form joint Gricean intentions. The nonconceptual way of establishing jointness in basic episodes makes achieving mutual intelligibility our own form of life, which we are able to discover in virtue of how it feels to act together. Contrast this picture again with the Shared Intentionality Hypothesis from the previous section. What that influential hypothesis claims is that we have a basic drive to engage in joint activity, of which communication is a slightly sophisticated form brought about in virtue of engagement in more basic forms, such as joint attention. In contrast, however, I have taken the communicative suggestion about joint attention seriously, which leads to the view that communication is a more fundamental form of joint activity than joint attention.

So, to go out on a limb, if all joint action—including, for instance, walking together—holds partly in virtue of communication,²⁴ then my picture of joint attention may also shed light on joint action in general. The general picture is thus that there are two types of joint interaction: one driven by the social need, and another that holds partly in virtue of sophisticated intentions. Hence, an interesting feature of the resulting account is that it provides a *non-unified* account of jointness. In some cases we are drawn together in a committal way by our social need. Yet we may do things together with people who we would like nothing to do with. Consider the vicious back-and-forth of a presidential debate. These other cases hold in virtue of sophisticated intentions. Nonetheless, the committal form remains developmentally basic.

I do not develop this promised general account of joint action in my dissertation. I leave that project to future work. I focus instead on developing the insights that this investigation into joint attention has revealed about the nature of communication, especially with regard to the use of demonstratives. Recall, as was discussed in §1.2, that the communicative conception of joint attention implies the falsity of the commonplace view that demonstrative understanding sometimes holds in virtue of joint attention. In chapters 2 and 3, I defend that provocative implication of the communication, and in so doing I develop an account of what demonstrative understanding does involve. My account draws upon the joint-activity-based Griceanism already introduced in this chapter, and because it do so it faces a challenge that any broadly Gricean account of communication faces when applied to linguistic communication, which is to explain what is distinctive about the use of language in communication. So, in the final chapter 4, I elaborate my response to that challenge by developing a formal semantics for demonstratives, and related expressions, which reveals how their distinctively linguistic contribution is a constraint on how they may be used in communication.

²⁴This proposal is suggested by Milward & Carpenter (2018) and Roessler (2020).

Chapter 2

Demonstratives and Joint Attention

Imagine two friends sitting on a park bench, commenting on the passing show. One points to an old flower bed, a single rose remaining in it, proclaiming, “That rose has a troubled beauty.” The other looks over, perhaps with an eye roll. As in this example, with the utterance of “that rose,” demonstratives are often used with pointing gestures to talk about nearby objects.¹ So it is *prima facie* plausible that, at least in a certain range of cases, demonstrative understanding involves joint attention to the relevant object—as was already mentioned in §1.2 of the previous chapter. Moreover, as we are about to see in more detail, this connection between joint attention and demonstrative understanding has found theoretical motivation in Campbell’s (2002) work, and more generally from the Neo-Fregean tradition.

In the previous chapter, however, an account of joint attention was built around the suggestion that any episode of it holds partly in virtue of communication. So, in an episode of joint attention where the relevant communication involves the use of a demonstrative, demonstrative understanding is the linguistic communication that helps sustain joint attention. Hence, joint attention cannot also be a part of how that demonstrative understanding holds; if my account of joint attention is correct, then joint attention is never a part of demonstrative understanding.

There is thus a tension to be resolved, and I do so in this chapter by undermining the motivation for the claim that demonstrative understanding ever requires joint attention, alongside providing positive reason against that claim. To begin, let me state clearly the alleged connection between demonstratives and joint attention.

¹I take the category of demonstratives to include complex, such as “that rose” in the example, and simple ones, such as “this” or “that” used without an accompanying noun-phrase complement. For the purposes of this chapter and the next, I also include third-person pronouns, such as “he” and “it.” in the category of demonstratives, though in the fourth chapter we will see reason for distinguishing them from this category.

Consider the claim that any utterance of a demonstrative requires for its understanding that the speaker and a hearer jointly attend to its referent. This crude claim faces an immediate problem, since anaphoric and bound uses of demonstratives do not have referents in the relevant sense.² So we must restrict attention to deictic uses of demonstratives. Yet deictic demonstratives may be used to communicate about objects distant in time or space, and even abstracta. Consider the following.

- (2) That hominid who discovered how to start fires must have gotten pretty lucky. (Inspired by King 2001)
- (3) That first baby born in the 22nd century will either have an existence incredibly long or incredibly short. (Inspired by Kaplan 1989)
- (4) Scientists have been trying to determine the material constitution of this planet [pointing to a galaxy map], but it is too distant to reliably tell.
- (5) That highest prime continues to elude us!

If need be, imagine (5) uttered by one of a group of amateur mathematicians.

Furthermore, even if the object is perceptually available to the interlocutors, joint attention is not always required for understanding. Imagine you are a spy hiding in the bushes and overhear me soliloquizing about the sole rose in my garden. As Campbell (2002) notes, in such a case you can understand my uses of demonstratives, yet there is no joint attention. So we must restrict our attention—as Campbell suggests—to ordinary conversation.

Thus, for the defender of a joint attention requirement on demonstrative understanding, the relevant thesis must concern a certain type of use of deictic demonstratives. Call the type for which the thesis putatively holds *PERCEPTUAL DEMONSTRATIVES*. Thus the focus of this chapter is the following thesis.

The Thesis Perceptual demonstratives require for their understanding joint attention to the referent on behalf of the interlocutors.

²These different uses of demonstratives are discussed in the final chapter. But here is some simple data to illustrate the taxonomy. The first is deictic; the second, anaphoric; the third, bound.

- (i)
 - a. *Context: A dog walks into the room.*
That dog is beautiful.
 - b. A woman walked in. That woman was happy.
 - c. Every guilty senator testified to congress before that senator was harangued.

Given the considerations above, a perceptual demonstrative is a deictic demonstrative used within ordinary conversation, where the referent is an object perceptually available to the interlocutors. This characterization should be clear enough to proceed.

As a brief aside, note that when philosophers talk of perceptual demonstratives they often mean demonstrative concepts. A DEMONSTRATIVE CONCEPT is a way of thinking of an object that is intimately tied to perceptual experience.³ But when I talk of demonstratives simpliciter I mean (token) natural language expressions, or their use (the distinction between token demonstratives and uses of demonstratives will only become important in the final chapter and conclusion). So perceptual demonstratives are a type of use of natural language demonstratives. Since thought is relevant to linguistic understanding, in what follows I do talk of demonstrative concepts, but I do so with explicit qualification.

The Thesis is endorsed by Campbell (2002), Dickie & Rattan (2010), and Seemann (2019). These authors may not explicitly refine it as I have done, yet it should be clear that it must be refined in some such way. The most explicit motivation for the Thesis is given by Campbell (2002), and I consider that argument in the next section. Afterwards, I develop and evaluate the background motivation for the Thesis that arises from the Neo-Fregean tradition that all these authors work within.⁴

2.1 Coordinated Attack and Common Knowledge

Campbell's (2002) argument is based on comparing two cases of potential rationally coordinated action towards a target. First, there is the famous case of two generals, standing on opposing hills, overlooking an enemy's encampment in the valley between them (Fagin et al. 1995). They are tasked with destroying the enemy, and can only succeed if they send their armies at the same time. The payoff for their success is very high, but the outcome in which only one sends their army is disastrous: the opposing army will destroy that single general's army, then shortly after easily destroy the other's in another one-on-one battle. So, if only one general attacks, the war is surely lost. The value of the final outcome, where both generals withhold from attacking, is not nearly as disastrous as sure defeat, but also not nearly as great as sure victory.

In attempting to coordinate their attack, the generals can only communicate by sending a

³Important accounts of demonstrative concepts in this vein are given by Evans (1982, 1985), Campbell (1997, 2002), Smithies (2011, 2019), and Dickie (2011, 2015).

⁴In addition to the authors mentioned in this paragraph, the tradition has also been importantly elaborated by Evans (1982, 1985), McDowell (1998a,c), and Heck (1995, 2002).

messenger to carry a note through the intervening valley. So one of the generals, *A*, might send a note telling the other that they should both attack at noon the next day. But, as the generals are aware, there is a chance that any messenger will be captured by the enemy. Suppose that the other general, *B*, does receive the note; he will presumably reason as follows: *A is not sure whether this message got through, and since there is much to lose if both of us do not attack at once, I will send a note confirming that I received A's note.* Now, if *B*'s response does not get through, *A* will presumably reason as follows: *I have not heard anything back and surely B would have tried to send a response if they received my original message, so either (i) my original message did not get through or (ii) B's response did not; but (i) is more likely than (ii), since the likelihood of a message not getting through (mine) must be higher than the likelihood of a message not getting through (B's) and one getting through (mine).*⁵ So if *A* does not receive a response from *B* to his original response, *A* will not attack, since it is more likely than not that *B* has not been made aware of when to attack with him. Suppose, on the other hand, that *A* does receive the response. He will presumably realize that *B* will reason in the same way that he, *A*, would upon not receiving a response to the original message; thus, *A* will send a response to *B*'s response. Similar reasoning can be iterated regardless of how many message are successfully sent back and forth.

Now compare the predicament of the generals with that of two hunters, who have been searching for a beast of prey that has been terrorizing a nearby village. They are now hiding in some bushes, having gotten the drop on the animal. If the hunters fire at the same time, they will kill the animal, bringing great benefit to themselves and their people. But, if only one fires, the animal will become merely aggravated, causing it to notice the hunters and attack. The hunters will not be able to adequately defend themselves in that case, given how swift the animal is. Finally, both may choose not to shoot, and wait for another day to find the animal again. Just as for the generals, this final outcome has a value between that of both hunters firing and only one of them firing.

Campbell observes that it is possible for the hunters to rationally undertake joint action towards their target. A natural way this might go involves communication with a perceptual demonstrative. One hunter looks at his partner, uttering "Now is the time to shoot that animal." The other understands, and they shoot together. By reflecting on this case of the two hunters, the following general principle receives intuitive support, and I take the principle as

⁵Here is the reasoning that (i) is more likely than (ii) in a bit more detail. Suppose there is single, positive probability ϵ of a message not getting through. The probability of (i), *A*'s message not getting through, is ϵ . The probability of (ii), *A*'s message getting through but *B*'s not, is $(1 - \epsilon)\epsilon$, which is less than ϵ .

the first premise of Campbell's argument for the Thesis.

Premise 1 The understanding of a perceptual demonstrative is sufficient for the possibility of rationally coordinated action towards the relevant object.

Furthermore, it is intuitive that the hunters, in rationally coordinating their action towards the target, must engage in joint attention towards the referent. In uttering "Now is the time to shoot that animal," it is natural to imagine the speaker looking at his partner then the animal, and the partner following the speaker's gaze. Moreover, in contrast with the hunters, it seems that the generals neither have the possibility of rationally coordinated action towards their target, nor the ability to engage in joint attention towards the enemy's camp. Thus, reflecting on the cases of the hunters and the generals intuitively supports the following generalization, which I take as the second premise of Campbell's argument.

Premise 2 The possibility of rationally coordinated action to some object requires joint attention to the object.

We have thus arrived at Campbell's argument for the Thesis, as it follows as conclusion from the premises just given.

Conclusion Understanding of a perceptual demonstrative requires joint attention to the relevant object.

In response to this argument, I do not dispute that the hunters but not the generals are able to engage in joint attention to their respective targets. What I argue is that there can be understanding of perceptual demonstratives between the generals. Thus, the case of the generals in fact provides a positive argument against the Thesis, in the form of a counterexample. Moreover, the way in which this point—that the generals can communicate with perceptual demonstratives—defuses Campbell's argument for the Thesis depends on whether Campbell is right that for the generals rationally coordinated action towards their target is impossible. If Campbell is right in that respect, then his argument is unsound because premise 1 is false. But if he is wrong—as recently argued by Lederman (2017)—then premise 2 of the argument is false. I do not take sides in this dispute between Lederman and Campbell; instead, in the rest of this section I explain how there can be understanding of perceptual demonstratives between the generals.

It has recently been pointed out that it is simply an intuitive fact that the generals are able to communicate (Jankovic 2014, Harris 2019). Imagine again one of the generals sending

a note to the other. It may read as follows: “Let us attack that camp at noon tomorrow. Also, do you see that embankment to the left of the camp? When we attack, let us meet there.” Regardless of whether the generals are able to rationally attack on the basis of such notes, it is pre-theoretically desirable to explain how there can be understanding, including of the contained demonstratives. Moreover, the authors just cited, who pointed out this intuitive fact, use it to argue against standard Gricean accounts of communication, which appeal to common knowledge. I explain in a moment how my refined Griceanism, introduced in §1.2 of the previous chapter, accounts for the generals’ demonstrative understanding. First, however, let me address a couple of potential responses that the defender of the Thesis might make to the counterexample of the generals.

Recall that the way I have characterized perceptual demonstratives involves the qualifications that they are deictic demonstratives (a) where the referent is perceptually available to the interlocutors and (b) used in ordinary conversation. With regard to (a), the enemy camp is perceptually available to each of the generals, for both can attend it—though, crucially, the generals are not in a position to have mutually responsive experiences of the camp, which is precisely why they cannot *jointly* attend to it. Therefore, the best response, on behalf of the Thesis, to the intuitiveness that there can be demonstrative communication between the generals is to rely on qualification (b) and claim that the generals are not engaging in *ordinary* conversation. But, it seems that the main reason one might hold that the understanding here taking place is non-ordinary is if one accepts that some sort of common knowledge is essential for capturing communication’s characteristic openness, as in standard Griceanism. Common knowledge of the content of a message, on the basis of its being sent, cannot be obtained by the generals, for the reasons outlined above in the presentation of their predicament. But, as I am about to explain, my refined Griceanism view of communication explains the openness of communication in a way that is independent of common knowledge.⁶ Thus, there is no theoretical pressure to accept the orthodoxy’s counterintuitive verdict about the generals’ ability to communicate with perceptual demonstratives.

Let me now present my refined version of Griceanism, applied to the case of a perceptual demonstrative.

Refined Gricean Communication (perceptual demonstratives) There is communication;_{ig}

⁶In the previous chapter, two types of communication were identified: a refined Gricean account that appealed to sophisticated intentions and their recognition, and a simpler but related form that appealed to a basic social motivation. Here and for the rest of this dissertation, I focus, mostly for simplicity, on the refined Gricean account.

(in the joint-activity-based Gricean sense) between speaker *S* and addressee *A* of an uttered perceptual demonstrative just in case:

1. *S* intends that *A* look at some object *o* (*the informative intention*),
2. *S* intends to engage with *A* **in the joint activity of making sense of one another's actions**, where successful engagement in this activity generates *A*'s recognition of the informative intention (*the primary communicative intention*)
3. *S* intends that this recognition serve as at least a part of *A*'s reason for producing that effect (*the secondary communicative intention*), and
4. the primary communicative intention is fulfilled (*the uptake condition*).

In making sense of this application, it must be noted immediately the background Gricean framework is that the speaker's informative intention is an attempt to influence the addressee's mental life. Therefore, it is not necessary in an informative intention that the speaker attempt to influence the addressee's belief, since there is more to mental life than belief.⁷ So I have proposed that for perceptual demonstratives the speaker's informative intention is that the addressee look at some object, and, more generally, I propose that for deictic demonstratives the informative intention is that the addressee think of some object. So, to use slightly better terminology, I follow Bach (1992) and others in calling the informative intention in these cases the REFERENTIAL intention. Thus, in brief, the refined Gricean account holds that there is understanding of a deictic demonstrative, or more specifically a perceptual demonstrative, when the addressee recognizes the speaker's referential intention, where that recognition is generated by the joint activity of mutual intelligibility.

As mentioned at the end of 1.2 of the previous chapter, one might worry that my account does not avoid a requirement of common knowledge, since it holds that understanding involves joint action, and joint action, the worry presupposes, requires common knowledge of some sort. But let me explain how my account of communication as joint activity does not require common knowledge of any sort. Following Tenenbaum's suggestion about the nature of joint action, the interlocutors are engaging in the joint activity of mutual intelligibility when they are both making sense of each other, and each are doing so because both intend to engage in the joint activity of making sense of one another (it is crucial that this "because" is not understood in a merely causal sense, but in the sense in which an agent's action is explained qua intentional action). There is simply not a requirement of common knowledge

⁷This point has been recently emphasized by Moore (2017).

added regarding the intentions to act together. Furthermore, there is reason to not add such a requirement—independent from theorizing communication as a form of joint activity—since cases can be given in which intuitively there are other kinds of joint action but in which there cannot be such common knowledge (Jankovic 2014).⁸

Finally, let me mention an issue for any broadly Gricean account of communication, including my own. The Gricean approach focuses on face-to-face conversation as the paradigm of ordinary communication, as it primarily attempts to explain what is involved in the understanding of an utterance between a speaker and addressee. But there can be communication, in some sense, in cases where there is no speaker and *addressee*. Recall, for instance, the case mentioned in the introduction of this chapter of a spy overhearing another speaking. Or consider reading a book by a long dead author, such as Plato's *Republic*. In response, the Gricean should rely upon the notion of ordinary, or paradigm, communication, and claim that the cases just mentioned are not of primary concern, since they are not paradigm cases of communication in the way that face-to-face communication is. My refined Griceanism shares in the overall Gricean focus on face-to-communication, since face-to-face conversation is clearly a joint activity, something the interlocutors do together. And cases such as those just mentioned—the spy, reading a book—there is clearly no joint activity. Accordingly, I endorse the response just given on behalf of the Gricean, which is that such cases are nonparadigm—though in the final chapter I outline a distinctive aspect of linguistic understanding that may be present regardless of whether one is engaged in a joint activity, which thus goes some of the way in accounting for what is happening in the case of, for instance, reading a book.

The reason I have raised this issue is to emphasize is that the kind of paradigm communication in face-to-face conversation and what is occurring in reading a book are at opposite ends of a scale, where the sense in which communication consists of understanding, as a joint activity, slowly tapers off. In between those endpoints, but closer to the side of face-to-face conversation, is the case of the letters that the generals are sending—alongside cases of phone calls, texts, and e-mail exchanges. Thus, the understanding that the generals may have of each other on the basis of their letters does partake in the paradigm form that is the target of my refined Griceanism, despite the fact that the sense in which it is a joint activity is slightly less vivid than the paradigm case of face-to-face conversation.

⁸Such cases are similar to that of the two generals, in which joint action holds but in which common knowledge of the relevant intentions is not possible. I am not going to present any such cases, given considerations of space, but I note they are not of joint actions that are targeted at particular objects; thus, the possibility of certain kinds of joint action for individuals in circumstances like that of the generals is neutral on whether the generals may rationally coordinate their attack on the enemy.

2.2 Transparency and Perceptual Demonstratives

Let me now move on from Campbell's explicit argument and uncover some background motivation for the Thesis—that perceptual demonstratives require joint attention for their understanding—from the Neo-Fregean tradition. The tradition in question goes back to Frege (1960a,b), and emphasizes the notion of transparent communication as a successor to Frege's claim that sameness of sense is required for referential communication. In brief, the argument is that if there is any transparent demonstrative communication, then it should occur with simple perceptual uses of demonstratives, and in such cases joint attention is precisely what grounds the direct rational connection.

I begin unpacking this argument by considering a passage from Frege.

...Herbert Garner knows that Dr. Gustav Lauben was born on 13th September, 1875 in N.N. and this is not true of anyone else; against this, suppose that he does not know where Dr. Lauben now lives nor indeed anything about him. On the other hand, suppose Leo Peter does not know that Dr. Lauben was born on 13th September 1875, in N.N. *Then as far as the proper name "Dr. Gustav Lauben" is concerned, Herbert Garner and Leo Peter do not speak the same language ...*

Accordingly, with a proper name, it depends on how whatever it refers to is presented. This can happen in different ways and every such way corresponds with a particular sense of a sentence containing a proper name. The different thoughts which thus result from the same sentence correspond in their truth-value, of course; that is to say, if one is true then all are true, and if one is false then all are false. Nevertheless their distinctness must be recognized. *So it must really be demanded that a single way in which whatever is referred to is presented be associated with every proper name.* (Frege 1960b: 297–298, emphasis added).

From the first part of this passage we may take Frege's insistence that there are different languages as simply the claim that there is no understanding between Herbert Garner and Leo Peter with their use of the proper name "Dr. Gustav Lauben." In the second part of the passage, Frege suggests the following explanation for that lack of understanding: a proper name has a single way of thinking of its referent—a sense—associated with it, which determines that understanding requires that both interlocutors think of the referent in that way. Applied to demonstratives, the Fregean suggestion here is that a token demonstrative is associated with a way of thinking of its referent.

In this section, however, I set aside theses about the semantic CONTENT of (token) demonstratives: what they contribute to determining the proposition expressed by their containing sentence. I take instead from this passage the following claim about demonstrative understanding: understanding a demonstrative involves the interlocutors thinking of the relevant object in appropriately related ways.⁹ I use the terminology of “coordination” to unpack the two ways this appropriate relation might hold (Fine 2007). First, the appropriate relation might hold in virtue of the interlocutors reflecting upon the differing ways they are thinking of the referent, and each establishing on that basis that they are thinking of the same object. Such a case is one of INDIRECT COORDINATION in referential communication. But there may also be cases of DIRECT COORDINATION in referential communication where the differences in perspective are rationally irrelevant in such a way as to mitigate the need for reflective establishment of knowledge of co-reference, and this more intimate connection between interlocutors’ ways of thinking is the notion of TRANSPARENT referential communication that was mentioned at the very beginning of this section, in my overview of the Neo-Fregean argument for the Thesis.

So here is the insight I wish to take from Frege’s passage.

The Neo-Fregean Insight Understanding of a demonstrative holds in virtue of an appropriate relation between the interlocutors’ ways of thinking of the referent, which consists of either direct or indirect coordination.

It will soon become apparent how, in relation to the overall argument I am presenting in this section, the Neo-Fregean must unpack the appropriate relation in the way I have done here, in terms of there being justification of some kind had by the interlocutors for taking each of their ways of thinking of the referent to be of the same object.

⁹I note that this claim seems weaker than the claim about demonstrative understanding generated by traditional Fregean view of demonstrative content, in the paragraph above, that a demonstrative’s content is a single way of thinking of its referent. A single way of thinking of an object, as content, determines that understanding requires that the interlocutors both think of the object in that single way; hence, the traditional Fregean view seems to hold that the appropriate relation must only be that of identity. Nonetheless, it may be that the way of thinking associated with a referential expression, as its sense, is a relatively coarse-grained way of thinking, so that understanding interlocutors may be thinking of the referent in the same way, as well as under more determinate, appropriately related ways of thinking (I take it that this something similar to this proposal is suggested by Dickie & Rattan (2010) when they distinguish between ways of thinking, which are the senses had by (token) demonstratives, and the more fine-grained modes of presentation that interlocutors must think of the referent under). In the next section, however, I reject the Neo-Fregean Insight, so I do not dwell on the question of whether, in addition to the (alleged) appropriately related ways that understanding interlocutors must think of the referent of a demonstrative, there is something, a sense, that the expression has, which somehow binds together or otherwise provides the basis for the various appropriately related ways of thinking; as mentioned in the main text, I set aside questions regarding the semantic content of (token) demonstratives.

Moving on, note that, in his passage, Frege presents differing ways of thinking as being based on differing descriptive information. But when we consider perceptual demonstratives we see that we must not conclude that descriptive ways of thinking are the only ones. Consider the following passage from Moore.

Can we say “that thing” = “that thing at which I am pointing” or “the thing to which this finger points” or “the nearest thing to which this finger points”? No, because the prop[osition] is not understood unless the thing in question is *seen*. (Moore 1962: 158)

To see what Moore is getting at here, imagine a speaker saying something about an object in her perceptual environment, and using the demonstrative “that thing” in doing so. An addressee might employ a description and think, “Oh, she is talking about that thing to which she is currently pointing.” But, intuitively, employing a description in that way does not allow the addressee to understand the speaker’s utterance; instead, the addressee must look at the object to which the speaker is referring with the demonstrative. So here is the upshot of Moore’s passage.

Moore’s Insight Understanding of a perceptual demonstrative requires that the interlocutors each attend to the referent.

Since deployment of a demonstrative concept plausibly requires attention to the relevant object, and not the use of identifying descriptions, we may state Moore’s Insight as the claim that the understanding of a perceptual demonstrative requires that the interlocutors think of the referent by deployment of demonstrative concepts.¹⁰

Moore’s and the Neo-Fregean Insight shape the background Neo-Fregean motivation for the Thesis, that perceptual demonstratives require joint attention for their understanding. To see how they do so, note that Moore’s Insight does not by itself secure the Thesis, for joint attention is more than parallel individual attention. So, the motivation in question for the Thesis involves upgrading Moore’s insight to the stronger claim that joint attention is required, and does so by drawing upon the Neo-Fregean Insight and related ideas regarding the role of transparency of referential communication. More specifically, by drawing upon an analogy between demonstrative communication and demonstrative thought, there is motivation for the claim that transparency is sometimes required for demonstrative understanding, and in

¹⁰Additional motivation and elaboration of Moore’s Insight is in Evans (1982), especially chapter 9, and Campbell (2002), especially the introduction and chapter 1.

particular that this transparency applies to perceptual demonstratives in such a way that joint attention is required for their understanding.

Let me now explain what matters about demonstrative thought, to which demonstrative understanding is to be analogized. Consider the following two scenarios involving Elijah, who is trying to sleep in his cave but for a bothersome fly.

- i Elijah is watching the fly as it roams around the cave. At one point in the episode of tracking, he forms the belief he would express with the sentence “That fly is black.” At a later point in the episode, he forms the belief he would express with the sentence “That fly is large.”
- ii Elijah sees the fly in front of him. He forms the belief he would express with the sentence “That fly is black.” A bit later on, having lost track of the fly, he sees the fly again, but does not realize it is the same fly. At this later time, he forms the belief he would express with the sentence “That fly is large.”

Now suppose that our protagonist were to make the following inference, purely between beliefs, in both scenarios i and ii.

$$\frac{\text{That fly is black} \quad \text{That fly is large}}{\text{There is something that is both black and large}} \quad (\alpha)$$

In both i and ii, the occurrences of “That fly” in the premises are co-referential. However, only in i is his inference justified. In ii, if Elijah is to infer that there is something both black and large, he must realize that that fly, seen at t_1 , is identical to that fly, seen at t_2 . That is, he must add an identity premise to his inference, and the resulting valid inference is as follows.

$$\frac{\begin{array}{l} \text{That fly}_{t_1} \text{ is black} \quad \text{That fly}_{t_2} \text{ is large} \\ \text{That fly}_{t_1} \text{ is that fly}_{t_2} \end{array}}{\text{There is something that is both black and large}} \quad (\beta)$$

One might wonder whether in i Elijah is in fact justified in inferring the relevant conclusion, without the aid of an identity premise; what is wrong with the suggestion that, in both i and ii, he must employ an inference like that in β ? In order to answer this question, let us first get a bit more precise about the notions of coordination in thought.

Two demonstrative beliefs B and B' , which are about the same object, are DIRECTLY COORDINATED just in case the subject has justification for taking B and B' to be about the same object that does not rest upon justification for the belief that the demonstrative ways of thinking

employed in each of B and B' are co-referential. Two demonstrative beliefs B and B' , which are about the same object, are INDIRECTLY COORDINATED just in case the subject's justification for taking B and B' to be about the same objection does rest upon justification for the belief that the demonstrative ways of thinking employed in each of B and B' are co-referential. The inference in α , if justified, is a case where the subject has directly coordinated demonstrative beliefs. If an identity premise must be added, as represented in β , the subject's demonstrative beliefs are indirectly coordinated.

Now, here is Campbell's (1987) argument that there must be directly coordinated beliefs.¹¹ Suppose, for reductio, there is only indirect coordination. Then, for every inference as in α we would have to add an identity premise asserting the co-reference of the demonstratives in the original premises, as in β . However, there would then be the further question of whether the occurrences of the demonstratives in the original premises were co-referential with the occurrences in the identity premise. Since, by supposition, there is no direct coordination, we must add two additional identity premises. We are thus launched on an infinite regress. No justification, however, can be based on an infinite regress. Hence, given that we sometimes are justified in taking two demonstrative beliefs to be about the same object, we must reject the supposition: there must be directly coordinated demonstrative beliefs.

The relations of coordination explain why certain inferences are justified and others not. More generally, their role is to explain how a subject's activity of grouping together beliefs as about the same object is in some cases immediately justified and in others not. In addition, we may talk of demonstrative concepts, which feature beliefs such as those above, as being directly coordinated or not. So the upshot of Campbell's argument is that if demonstrative concepts are ever indirectly coordinated, then there must also be cases in which they are directly coordinated. And, finally, it is plausible that unbroken attention to an object over time plays the role of providing such direct coordination. Hence, we get an answer to the question above of whether, in scenario *i*, Elijah must establish an additional identity premise: he is not required to do so, since his attention to the fly grounds the direct coordination between the demonstrative concepts in the various beliefs he forms.

With this result about demonstrative thought in hand, consider the following argument for the Thesis.

Premise 1 There are cases in which direct coordination is required for demonstrative understanding.

¹¹Similar but less vivid arguments for this conclusion are given by Strawson (1971b) and Evans (1982, 1985).

Premise 2 If there are cases in which direct coordination is required for demonstrative understanding, then such cases include those of perceptual demonstratives.

Premise 3 If direct coordination is required for the understanding of a perceptual demonstrative, then joint attention to the referent is so required.

Conclusion Understanding of a perceptual demonstrative requires joint attention to the referent.

Again, the justification for each premise of this argument comes by analogy with the argument just given about demonstrative concepts. With regard to premise 1, the relevant upshot of the analogy with thought is that if there is ever to be demonstrative understanding, which we may assume there is, then there must be cases of demonstrative understanding that hold in virtue of direct coordination—since, by the Neo-Fregean Insight, demonstrative understanding involves either direct or indirect coordination, and there cannot be only indirect coordination. So, taking a final step, there are cases in which direct coordination is required for demonstrative understanding.

Here is the motivation for premise 2: if there is to be direct coordination required for the understanding of a use of a demonstrative, then—by analogy with demonstrative concepts and their relationship to attention—it is plausible that ones tied to perception are where it is required. And, for premise 3, since demonstrative understanding in ordinary conversation is something two individuals do together, it is plausible that it is not mere parallel individual attention that grounds the direct coordination, as in lone thought; rather, to upgrade Moore's Insight, it is *joint* attention, between the interlocutors, that does so.

2.3 Frege and Felicitous Underspecification

In this section, I respond to the argument just given for the Thesis, that perceptual demonstratives require joint attention for their understanding. My response to the argument rejects the first premise, that direct coordination is sometimes required for demonstrative understanding. Again, the motivation for that premise rests upon the Neo-Fregean Insight, that demonstrative understanding holds in virtue of either direct or indirect coordination—that is, justification of *some kind* for the interlocutors taking their ways of thinking of the relevant object to be co-referential is required. Given the Neo-Fregean Insight, the assumption that demonstrative understanding always consists of indirect coordination leads to a regress

analogous to that given for demonstrative concepts in thought; hence, there must be some demonstrative understanding that holds in virtue of direct coordination. In what follows, I defuse this motivation for the first premise by targeting the Neo-Fregean Insight. I begin, in §2.3.1, by developing a direct challenge to the Insight. Afterwards, in §2.3.2, I address its standard motivation.

2.3.1 Challenging the Neo-Fregean Insight

My challenge to the Neo-Fregean Insight is based on cases of demonstrative understanding where the interlocutors have not settled on one or another of several potential referents. As noticed by Dummett (1973: 74), one can understand an utterance of “That is a book” without settling on whether the demonstrative refers to the particular book token or the book type of which it is a token. This phenomenon has been recently revived by King (2014b, 2017), who labels it “felicitous underspecification.”¹² Here are some vivid cases from King’s work, where the first is an analogue of Dummett’s.

Sports Car At a dealership, John and Mary stop in front of a row of some new model of sports car. The salesman gets into one and drives up to them. Consummed, John exclaims,

(6) That’s a beautiful car!

Intuitively, there can be understanding even if John and Mary do not both fix on either the car type or token.

Package John hears the characteristic, now hourly sound of a delivery truck stopping on his street. He runs out to see his new computer being delivered, in a big brown box. He exclaims to Mary,

(7) It’s arrived!

¹²King argues that this phenomenon arises for many other context-sensitive expressions. See also Blackburn (1988), Richard (2004), Glanzberg (2009), Buchanan (2010), von Fintel & Gillies (2011), MacFarlane (2016), Bowker (2019), Dickie (2020), Szabó (2020), and Charlow (forthcoming) for related insights and discussion. Cases that suggest it even arises for proper names are given by Schiller (2019). In this dissertation, I only focus on demonstratives.

There can be understanding even if John and Mary do not both fix on either the package or its contents.

Tablets John and Mary have ordered from Amazon six tablets of the same make and model. The package arrives and, after opening it and realizing they are all of some other make and model, John mutters,

(8) That's not the tablet I ordered.

There can be understanding even if John and Mary do not both fix on the tablet kind or any one of the tokens.

Co-Pilots John and Mary are watching an airshow where the planes are co-piloted. Noticing one of the planes start smoking, John says,

(9) He's in trouble.

There can be understanding even if John and Mary do not both fix one or the other pilot.

I have given several cases because there might be resistance to recognizing genuine felicitous underspecification. One might even be tempted to explain them away. But there are many such cases; King provides even more. Moreover, these particular cases provide might seem contrived, but this phenomenon—where it does not matter for understanding whether one of several potential referents are settled upon—seems to be a relatively common feature of everyday communication.

What is crucial for present purposes is that in cases of felicitous underspecification there may be understanding without the interlocutors thinking of the same referent, and even without either interlocutor thinking of any object at all. Thus, such cases show that there can be demonstrative understanding without the interlocutors deploying ways of thinking of any object. So these cases provide counterexamples to the Neo-Fregean Insight, since there is understanding in them without an appropriate relation, of any kind whatsoever, between ways of thinking had by the interlocutors.

The problem that these cases raise for the Neo-Fregean Insight can be sharpened by comparing it to a slightly different problem, which concerns demonstrative understanding where the interlocutors are having empty singular thoughts. A SINGULAR THOUGHT is a mental state

whose having of content depends on the existence of some particular object. Suppose that John is walking through the woods and sees a bear in front of him. He forms the belief that that bear is dangerous, as well as the intention to get away from it. These mental states are not about the bear by having descriptive content that the bear uniquely satisfies; the relevant belief is not, for instance, that there is one and only one bear-like thing standing in those bushes, which is dangerous. Such a belief would be false if there is no bear. The relevant belief is that *that bear* is dangerous; it is about the bear in a direct way, and neither true nor false if there is in fact no bear. Accordingly, an EMPTY singular thought is one that has no content, since there is no relevant object. So, to develop our running example, suppose that John and Mary are walking through the woods, and it is now well past dark. John is anxiously looking around, and thinks he sees a bear—though there is no bear, only a misshapen tree behind some foliage that the wind has rustled. He exclaims to Mary, “That bear is about to attack us! Let’s run!” Mary looks and thinks she sees a bear as well. If in a case such as this there were understanding of the uttered demonstrative, then there would be something we might call the “problem of empty communication,” constituted by explaining how it is that there is such understanding, despite the fact that the interlocutors’ respective ways of thinking are empty. Cases that support this problem of empty communication are ones where there seems to something like coordinated hallucinations on behalf of the interlocutors.

It is not clear, however, that there is understanding in such cases. Thus, it is not clear whether there is a genuine problem of empty communication (in contrast, it is intuitively clear that in cases of felicitous underspecification there is understanding). More importantly, even if there were understanding in cases of empty communication, the problem such cases would raise is not the same as that raised by cases of felicitous underspecification. For the problem of empty communication involves the issue of how two empty ways of thinking, between interlocutors, may be appropriately related. But, again, in felicitous underspecification there may be understanding without either interlocutor thinking of any object at all. So the issue is not of how there can be coordination between empty singular thoughts.

Let me conclude this subsection by considering a rejoinder on behalf of the Neo-Fregean Insight. Recall that the Neo-Fregean Insight elaborated the talk of an appropriate relation between ways of thinking in terms of the interlocutors having justification for taking their ways of thinking of some object to be co-referential (that is, to in fact be of the same object). So it might be proposed that the Neo-Fregean Insight can account for the understanding in felicitous underspecification, since it might be that all that is required for demonstrative understanding is that the interlocutors have justification for the belief that theirs and the other’s

ways of thinking are co-referential. In order for an individual to have justification for a belief, it is not necessary for that individual to in fact form that belief, so the fact that the interlocutors do not form such a belief does not mean it cannot be justified for them.¹³ Moreover, a belief can be justified without being true, which allows for the theoretical possibility that in felicitous underspecification the interlocutors do have the relevant justification.

Focusing on the addressee's justification, for simplicity, here is how this proposed account of felicitous underspecification is most naturally elaborated.¹⁴ Again, the relevant belief of the addressee—that theirs and the speaker's ways of thinking of some object are co-referential—is not true. The belief is not true because neither the speaker nor addressee are thinking of any particular object whatsoever (as referent of the demonstrative). The fact that there are no ways of thinking might make the addressee's belief not true by making the relevant belief lack a truth-value, or by making it false; but the proposal that follows is neutral on which of those two alternatives is right.¹⁵ What is crucial is that a belief that is not true is nonetheless justified when the believer is merely unlucky that the belief is not true. Thus, the supporter of the Neo-Fregean Insight should claim that, in demonstrative communication where the speaker and addressee are not thinking of any object, the addressee has merely gotten unlucky. So, for cases of felicitous underspecification, the Neo-Fregean proposal under consideration is that the addressee does have justification for taking theirs and the speaker's way of thinking of some object to be co-referential, since the addressee is merely unlucky that neither they nor the speaker are thinking of any particular object.

In response to this proposal, it must be acknowledged that there is some plausibility to the claim that the addressee is merely unlucky with regard to the speaker: the addressee, in order to be justified in their belief, might not be required to rule out that the speaker is not deploying a way of thinking of some object. But it is plausible that the addressee is required to rule out that they themselves are not deploying a way of thinking, since whether they are deploying a way of thinking is something under their control. That is, it is not plausible that an individual

¹³Here I am relying on the distinction between propositional and doxastic justification, as in the previous chapter (see footnote 6 there), and allowing that the relevant justification be propositional.

¹⁴This elaboration is based on a proposal concerning empty singular thought given by Dickie (2014), which concerns how a single individual can have justified demonstrative beliefs that have no referent.

¹⁵Whether the addressee's relevant belief lacks a truth-value or is false depends on whether it is *de re* or *dicto*, respectively, in attempting to be about the (non-existent) ways of thinking of the object. If the relevant belief is *de re*, then it holds that *those particular ways of thinking had by the speaker and addressee* are about the same object, so lacks a truth-value if there is no such ways of thinking. If the relevant belief is *de dicto*, then it holds that there is unique way of thinking of the object had by the speaker, and another had by the addressee, which are about the same object—and that belief is false in the case at hand. As mentioned in the main text, I take it that the present considerations are neutral with regard to which of these alternatives is correct.

can be merely unlucky that they themselves are not at least attempting to think about some object. Thus, in cases of felicitous underspecification, there is understanding without the addressee having justification for taking theirs and the speaker's ways of thinking to be co-referential. So, to conclude, cases of felicitous underspecification constitute counterexamples to the Neo-Fregean Insight, even the weaker version that merely concerns justification for a belief and does not require any type of appropriate relation between ways of thinking the interlocutors are in fact deploying.

2.3.2 Pathways of Thinking

In the next chapter, I account for the understanding present in felicitous underspecification. Before turning to that task, the main purpose of the present subsection is to discuss the standard motivation for the Neo-Fregean Insight, and briefly outline how it may be undermined on my Gricean approach.

In the previous section, the Neo-Fregean Insight was motivated with a passage from Frege. Let me reintroduce that motivation by considering a case, from Loar (1976), that more vividly illustrates the type of case hit upon in Frege's passage.

Television John and Mary are at home with the television on in the background. They are talking about a man they regularly see on the train during their morning commute. That same man happens to be on the news program playing on the television. Mary notices the man on the television, and utters,

(10) That man [pointing at the TV] is well-dressed.

Enthralled in their previous conversation topic, John does not notice Mary's pointing gesture or what is happening on the television. He assumes she is continuing to talk about the man they see on their commute, and interprets Mary's utterance accordingly. Intuitively, John has not understood Mary, despite the fact that they are thinking of the same object as referent.

In so-called "Loar cases," such as Television, it seems that the lack of understanding is because the interlocutors are not thinking of the relevant object, in this case the man, in appropriately related ways. Hence, Loar cases support the general claim that the holding of such an appropriate relation is required for the understanding of any (deictic) demonstrative. Finally, the

Neo-Fregean Insight unpacks that appropriate relation in terms of justificatory coordination between the interlocutors' ways of thinking of the relevant object.

The Neo-Fregean Insight, I stress, is a claim about what is involved in the understanding of demonstratives. It does not directly concern demonstrative content. But let me say a few words about Loar cases and how they are brought to bear upon the topic of demonstrative content. Loar cases are typically used to argue against the popular Neo-Russellian view, according to which the content of a token demonstrative in context is simply an object, the token demonstrative's referent.¹⁶ The argument against Neo-Russellianism from Loar cases appeals to the principle that a token demonstrative's content determines what is required for its understanding. If Neo-Russellianism is correct, then, in *Television*, the content of the token demonstrative "that man" is simply the particular man himself: the one being interviewed on the television and who the interlocutors know from their morning commute. The only requirement upon understanding that that particular man can determine is that the addressee think of that man, in any which way.¹⁷ So, given the crucial principle, Neo-Russellianism predicts that the addressee thinking of the particular man is sufficient for there to be understanding of the token demonstrative. But, in *Television*, John has not understood Mary's utterance, despite the fact that he does think of the particular man.

Alongside their use in arguing against Neo-Russellianism, Loar cases are typically taken to support a Neo-Fregean view about demonstrative content, according to which a token demonstrative has, as content, a sense: something which determines a limited collection of ways of thinking of a referent. A way of thinking of a referent that is relevant for demonstrative content is something more fine-grained than a referent. Consider, for instance, the different ways of throwing a baseball—or, as famously analogized by Evans (1985), the different ways of dancing. The Neo-Fregean thus explains, in virtue of the token demonstrative's content, how more is required for understanding the utterance than that the addressee think of the particular man: in *Television*, it is claimed, the addressee must think of the particular

¹⁶See Kaplan (1989) for a seminal development of Neo-Russellianism about demonstratives. Some important elaborations of the argument against it, which, as discussed in the next paragraph, typically goes with an argument for a Neo-Fregean view, are Loar (1976), Evans (1982), Récanati (1993), Heck (1995, 2002), Bezuidenhout (1997), Paul (1999), and Dickie & Rattan (2010).

¹⁷The claim in this sentence might seem obscure. So here is the present argument stated more clearly. The crucial principle entails that the requirements upon understanding a token demonstrative supervene on its content in context. Now, compare with *Television* another case, *Television'*, in which all is the same except that Mary in fact keeps talking about the man as seen on their commute, and does not gesture towards the television. In both *Television* and *Television'*, the same man is the content of the uttered demonstrative, but there must be different requirements upon understanding, since John's same activity satisfies them in *Television'* but not in *Television*. Thus, if the crucial principle is correct, Neo-Russellianism is not.

man in a way determined by the sense that is the content of the token demonstrative.

I reject both the Neo-Fregean and Neo-Russellian accounts of demonstrative content. Fully unpacking the reasons against those accounts, as well as my alternative, will take until the very end of this dissertation. In this section, I have already begun along that path by criticizing what the Neo-Fregean says about demonstrative understanding, since, as was just explained, that claim about understanding is taken to support the Neo-Fregean view about content. In the previous subsection, I gave a direct argument against the Neo-Fregean Insight, and in the remainder of the present subsection I suggest that something else is going wrong with Loar cases. Let me stress, however, that it is not as if the present considerations vindicate Neo-Russellianism. Taken together, Loar cases and cases of felicitous underspecification suggest that both more and less is required for demonstrative understanding than what is held by Neo-Russellianism. So these two types of cases are mystifying from the Neo-Russellian point of view.

What I suggest about Loar cases, very programmatically, is that the lack of understanding therein concerns not the way the addressee thinks of the referent, in relation to the speaker's, but instead the way that the addressee *comes to think* of the referent.¹⁸ This suggestion about Loar cases may be properly called "Gricean" because any Gricean account of understanding holds that communicative success involves the addressee recognizing an aspect of the speaker's mental life *in an appropriate way*; and if recognizing that aspect of the speaker's mental involves thinking of an object—as recognizing a referential intention does—then plausibly such an account requires that an addressee come to think of the object in an appropriate way.

In terms of my refined Griceanism, the addressee must recognize the speaker's referential intention—in thinking something such as "Oh, the speaker wants me to think of that object"—by virtue of the joint activity of making sense of one another. But, intuitively, in *Television* the interlocutors are not engaging in that joint activity, since John, the addressee, does not properly consider Mary's demonstration. Thus, the problem in *Television* is that the activity that John is engaging in with Mary, in his coming to think of the relevant object, does not amount to the proper joint activity that it must in order for there to be understanding. So the problem concerns the way that John comes to think of the relevant object, not the nature of the concept by which he is thinking of the object.

Support for this Gricean treatment of Loar cases may be found on the basis of further re-

¹⁸I take it that this general approach, that of distinguishing ways of thinking of objects with ways of coming to think of objects, is the background insight guiding Buchanan's (2014) Gricean treatment of Loar cases. But see Peet (2017) for issues with Buchanan's way of spelling out the approach.

flection upon cases of felicitous underspecification. Reconsider Co-Pilots, a case of felicitous underspecification from the previous subsection, but with a twist.

Modified Co-Pilots The day is hazy, and atmospheric refraction makes one and the same plane appear in two places at once, though John and Mary believe the two hazy images correspond to separate planes. Now, suppose that John points to one image in uttering (9), “He’s in trouble.” Mary, however, mistakenly thinks he is pointing to the other. In this case, there is intuitively misunderstanding, despite the fact that it is not required that Mary thinks of the same object as John, for there can be understanding even if neither John nor Mary think of the plane or either co-pilots as the referent of the demonstrative.

Now, I suggest that the lack of understanding in Modified Co-Pilots is due to the fact that the interlocutors are not engaging in the joint activity of making sense of one another, since Mary, now the addressee, does not properly consider John’s demonstration. But note that the reason there is lack of understanding here seems similar to that in the standard Loar case, Television, despite the fact that here the interlocutors are not thinking of any particular object. Thus, it seems that the basis for the lack of understanding in Television is not that there fails to hold an appropriate relation between interlocutors’ ways of thinking of the relevant object. Rather, since there seems to be a similar failure of understanding in both Television and Modified Co-Pilots, the lack of understanding in Television is due to an activity that is also present in Modified Co-Pilots, which I propose is the addressee’s failed attempt at making sense of the speaker. In Television, where the addressee does think of a particular object, that activity is part of how the addressee comes to think of the object. Thus, considering further cases of felicitous underspecification shows that they can be modified in such a way as to motivate the Gricean treatment I propose for Loar cases.

Finally, I wish to address a loose end of the present chapter. In the previous section, I mentioned a passage from Moore and a corresponding Moorean Insight, that understanding of a perceptual demonstrative—a deictic demonstrative used in ordinary conversation to talk about an object perceptually available to the interlocutors—requires that the addressee attend to the referent. This insight was motivated by the following case, suggested in the passage from Moore. The speaker is looking at an object and points at it in using a demonstrative. But the addressee responds by thinking something along the lines of “Oh, the speaker wants me to look at the object at which she is pointing.” That is, the speaker uses descriptive material in specifying an object in order to recognize the speaker’s intention—instead of looking at

the object to do so.

The case motivating Moore's Insight is a Loar case in that the interlocutors are thinking of the same object, yet there is intuitively no understanding. So let me explain how the present Gricean approach to Loar cases accounts for it. The addressee's use of descriptive material in specifying the object suggests a disengaged, third-personal way of making sense of the speaker; the addressee seems not to be engaging in the joint activity of making sense of one another, but instead merely reflecting observationally upon what the speaker is doing. Imagine that you are speaking to someone and you make a demonstration, but they do nothing overt to consider it. In such a case the problem is not only that the addressee seems not doing their part in making sense of what you are doing, but, moreover, the addressee's action—that of not responding in any way to your gestures—will also be unintelligible to you, for if someone does nothing to respond to your actions that are overtly directed towards them, they will not seem intelligible to you as a fellow intentional agent. Hence, in such a case there is lack of mutual intelligibility, and hence, on my refined Griceanism, a lack of understanding of the utterance.

Suppose, in contrast, the addressee follows your demonstration, looks at what you are pointing at, and thinks "Oh, he wants me to look at that object." Such a case is one where an intention is attributed on the basis of the joint activity of achieving mutual intelligibility, and there is intuitively demonstrative understanding present. But, I suggest, there is also that joint activity, and intuitively understanding, if the addressee follows the gesture in the same way, and thinks "Oh, he wants me to look at the object at which he's pointing." So, I conclude, the intuitive lack of understanding in the case suggested in Moore's passage does not support Moore's Insight, nor more generally the Neo-Fregean Insight, and instead can be accounted for on my refined Griceanism.

2.4 Looking Ahead

The main goal of this chapter was to defuse existing motivation for the Thesis, that joint attention is required for the understanding of perceptual demonstratives. I began, in §2.1, with Campbell's (2002) explicit argument for it. Next, in §2.2, I developed further motivation for the Thesis, which was found in the background Neo-Fregean tradition. One crucial part of that background motivation I called the Neo-Fregean Insight, that demonstrative understanding involves justificatory coordination between the interlocutors' ways of thinking of the referent. So I focused, in the final §2.3, on criticizing that Insight, both by directly arguing against it in

§2.3.1, on the basis of cases of felicitous underspecification, and undermining its motivation in §2.3.2, by suggesting an alternative, Gricean account of Loar cases. The Gricean suggestion about Loar cases requires further development. But I leave that task to future work. In the next chapter, I turn to accounting for felicitous underspecification. But, before doing so, let me mention a final issue that arises from the present chapter's discussion of the relative priority of demonstrative understanding and joint attention.

Recall from the introduction of this chapter that the guiding purpose here in defusing motivation for the thesis arose from the account of joint attention in the first chapter. On the communicative conception of joint attention there defended, communication helps bring about joint attention—and there I supported a simple model of how communication helps bring about joint attention. Recall the distinction, drawn from the psychological literature, between primary and secondary intersubjectivity (Rochat & Striano 1999, Adamson & Russell 1999, Hobson 2002, Striano & Bertin 2005). Throughout the first year of life, infants engage in direct one-to-one communication with others, which is primary intersubjectivity. Then, around the first birthday, joint attention begins, which is the dawning of secondary intersubjectivity, which itself involves interacting with others in relation to the world. Crucially, the social engagement in primary intersubjectivity plays a role in establishing secondary intersubjectivity.¹⁹ I thus propose that the way in which primary intersubjectivity plays that role with regard to joint attention is that joint attention crucially involves the interlocutors making sense of one another. Furthermore, I propose that successful communication involves the interlocutors making sense of one another, which is to say that communication involves primary intersubjectivity.

Thus, since communication with perceptual demonstratives is a form of communication, communication with perceptual demonstratives is one way that communication can help sustain joint attention. But one might think that there is a special connection between communication with perceptual demonstratives and joint attention, in contrast with other forms of communication—even if the Thesis, that joint attention is required for the understanding of perceptual demonstratives, is rejected. Children begin using perceptual demonstratives at a very young age: a child might point and exclaim “That!,” or utter “That truck is my favourite.” As covered in the previous chapter, joint attention also begins at a very young age, and many such episodes of joint attention are initiated by uses of perceptual demonstratives. I do not have a satisfying account of what might be the special relationship between the use of percep-

¹⁹For recent empirical support for this claim see R  ther & Liszkowski (2020), who found that among infants a higher degree of primary social interaction predicted success on secondary subjectivity tasks, but not vice-versa.

tual demonstratives and joint attention, beyond the fact that from understanding of a perceptual demonstrative it is a short step to achieving joint attention: the communicative condition is already in place, and the communication concerns an object that is perceptually available to both the speaker and addressee; all that must be added for joint attention are mutually responsible experiences of the object. Thus, my treatments of joint attention and communication plausibly explain how understanding of perceptual demonstratives often leads to joint attention, but I leave the task of making sense of a deeper connection between the two phenomena to future work.

Chapter 3

Underspecification in Demonstrative Communication

This chapter addresses the phenomenon of felicitous underspecification, as it arises for assertions containing deictic demonstratives. Here are some instances of the phenomenon, from King (2014a, 2017), which were introduced in §2.3.1 of the previous chapter.

Sports Car At a car dealership, John and Mary stop in front of a row of some new model of sports car. The salesman gets into one and drives up to them. Consumed, John exclaims,

- (1) That's a beautiful car!

Intuitively, there can be understanding even if John and Mary do not both fix on either the car type or token.

Package John hears the characteristic, now hourly sound of a delivery truck stopping on his street. He runs out to see his new computer being delivered, in a big brown box. He exclaims to Mary,

- (2) It's arrived!

There can be understanding even if John and Mary do not both fix on either the package or its contents.

Tablets John and Mary have ordered from Amazon six tablets of the same make and model.

The package arrives and, after opening it and realizing they are all of some other make and model, John mutters,

(3) That's not the tablet I ordered.

There can be understanding even if John and Mary do not both fix on the tablet kind or any one of the tokens.

Co-Pilots John and Mary are watching an airshow where the planes are co-piloted. Noticing one of the planes start smoking, John says,

(4) He's in trouble.

There can be understanding even if John and Mary do not both fix one or the other pilot.

In general, in a case of felicitous underspecification there is understanding despite neither of the interlocutors settling on one or another particular referent.

Let me introduce a few theoretical notions in order to draw out what I take to be the central puzzle regarding felicitous underspecification. The Gricean approach to communication that has been defended in the previous two chapters posits certain speaker intentions, and these intentions plausibly determine what a speaker means in making an utterance. So I suggest that *WHAT IS MEANT* by an assertion is a particular, determinate claim given by what the speaker intends to say with it. But one may also ask about what the interlocutors take the speaker to mean to say. So I suggest that *WHAT IS SAID* by an assertion is a particular, determinate claim given by what it is common ground that the speaker means by their utterance. The notion of *COMMON GROUND* will be elaborated as this chapter unfolds, but for preliminary purposes it can be taken as the information mutually assumed by the interlocutors for the purposes of conversation.

In cases of felicitous underspecification, the speaker does not intend to say anything. In *Sports Car*, for instance, the speaker is not intending to say of either the car type or token that it is beautiful. So there is nothing that is meant. And in such cases neither is it common ground that there is a certain claim that the speaker has intended to make. In *Sports Car*, for instance, the common ground does not settle whether the speaker intends to say of either the

car or token that it is beautiful. So there is nothing that is said. Yet, in such cases, there is understanding. So the central puzzle of felicitous underspecification is of how there can be understanding where there is nothing meant or said.

Furthermore, a satisfactory answer to this puzzle must also explain the limits of when such understanding is possible. Consider the following case.

Cluttered Shop John and Mary are in a cluttered shop, surrounded by bric-à-brac. John utters the following, with a vague gesture and no thought about what he means to speak about.

(5) #That's a fine red one.

Here there is nothing meant or said, and also no possibility of understanding.

I use the hashtag “#” in this chapter to indicate that the utterance is unsuccessful, or infelicitous, in such a way that prohibits the possibility of understanding. So Cluttered Shop is a case of *infelicitous* underspecification, and the correct account of how there may be understanding without anything meant or said must not predict that there may be understanding in such a case.

Now let me say a little bit more in general about the notions of what is meant and what is said, in order to provide a preliminary outline of my solution to the puzzle of felicitous underspecification. What is meant and what is said by an utterance may often line up, since it may often be that the common ground accurately singles out what the speaker intends to say. But the notions may also come apart, in the following two ways. First, it may be that the common ground does not entail that there is something in particular that the speaker means, since there might be multiple claims about what the speaker means that are consistent with what is mutually assumed by the interlocutors. Second, and more radically, it may be that the true claim about what the speaker means—if anything at all—is not even consistent with the common ground, since what is mutually assumed need not be consistent with the way things actually are. And here is how the possibility of what is meant and what is said coming apart in these two ways sheds some initial light on how the present puzzle may be resolved. In cases of felicitous underspecification, the common ground is consistent with a number of hypotheses about what the speaker meant, and is not consistent with the truth that the speaker meant nothing. Since there are a number of eligible hypotheses, there is also nothing that is said—but, as will be elaborated, the eligible hypotheses may be related in such a way that something determinate can be done on their basis.

In this chapter, I develop the proposal just outlined. In §3.1, I begin by showing how a collection of well-known pragmatic principles, from Stalnaker ([1978] 1999), are satisfied in cases of felicitous underspecification—and also that they rule out cases of infelicitous specification. In the sections that follow, §3.2 and §3.3, I elaborate and defend the proposal that there is understanding of an assertion just in case Stalnaker’s principles are satisfied. I conclude by comparing the resulting view of understanding to my Gricean account of understanding from the previous chapters.

3.1 Felicitous Underspecification and its Limits

In a seminal paper, Stalnaker ([1978] 1999) proposed three “essential conditions of rational communication” (88). All three principles are constraints on what the common ground of a conversation must be, given a speaker’s assertion. As already mentioned, I take for preliminary purposes the common ground to be the information mutually taken for granted by the interlocutors for the purposes of the conversation. This initial characterization makes it clear that a common ground determines a *CONTEXT SET*: the set of maximally specific ways that the world might be—i.e. possible worlds—that are consistent with the common ground.

My focus in this section is on how the principles are satisfied in cases of felicitous underspecification, and how they are not in cases of infelicitous underspecification. One of the principles, called “Uniformity,” is of central concern.¹ It holds roughly that any two hypotheses about what the speaker means, consistent with the context set, must lead to the same update on the context set. Here is the principle stated more formally.²

Uniformity Given a speaker’s utterance, each world w in the context set C must determine a proposition as what the speaker means, in w , and any two such determined propositions P and P' must be *UPDATE-EQUIVALENT* on C , in that $C \cap P = C \cap P'$.

¹The terminology of “Uniformity” is from Hawthorne & Magidor (2009), but note that the principle I am endorsing here corresponds to what they call “Weak Uniformity,” as it involves a qualification that Stalnaker ([1978] 1999) mentions in footnote 13. I note also that Hawthorne & Magidor criticize this principle and its motivation, but for responses see Stalnaker (2009), Almotahari & Glick (2010), and Kirk-Giannini (2018).

²Here is a brief explanation of the standard notation and terminology I here employ. The variables “ w ,” “ w' ,” ... range over maximally specific ways that the world might be, i.e. possible worlds. The variables “ P ,” “ P' ,” ... and other capital letters range over less specific ways that the world might be, i.e. sets of possible worlds. I talk of sets of possible worlds as propositions, but for purposes of this chapter I need only assume that a proposition, strictly speaking, determines a set of possible worlds. Finally, I use the standard set-theoretic operation of intersection, denoted by “ \cap ,” which takes two propositions and determines what they have in common: $P \cap P'$ is the set of worlds that are in both P and P' .

Crucially, Uniformity allows there to be indeterminacy in the context set as to what the speaker means, as long as there is a determinate suggestion as to how the context set may be updated—where the determinate suggestion is guaranteed by update-equivalence.

Uniformity may be equivalently stated as the requirement that every hypothesis about what the speaker means, consistent with the context set, leads to the same update on the context set. That equivalence holds because every one of a set of propositions has the same intersection with a given context set just in case every pair of propositions in the set of propositions is update-equivalent on the context set. So for simplicity I will speak of not only pairs of propositions being update-equivalent (on context sets), as the notion of was introduced in Uniformity, but also more generally sets of propositions being update-equivalent.

Here is how the four cases of felicitous underspecification from the introduction satisfy Uniformity. In Sports Car, the speaker, John, has done enough, but only enough, to help make it common ground that he is either saying that the car kind is beautiful, or that the car token of that kind, in front of them, is beautiful. But it is a manifestly plausible mutual assumption that a kind is beautiful just in case its instances are. See figure 3.1 for an illustration of how Uniformity holds in Sports Car. In Package, the common ground has it that John is either




Figure 3.1: The solid rectangle represents the common ground. The dotted rectangles represent the different claims the speaker could be making. T represents the worlds where the car token is beautiful (\bar{T} where it is not), and K the worlds where the car kind is.

saying that the package arrived, or its contents. But it is a reasonable joint assumption that a package arrives just in case its contents do. In Tablets, the common ground does not settle whether it is this or that particular tablet that John is speaking of, or even the kind, when he says that it is not the one he ordered. But it is plausibly part of the common ground that all the tablet tokens are of the same kind, so one is of the wrong kind just in case the others are, as well as the kind itself. In Co-Pilots, the common ground holds that John may be saying either that one co-pilot or the other is in trouble. But it is a plausible assumption that, when

a plane is threatening to go down, one co-pilot is in trouble just in case the other is.

Here are Stalnaker's two other principles, in addition to Uniformity.

Informativeness Given a speaker's utterance, each world w in the context set C must determine a proposition P as what the speaker intends to say, in w , and any such determined P is true in some but not all of the worlds in C .

Contentfulness Given a speaker's utterance, each world w in the context set C must determine a proposition P as what the speaker intends to say, in w , and any such determined P must have a truth-value in every world of C .

Together, the three principles guarantee that, given a speaker's assertion, the context set itself determines that there is a single way of updating that same context set (Uniformity), where that update is a genuine reduction of the context set (Informativeness), and where it is clear, for each world in the context, whether it survives that reduction or not (Contentfulness).³

It should be apparent that the two principles other than Uniformity are satisfied in the cases under discussion. In each of the cases, any proposition the speaker could be intending to say, given the context set, is one that is plausibly not already taken for granted by the interlocutors. So Informativeness is satisfied. And in each of the cases no proposition the speaker could be intending to say, given the context set, is one that plausibly has truth-value gaps. So Contentfulness is satisfied.

Thus, based on the cases under discussion, I propose in general that in cases of felicitous underspecification Stalnaker's three principles are satisfied. Now I argue that Stalnaker's principles are not satisfied in cases of *infelicitous* underspecification. In *Cluttered Shop*, again from the introduction, it is not a reasonable joint assumption that one of the red things in the shop is fine just in case all the other ones are. So my focus remains on the principle Uniformity, which, again, requires that the every hypothesis about what the speaker means *that is consistent with the context set* leads to the same update *on the context set*. More importantly than accounting for the particular case of *Cluttered Shop*, Uniformity places a general limit on the

³Following Stalnaker ([1978] 1999: 89–90), my talk here of there being a *clear* way of updating the context set is meant to recapitulate Contentfulness. Given Contentfulness, it must not only be that each of the propositions that could be meant are true in the same worlds of the context set—and thus provide a determinate update—but it must also be that none of the propositions that could be meant lack a truth-value in any world of the context set. If one of the propositions did lack a truth-value in a world of the context set, then the question arises of whether that world really should not survive the update, since the world does not make the proposition false. Hence, it would not be clear whether that world should survive the update or not. I wish to note, however, that the considerations that follow in this chapter do not crucially upon Contentfulness in the way that they rely upon the other two principles; hence, it is fine if Contentfulness cannot be ultimately motivated.

possibility of felicitous underspecification, given the way that it does not allow aspects of conversational context, other than common ground, to help make underspecification felicitous. Let me explain.

King (2017) proposes that contextual goals may play a role in allowing for felicitous underspecification. His discussion is brief and his proposal highly schematic, but we may refine his appeal to contextual goals by looking to Roberts's influential identification of discourse and domain goals (Roberts 2005a, 2012a,b). DISCOURSE GOALS provide the topics of conversations, and play a core role in conversation's function of enabling joint enquiry. Roberts takes discourse goals to be questions under discussion (QUDs). Accordingly, each context has a QUD, which can be either implicit or given explicitly. A QUD—for instance, what is John's favourite cheese?—structures the joint inquiry endemic to conversation. On the other hand, DOMAIN GOALS are practical interests of interlocutors, which they have agreed to work together to satisfy through their conversation. They are particular things that individuals want out of the world—our “domain.” They also are either implicit or explicit. Suppose, for instance, that John is lost in a foreign city and needs to find the train station, so strikes up a conversation with a stranger. More examples of both of these kinds of goals are given below alongside pertinent data.

I now argue that appealing to either domain or discourse goals overgenerates felicitous underspecification. Underspecific cases are in fact more constrained in their felicity than they would be if contextual goals played a role in permitting it. Consider first domain goals, with the following case.

Hammer A mechanic and his assistant are working together to repair a car. The mechanic gestures vaguely at his table of tools, without looking, and says to his assistant,

(6) #I need that tool now.

This utterance is infelicitous, even if the mechanic is at a stage in the repair such that a certain hammer is required.

The mechanic's vague demonstration makes it compatible with the common ground that he may be intending to speak about one or the other of several tools on the table. Given, however, the mechanic and assistant's coordination on the practical goal they are trying to achieve, there is a specific tool on the table such that the mechanic needing it is uniquely conducive to that goal. So it seems that domain goals could play a role in narrowing down the hypotheses

about what the speaker is intending to say, in this case to a single one. Yet the utterance in Hammer is infelicitous.⁴ And Uniformity is violated: the different hypotheses suggested by the demonstration—that the mechanic needs this or that tool on the table—are not update-equivalent, since it is not a plausible common assumption that the mechanic needs one tool just in case he needs another. Thus, domain goals do not play a role in narrowing down the potential hypotheses about what the speaker is intending to say. See figure 3.2 for an illustration of how update-equivalence is not satisfied in this case.




Figure 3.2: T represents worlds where one certain tool is needed by the mechanic, and T' represents worlds where another is. For this diagram, I have assumed, for simplicity, there are only two relevant tools.

Taking a step back, it is theoretically unclear why domain goals should even be expected to aid with felicity. They are simply interests the interlocutors happen to have, and, though commonly accepted, are not integral to conversation as such. QUDs, on the other hand, as discourse goals do have a more intimate connection to the nature of conversation. Hence, it is more plausible that they may help refine King’s appeal to contextual goals. Yet Dickie (2020) shows, with the following cases, that QUDs do not play this role.

Good Student John and Mary are professors, and Mary visits John’s graduate seminar. John has previously mentioned to her that even one good student can make a seminar worthwhile, to which she has expressed agreement. The seminar that Mary visits ends up going terribly, with much awkward silence and off topic comments. As the students

⁴A very similar case to Hammer is one where it is common ground that the mechanic needs a certain hammer at the point in the repair that he is at. In such a case, it is plausible that it also common ground that the mechanic intends to say that he needs that hammer now with (6). My proposal holds that in such a case there is felicity, and I think that lines up with intuition. Moreover, it is also consistent with my proposal that in such a case the domain goals operate in virtue of how they influence the common ground. The claim I am making in this paragraph, on the basis of Hammer, is that domain goals do not by themselves allow for felicitous underspecification by narrowing down the relevant hypotheses about what the speaker means.

are leaving, Mary asks John whether this seminar, on the whole, is worthwhile. In response, he gestures towards the corner of the seminar table, where several students sat taking notes the whole time. He utters,

(7) #She's a really good student—it's a pity she didn't say anything.

The relevant QUD here is explicit: is this seminar worthwhile? John's answer, despite the fact that it is underspecific, provides a clear answer to that question (Yes). Yet the utterance is infelicitous.

Miser John is complaining to Mary that food has become so expensive since the recession. She responds by accusing John of being a “miser” who's always thought everything is overpriced. John retorts as follows, while gesturing towards a shelf behind Mary, which contains dozens of knickknacks, but withdrawing his arm quickly before she turns to look:

(8) Nonsense! #I think that's cheap.

The relevant QUD here, though implicit, is plausibly whether John thinks anything cheap. And the utterance, despite the fact that it is underspecific, provides a clear answer to that question (No). Yet there is infelicity.⁵

In both these cases, Uniformity is violated. In Good Student, John could be speaking of one or the other student, according to the common ground. But it is not a plausible joint assumption that one student is good just in case the other is. In Miser, John could be speaking about one or another knickknack, according to the common ground. But it is not a plausible joint assumption that one knickknack is cheap just in case another is.

Finally, it must be said that Dickie takes her cases to do something very different than provide support for Uniformity—a point to which I return in the next section.

3.2 Understanding and the Commonplace Effect

Recall from the introduction that there is understanding in cases of felicitous underspecification, despite there being nothing that is meant or said. Thus, such cases raise the puzzle of

⁵I have modified the case slightly from how Dickie originally presents it in order to make the utterance more directly responsive to the question under discussion. With my modification, infelicity intuitively remains.

how there can be understanding while there is nothing that is meant or said. I also proposed that fully answering that puzzle involves accounting for the limits of felicitous underspecification. But we also saw in the previous section that the satisfaction of Stalnaker's principles is necessary and sufficient for the felicity of underspecification. Thus, over this section and the next, my plan is to resolve the puzzle of felicitous underspecification by explaining how it is that the satisfaction of Stalnaker's principles is necessary and sufficient for understanding. In this section, I begin by elaborating the connection between Stalnaker's principles and understanding.

To begin, consider how Stalnaker provides background motivation for his principles.

How does an assertion change the context? There are two ways, the second of which, I will suggest, should be an essential component of the analysis of assertion. I will mention the first just to set it apart from the second: The fact that a speaker is speaking, saying the words he is saying in the way he is saying them, is a fact that is usually accessible to everyone present. Such observed facts can be expected to change the presumed common background knowledge of the speaker and his audience in the same way that any obviously observable change in the physical surroundings of the conversation will change the presumed common knowledge. (86)

So Stalnaker begins by distinguishing two ways in which an assertion may influence the common ground. The first, "commonplace" effect—as he later calls it—is the effect an assertion has on the common ground in virtue of it being a mutually observed event. The addressee, for instance, need not accept the assertion in order for the commonplace effect to take place. Stalnaker, however, sets aside the commonplace effect, and focuses instead on what he calls the "essential" effect of an assertion.

Once the context is adjusted to accommodate the [commonplace effect], how does the CONTENT of an assertion alter the context? My suggestion is a very simple one: To make an assertion is to reduce the context set in a particular way, provided that there are no objections from the other participants in the conversation. The particular way in which the context set is reduced is that all of the possible situations incompatible with what is said are eliminated. To put it a slightly different way, the essential effect of an assertion is to change the presuppositions of the participants in the conversation by adding the content of what is asserted to what is presupposed. This effect is avoided only if the assertion is rejected. (86)

The view here about the essential effect of assertion may be glossed as the traditional Stalnakerian view that assertion is fundamentally about sharing pieces of information about the way the world is.

Now we may return to Dickie's cases, *Good Student* and *Miser*, in order to wrap up a loose end from the previous section, and also ultimately point the way to my elaboration of the connection between understanding and Stalnaker's principles. Dickie takes her cases to show that this traditional Stalnakerian view about assertion is incorrect, and her criticism goes roughly as follows. Given that traditional view, the possibility of felicitous underspecification is mysterious, since in such cases there is not a determinate piece of information presented by the speaker. So, she argues, such a proposal is naturally elaborated—broadly as King does—by appeal to QUDs: assertion is about sharing information determinate enough given our discourse interests. But, as her cases show, this natural elaboration overgenerates felicity.

On the basis of her criticism, Dickie—and Szabó (2020) following in her stead—outline alternative pragmatic views. She suggests, very roughly, that linguistic communication is about establishing mutual rational focus on particular objects that inhabit our world. These joint relations of epistemic focus on particular objects are more fundamental than pieces of information about the way the world is as a whole. Since in using a demonstrative a speaker is attempting to engage in that joint focus with her audience, the speakers' actions in *Good Student* and *Miser* are inappropriate because they are not doing enough to initiate that joint activity. But there is no such problem in King's felicitous cases, since such an activity has been initiated, despite the fact that it does not decide between one among several candidates.

Szabó's suggestion is less radical, as he proposes simply adding the requirement that speakers make themselves intelligible alongside the goal of offering information that addresses the QUD. Accordingly, the problem with the speakers' action in *Good Student* and *Miser* is that they have not appropriately justified their answers. Szabó considers, in analogy, the oddness of responding to someone's request for help with a mere "No." For an even more direct comparison, suppose that John in *Miser* were to retort with "No, I think something is cheap." There is something that curt answer lacks. Similarly, John's response (8) provides the same answer—that he thinks something is cheap—but, since it is underspecific, it also does not provide a determinate reason.

These proposals are promising, and one may ultimately be correct. However, I do not dwell on them—However, I do not dwell on them here—though I do return to some aspects of Szabó's view in the next section, since it bears some similarity to what I propose there.⁶

⁶It should also be noted that Buchanan (2010) and Bowker (2019) provide treatments of phenomena highly

Dickie is certainly correct that the phenomenon of felicitous underspecification and its limits poses a challenge to the traditional Stalnakerian view that assertion is fundamentally about sharing pieces of information about the world. Nonetheless, as the previous section should have made clear, the phenomenon supports the three principles that Stalnaker proposes.

Moreover, the phenomenon also reveals that the traditional Stalnakerian view does not provide support for the three principles. Recall, from the introduction, that what is meant by an assertion is a particular, determinate claim that the speaker intends to say; and what is said is what is meant according to the common ground. The piece of information to be shared, according to the traditional Stalnakerian view, should be identified with either what is meant or what is said. But cases of felicitous underspecification involve perfectly successful utterances, according to the principles, but in which nothing is meant or said. Here is why. First, the three principles may be satisfied when there is nothing that is meant because they do not require that the common ground be consistent with the truth about what the speaker meant. In general, there should be no requirement that the common ground reflect the truth: it is possible that what the interlocutors take for granted, for the purposes of conversation, does not reflect how things actually are. Hence, the truth that the speaker means nothing need not be reflected in the common ground. Second, here is how the principles may all be satisfied while there is nothing that is said. The key principle, Uniformity, allows the common ground to be consistent with multiple hypotheses about what the speaker means, as long as there is update-equivalence among the eligible hypotheses. The other two principles also allow there to be multiple eligible hypotheses about what is meant, as long as each is true in some but not all worlds of the context set (Informativeness) and false in all the others (Contentfulness).

So the plan for the rest of this chapter is to find an alternative theoretical basis for Stalnaker's principles, in order to find a deeper explanation of felicitous underspecification and its limits. I set aside the issue of what the nature of the essential effect is—though note that the essential effect of an assertion is what takes place when the assertion is accepted by the addressee. I begin by focusing instead on the commonplace effect, which, recall, is the effect an assertion has on the common ground, preceding the essential effect. Stalnaker sets aside the commonplace effect in the course of motivating his principles, but that is misleading because his principles are in fact constraints on what the commonplace effect of an assertion must bring about; the principles state requirements that the context set must satisfy, given that an

related to the that of felicitous underspecification under discussion in this chapter. They are concerned with expressions other than demonstratives—though it should be noted that I have drawn upon a notion of update-equivalence in a similar way as Bowker does. They are both working in the Gricean tradition, as I am, though I set aside discussion of their work, given space constraints.

assertion has taken place—which is to say *given that the context set has been influenced by the commonplace effect of that assertion*. Together, the principles require that the commonplace effect of an assertion bring it about that the context set is consistent with only update-equivalent hypotheses about what the speaker meant (Uniformity), where the determinate update these hypotheses genuinely reduces the context set (Informativeness), and where none of the hypotheses lacks a truth-value in any of the worlds of the context set (Contentfulness). So, if the commonplace effect of an assertion satisfies the three principles, there is a clear way to update the context set on the basis of that assertion.

Thus, since the principles together require that the commonplace effect brings it about that there is a clear way to update the context set, what the commonplace effect is supposed to do is provide the basis for what is to happen if the assertion is accepted. But the notion of understanding an assertion corresponds precisely to that which provides the basis for its acceptance. Hence, the three principles are requirements upon understanding. As mentioned at the beginning of the previous section, Stalnaker describes his principles as “essential conditions of rational communication.” In contrast, the present considerations reveal that the principles are conditions in an account of understanding. Thus, my proposal concerning the connection between the principles and understanding is as follows.

Understanding and Stalnaker’s Principles There is understanding of an assertion just in case the commonplace effect of that assertion brings it about that the context set satisfies the three Stalnakerian principles.

In the next section, I theoretically motivate this connection between understanding and the principles, thereby completing my account of felicitous underspecification and its limits.

3.3 The Goal of Conversation

In the previous chapters, I developed an account of understanding by drawing upon a Gricean tradition that originates from Grice’s work on speaker meaning (Grice [1957] 1989, [1969] 1989). That tradition focuses on the intentional structure of a speaker’s communicative act. In cases of felicitous underspecification, however, there is a disquieting paucity of mental activity. Crucially, there is nothing that the speaker intends to say. Thus it is unclear how that tradition can help with the present considerations. Fortunately, there is another strand of the Gricean tradition, beginning with Grice’s ([1975] 1989) proposal that conversation is a cooperative enterprise. I wish to draw now upon this other aspect of Griceanism.

An immediate issue for the proposal that conversation is cooperative concerns what the relevant sense of “cooperative” could be. Individuals may converse despite having very different, or even opposed, reasons for doing so. Imagine, for instance, a trial in which the accused takes the stand and is cross-examined by the prosecution, or the vicious back-and-forth of a presidential debate. So “cooperative” must be meant in a weak and perhaps unfamiliar sense. But now consider engagement in joint activity: doing something together. As noted by Bratman (2014), joint activity always has a joint goal, despite the fact that participants may have ulterior motives. Imagine two roommates lifting a heavy chair together up the stairs to their apartment, each with the hope of being the one who gets to use it the most. The joint activity may even be explicitly competitive, as in a game of chess. Thus, what I take to be correct about Grice’s proposal about cooperativity is the claim that conversation is a joint activity, since that claim provides a plausible account of the way in which conversation is cooperative.

But the claim that conversation is a joint activity demands further elaboration. It may be that conversational episodes of communication are cases of joint action, but that the joint goal varies, depending on interests the interlocutors happen to have. Again, there are all sorts of reasons we converse—think of all the reasons why we might do anything. But I maintain that there is a purpose with a special status present in all episodes of communication. In fact, the claim that the transmission of knowledge fulfills this role has received widespread support, and may be fairly called the orthodoxy.⁷ But I propose that the goal of conversation is *MUTUAL INTELLIGIBILITY*: making sense of one another, as fellow agents, in virtue of one another’s actions. Let me now elaborate and defend this proposal by contrasting it with the orthodoxy.

To begin, it is worth getting clear on what it would mean for something to be the basic purpose of conversation. There is intuitive content to talk of the basic purpose, or the goal, of conversation, but, again, how do we know there really is such a thing, instead of only the particular practical interests interlocutors happen to have? There is a clue as to how to proceed in Szabó’s (2020) recent defence of the knowledge-centric status quo.

⁷See Evans (1982), Heck (1995), Williamson (2000), and Szabó (2020) for expressions of the orthodoxy. The latter argues that it should be attributed to Grice. But note that my claim that making sense of each other is the basic purpose of conversation is not novel, as other authors have hit upon it. According to Evans & McDowell (1976), “[w]hen we interpret a man’s language we are making him intelligible to us” (xxiii). This thought is further developed by McDowell (1998b,c), who claims that the basic constraint on theory of meaning is that it makes another’s action in speech intelligible as rational action. But note that, on my view, making sense of each other is at home in pragmatics, not semantics (also, some parts of McDowell (1998b) seem to support the knowledge-based orthodoxy). Heck (2002) claims that the basic purpose of conversation is rational engagement, which is when the interlocutors’ cognitive lives are brought “to bear upon” each other (16).

Even if the parties are ultimately interested in something else (for instance, they want to coordinate some joint action, get better acquainted, or simply avoid boredom), these goals are *usually* achieved, in part, by means of sharing (that is, making common) some private knowledge. (62, emphasis added)

From this passage, we get the following line of thought: interlocutors may bring this or that purpose to the table in conversation; but the basic purpose is the purpose by means of which any other purpose brought to the table is achieved...*usually*.

The “usually” qualification is crucial for Szabó’s motivation for the claim that knowledge-transfer is the central goal. For, as he notes elsewhere on the same page, the orthodoxy faces obvious counterexamples: pretence and coercion. Furthermore, knowledge-transmission is at home with the philosopher’s traditional focus on the assertion of declarative sentences, but not the myriad other uses of language. Yet we should avoid making the “usually” qualification because it leaves the matter at an unsatisfying level. Dropping that qualification leads to a constraint on basic purposes that is not satisfied by knowledge-transfer, but the resulting constraint also seems very strong. Is there really any candidate purpose that is always the ultimate means in conversation?

As hit upon in the previous section, Szabó (2020: §V) himself proposes that something like mutual intelligibility plays a special role in conversation. He says, however, that it “is not a part of the goal of conversation—it is a norm participants conform to as a means to achieve [the goal of knowledge-sharing]” (82). But here he seems to be misapplying the very principle to which he appealed in arguing that knowledge-transfer is the goal. Precisely because intelligibility is a means to sharing knowledge, the latter is not basic. We often do use language to share knowledge. Yet in such cases the transmission of knowledge is achieved by means of making sense of each other’s linguistic action. Furthermore, making sense of each other is the means for all the other things we do with language in conversation: asking questions, making commands, passing the time with the idle expression of preference or belief, and so on. First and foremost, the understanding underlying all these uses involves making sense of each other’s linguistic action. Hence, it is intuitively plausible that mutual intelligibility satisfies the strong, unqualified constraint on what it means for something to be the basic purpose of conversation. So, I conclude that, no matter what else we are attempting to achieve in talking to each other, it is ultimately done in virtue of us making sense of each other.

Let me defuse a couple potential counterexamples to the conclusion just reached. Imagine that John and Mary have gotten into a heated argument. Mary has become so angry that she wants John to leave her house immediately. One thing she may do is make loud, aggressive

noises—which may even be words—purely for the sake of causing John, in one way or another, to leave. Mary’s action here is what one would do to scare off a threatening nonhuman animal, and it is odd to imagine acting that way towards another person. Even if Mary succeeds in making John leave, in this case there is intuitively no communication. But also John and Mary are not engaging in the joint activity of making sense of one another, so this case is not a counterexample to the claim that engaging in such activity is necessary for communication. In the case under discussion, it would be much more natural for Mary to tell John to leave, a clear of communication, and even if Mary utters that command in a blind rage—John fuming as well and storming out after—John and Mary are nonetheless making sense of each other when she utters those words and John understands her and acts accordingly. So neither does modifying the case in that way provide a counterexample.

Thus, I claim that the basic purpose of conversation is that of finding mutual intelligibility. Now, following Evans (1982: 310–311) and Heck (1995, 2002), I hold that there is a tight connection between understanding and the basic purpose of conversation. The point of understanding is to serve the basic purpose of conversation, by making that basic purpose’s satisfaction possible. In particular, there is understanding of an assertion just in case the basic purpose of conversation has been enabled. Thus, the following claim is established.

Understanding and Mutual Intelligibility There is understanding of an assertion just in case the goal of making sense of one another has been enabled on the basis of that assertion.

Now, recall the claim Understanding and Stalnaker’s Principles, from the previous section, that there is understanding of an assertion just in case the commonplace effect of that assertion brings it about that the context set satisfies the three Stalnakerian principles. In order to establish that claim—given that Understanding and Mutual Intelligibility is now on the table—I turn to establishing a connection between the satisfaction of Stalnaker’s principles and the goal of mutual intelligibility. I do so by considering the nature of the common ground.

So far in this chapter the common ground has been characterized as the information mutually taken for granted by the interlocutors for the purposes of conversation. But now more can be said in refinement of this characterization. I propose that the common ground is the background information drawn upon by interlocutors in the joint activity of finding mutual intelligibility. Thus, if the common ground is updated, then the project of achieving mutual intelligibility has been served, since there is thereby more information that the interlocutors may rely upon in doing so. Therefore, *if there is a clear way of updating the context set on the*

basis of the commonplace effect of an assertion, then there is a clear way of contributing to the joint activity of making sense of one another on the basis that assertion. Moreover, I propose that the common ground exhausts the resources that interlocutors have in making sense of one another: since in conversation interlocutors are jointly engaged in achieving mutual intelligibility, the information drawn upon in making sense of one another must be information mutually taken for granted. Therefore, *there being a clear way of contributing to the joint activity of making sense of one another on the basis an assertion requires that there be a clear way of updating the context set on the basis of the commonplace effect of an assertion.*

Given the two claims emphasized in the previous paragraph, the following follows.

Mutual Intelligibility and Stalnaker's Principles The goal of making sense of one another has been enabled just in case there is a clear way of updating the context set on the basis of the commonplace effect of an assertion.

So, with this connection between mutual intelligibility and Stalnaker's principles, combined with Understanding and Mutual Intelligibility above, the desired conclusion has been established: there is understanding of an assertion just in case the commonplace effect of that assertion brings it about that the context set satisfies the three Stalnakerian principles.

Thus, my solution to the puzzle of felicitous underspecification is complete. We found in §3.1 that in cases of felicitous underspecification—where there is understanding, yet nothing either meant or said—Stalnaker's principles are satisfied, and also that in cases of infelicitous underspecification the principles are not. And in the present section we found a deeper theoretical explanation of the possibility of felicitous underspecification, as well as limits, by refining Grice's claim about the cooperative nature of conversation in terms of the joint activity of achieving mutual intelligibility. Crucially, that refined claim about the nature of conversation provided a general claim about the role of understanding, as well as a refined conception of common ground.

To conclude this section, let me say a little bit more about the present solution to the central puzzle concerning felicitous underspecification: that of explaining how there may be understanding of an utterance despite there not being anything that is meant or said. What is meant and what is said stand at opposite extremes on a scale of sociality. What is meant is something private, in that it is determined by the speaker's mental states. But what the speaker has said is fully public in such a way that the speaker is on the hook for having said what they said (Camp 2018). That the speaker said what they said can transcend the particular interaction between the speaker and addressee, as the addressee can take it from their interaction with

the speaker and share it with others. So what is said is something social, in a large-scale, collective sense of “social.”⁸ Understanding, in contrast, is neither private nor social. It is a joint interaction between interlocutors. So it is not entirely up to the speaker; hence it is not private. But it is merely a small-scale, interpersonal interaction, and need not give rise to anything public. All that is required for understanding, on my view, is that the interlocutors’ joint project of making sense of one another is served, for which it is sufficient that there be some way of updating the information that is drawn in that joint project.

3.4 Looking Ahead

The main goal of this chapter was to account for the understanding present in cases of felicitous underspecification, despite the fact that in such cases there is nothing meant or said. I also explained the limits of such understanding. I began, in §3.1, by showing how three principles proposed by Stalnaker—Uniformity, Informativeness, and Contentfulness—are satisfied in cases of felicitous underspecification—but also rule out cases of infelicitous underspecification. Next, in §3.2, I elaborated the connection between understanding and Stalnaker’s principles. Finally, in §3.3 I developed a Gricean proposal—related but not identical to that in the previous chapters—concerning the connection between understanding and the joint activity of making sense of one another, in order to provide a theoretical basis for the connection between understanding and Stalnaker’s principles.

In order to set the stage for the next, and final, chapter, I note that there has been no mention of the semantic properties of demonstratives in the accounts of demonstrative understanding so far in the past two chapters. So the plan for the next chapter is to consider what is distinctively linguistic about demonstratives, and related expressions, and develop a formal semantics on that basis. In the conclusion, I draw together the Gricean considerations

⁸My account of what is said, from the introduction, explains how it is out in the open in that way, since it holds that what is said is the particular, determinate claim that the common ground entails that the speaker meant to say. But I note that Camp (2018) suggests that characterizing what is said requires appealing to more than just common ground. She argues that cases of insinuation can be such that it is common ground that the speaker has made the relevant claim, yet it is not thereby the case that the speaker has said it, in the relevant sense, since the speaker is merely insinuating. So she suggests that, in order for something to be truly said, it must go on what she calls the “conversational scoreboard.” For the purposes of this chapter, it is open to me to accept Camp’s suggestions, since one may appeal to the conversational scoreboard but still hold that something is said *only if* it is common ground that the speaker means to say it. But I suspect, following Bach & Harnish (1979), that in cases of insinuation it is not the case that it is common ground that the speaker has made the relevant claim, and I do not think that Camp’s argument to the contrary is decisive—see page 55 of Camp 2018 for the argument. So I would like to leave open the possibility that there is no need for a notion of conversational scoreboard in order to characterize what said.

concerning demonstrative understanding in this and the previous chapter with the formal semantics of the next one.

Before moving on, however, I wish to conclude by comparing the conception of understanding that emerged from this chapter with my refined Gricean account of understanding from the previous chapters. As introduced in §1.2 of the first chapter, my refined Griceanism holds the following about utterance understanding. The speaker must have an INFORMATIVE intention to influence the addressee's mind in some way, and two COMMUNICATIVE intentions, where the PRIMARY communicative intention is to engage with the addressee in the joint activity of making sense of one another, where that activity is to generate the addressee's recognition of the informative intention; and the SECONDARY communicative intention is that the informative intention be fulfilled, in virtue of the addressee's recognition as specified in the secondary communicative intention. With all those intentions of the speaker's in place, understanding then holds just when the speaker's primary communicative intention is fulfilled, which is to say that the interlocutors engage in the joint activity of making sense of another, and their doing so generates the addressee's recognition of the speaker's informative intention.

My refined Griceanism attempts capture what is involved in understanding an utterance of any kind whatsoever, though for present purposes let me consider only the case of assertion, which has been the focus of this chapter. In the case of assertion, the informative intention is standardly taken to be for the addressee to form some belief. But I propose that a speaker *S*'s informative intention, in making an assertion, is of the following form: *S intends to say that P*. Thus, it need not be that the speaker is attempting to influence the addressee's mental life, but instead simply that the speaker intends to say something (McDowell 1998b). The present virtue of this modification of my refined Griceanism is that it lines up with my account of what is meant by assertion from the introduction to this chapter, according to which what is meant is what the speaker intends to say.⁹

The proposal defended in the present chapter is that there is the understanding of an assertion just in case the commonplace effect of that assertion brings it about that the common ground satisfies Stalnaker's principles. This proposal allows there to be understanding without anything that is meant. So the proposal entails that there can be understanding of an assertion without the speaker having the informative intention required by my refined

⁹Note that the present proposal about a speaker's informative intention and what is meant raises the question of how a speaker may gain the concept of *saying* that is deployed in the informative intention; however, I set that question of concept acquisition aside.

Griceanism—hence also without the presence of the communicative intentions it posits, since those communicative intentions make reference, in their content, to the informative intention. Thus, on the present proposal, a speaker need not have various sophisticated intentions: what matters for understanding is that an informative intention is represented with enough specificity *in the common ground*.

But now recall the motivation that was given in previous chapters for the necessity of the various intentions posited by my refined Griceanism: cases of sneaky intentions from §1.2 where the distinctive openness of communication is lacking, and Loar cases from §2.3.2 where there fails to be understanding of a demonstrative utterance despite the fact that the interlocutors are thinking of the same object as referent.¹⁰ In order for the present Stalnakerian account of understanding to be supported, it must ultimately be shown how the account can itself explain how there is a lack of understanding in those cases. I do not do so in this dissertation, but I conclude this chapter with a brief and programmatic suggestion. A crucial part of the explanation I provided in the previous chapters for the lack of understanding in those cases involved appeal to how understanding involves the joint activity of making sense of one another. Moreover, as was seen in the previous section, the connection between understanding and Stalnaker's principles was motivated by appeal to the connection between understanding and the joint activity of making sense of one another. Thus, I suggest, the present Stalnakerian account of understanding shares with the refined Gricean one the requirement that the interlocutors are engaging in that joint activity; hence, the present account should be able to ultimately explain in a similar way the lack of understanding in the cases from the previous chapters.

¹⁰Another potential issue for the present proposal arises from the case of the generals considered in §2.1 of the previous chapter. In that case, there can be understanding of utterances between two individuals, despite the fact that it cannot become common knowledge that either individual has made any such utterance, and hence also that it cannot become common knowledge that either has intention to say something with one of their utterances. The issue the case of the generals raises here, however, relies upon the claim that something cannot be a part of the common ground if it is not common knowledge between the interlocutors. But the information mutually taken for granted by the interlocutors, for the purposes of making sense of one another, need not be common knowledge.

Chapter 4

Saliency and Deixis in the Anaphoricity of Demonstratives

This chapter is about the formal semantics of demonstratives, e.g. “that man,” and definite descriptions, e.g. “the man.” What are the similarities and differences between the meanings of these expressions? My main goal is to make sense of how demonstratives are sensitive to demonstrations in a way that definite descriptions are not. This difference in demonstration-sensitivity is illustrated by (what I call) the “deictic contrasts” (Maclaran 1982).

- (1) *Context: John and Mary are in an art gallery with a number of paintings in front of them.*
 - a. That painting [pointing at a painting] is beautiful.
 - b. #The painting [pointing at a painting] is beautiful.
- (2) *Context: John is directing Mary as they are working together rearranging his living room.*
 - a. You will take that chair [pointing to one chair] and I will take that chair [pointing to another chair]
 - b. #You will take the chair [pointing to one chair] and I will take the chair [pointing to another chair].

In (1), a demonstrative may be used alongside a demonstration to pick out one of the paintings. But a definite description leads to infelicity (with or without the demonstration). Relatedly, in (2) type-identical demonstratives within a single sentence may be used, alongside multiple demonstrations, to talk about one and then another chair. But not so type-identical definite descriptions.

Alongside the deictic contrasts, consider (what I call) the “anaphoric contrast” (Roberts

2002).

- (3) A woman entered from stage left. Another woman entered from stage right.
- a. ...This woman was carrying a bouquet of flowers.
 - b. ...#The woman was carrying a bouquet of flowers.

The demonstrative continuation is felicitous but not so the descriptive one.

The main contribution of this chapter is thus a treatment of definite descriptions and demonstratives that makes sense of the deictic and anaphoric contrasts—a task which also involves accounting for the semantics of demonstrations. A good place to start might appear to be one of the two main approaches to definite descriptions. One of them emphasizes uniqueness: as a first gloss, a definite description “The *F*” presupposes that there is a unique *F*, and so may be used to say something about the unique *F*. The other approach emphasizes familiarity: a definite description “the *F*” presupposes that there is some *F* that is familiar to the conversational participants, and so may be used to say something about the familiar *F*. I begin in §4.1, however, by showing that both approaches fail to account for the ways in which definite descriptions behave in the contrasts above.

Fortunately, there is an alternative to these main camps, which combines uniqueness and familiarity (Roberts 2003, Barker 2004). This hybrid approach does account for the relevant behaviour of definite descriptions. Here is its central tenet.

Definite Descriptions Definite descriptions, like indefinites, e.g. “a man,” have existential main content. But they differ in their presuppositional content (indefinites have no presuppositions). A definite description “the *F*” presupposes that there is a unique familiar *F*, and if defined helps update the information about the unique familiar *F*.

With an account of definite descriptions in place, we may then ask how it may be extended to capture the meaning of demonstratives. I show that Roberts’s (2002) way of doing so fails to capture the behaviour of demonstratives in the anaphoric contrast, and I conclude §4.1 by arguing that the proper way forward involves providing a new theoretical understanding of the notion of familiarity, as well as an improved semantics for demonstratives and demonstrations.

Accordingly, in §4.2 I refine and extend the hybrid approach along those lines. So, in addition to the thesis in the previous paragraph, my overall proposal includes the following.

Familiarity A FAMILIAR *F* is a discourse referent that is salient or deictic. Roughly, a salient

discourse referent is one established by an overt indefinite, and a deictic one by the interlocutors' mutual perceptual environment.

Demonstratives A demonstrative “that *F*” presupposes that there is at least one familiar *F*, and if defined helps update the information about the most recently established familiar *F*. (Demonstratives also have the same existential main content as definite descriptions and indefinites.)

Demonstrations A demonstration presupposes that there is at least one deictic discourse referent, and establishes one of the deictic discourse referents as the most familiar.

The four theses just presented—between this paragraph and the previous one—represent in broad strokes the view I propose in this chapter. The way in which the proposal is filled in draws upon a dynamic semantics, in the style of Bittner (2009, 2011); accordingly, the talk of discourse referents, introduced in Familiarity, will be refined as the chapter proceeds.

For the central applications of my proposal, intensional resources are not required. But, in §4.3, I address two intensional phenomena. First, demonstratives do not only differ from definite descriptions with regard to demonstration-sensitivity, but also in that the former have deictic features: “that,” for instance, has the deictic feature of distality; “this” the deictic feature of proximity. Drawing on Roberts (2020), I account for how deictic features are perspective-sensitive in that they typically help locate the referent's position relative to the speaker: the distal demonstrative “that” is typically used to talk about objects relatively far from the speaker; the proximal demonstrative “this” to talk about objects relatively close. This perspective-sensitive content, however, is not-at-issue in that it is similar to the effect, for instance, of non-restrictive relative clauses. To capture this not-at-issueness, I draw on Murray's (2014) account of appositives, which itself builds on Bittner's dynamic system and crucially relies upon intensional resources.

The second intensional phenomenon I discuss is rigidity. The present focus on differences between demonstratives and definite descriptions might bring to mind, for a seasoned reader, Kaplan's (1989) influential insight that they differ with regard to rigidity. But I argue that such considerations should not in fact be given a central role in the study of the semantics of expressions such as demonstratives and definite descriptions. And I use that point as an opportunity to discuss more broadly the relationship between a dynamic semantics treatment of context-sensitive referential expressions such as demonstratives and definite descriptions, and the more traditionally philosophical notion of reference.

In conclusion, I briefly address some recent related work by Mandelkern & Rothschild (2020). After the conclusion, in the appendix, a formal system is presented.

4.1 Familiarity and Uniqueness

The two main views of definite descriptions emphasize either their implications of uniqueness or of familiarity, giving one or the other presuppositional status. I show in §4.1.1 and §4.1.2 that both these views fail to account for the infelicity of definite descriptions in the deictic and anaphoric contrasts. In §4.1.3, I consider a third, hybrid view that combines familiarity and uniqueness.

4.1.1 Presupposing Uniqueness

As initially glossed, the uniqueness view holds that a definite description “the F ” presupposes that there is a unique F , and if defined denotes the unique F .¹ This proposal can be implemented in a static, extensional framework. Here is the lexical entry for the definite determiner (see this footnote² for a brief explanation of the formalism from the lambda calculus used throughout this chapter).

$$(4) \quad \llbracket the \rrbracket = \lambda P_{et}.\exists!x(P(x)).!x((P(x)))$$

The meaning of the definite determiner takes in a property P , presupposes that there is a unique object that is P , and if defined denotes the unique object that is P .

Though the semantic entry is correct for present purposes, the initial gloss that motivated it is not the whole story. Consider (5).

¹Prominent versions of the overarching uniqueness view are in Russell 1905, Strawson 1950, Evans 1977, Heim 1990, Kadmon 1990, King 2001, Elbourne 2005, 2013. I group together here authors who disagree about whether uniqueness is asserted or presupposed, which is roughly the distinction between E- and D-type variants of the view. The version I present here section is E-type, where uniqueness is presupposed, but my criticism below applies equally well to the D-type variant.

²In the lambda calculus, a term for a function begins with a series of variables with suffixes and each prefixed by a lambda “ λ .” This initial part of the term represents what the function takes as its arguments, and the suffixes represent what type of thing each argument is. So in (4), for instance, the overall function takes as argument something of type et , which itself is a function from entities to truth-values (i.e. a property). After the initial part of the term, which ends with a period, there may be definedness conditions added for the function. So in (4), for instance, the overall function is a partial function, requiring that the property argument have a unique instance. Note that definedness conditions here model semantic presuppositions. Finally, the last part of the function term, after the second period, represents what the function outputs (if the definedness conditions are met). So in (4), for instance, the function outputs the individual that is the unique instance of the property taken in as argument.

- (5) *Context: A student makes a comment to another at the beginning of class.*
The teacher is happy today.

Here the definite description is felicitous, even though there is manifestly more than one teacher in all of existence. But, if the definite determiner's meaning takes as argument a single property P , how can that felicity be?

Following the influential insight of Stanley & Szabó (2000), we may take a cue from the fact that definite descriptions may be bound, as in (6).

- (6) Every rich person's house has sleeping pills in the master bathroom.

In order to make sense of how "the master bathroom" may co-vary with "every rich person's house," Stanley & Szabó posit covert variables in the definite description's syntax, as follows.

- (7) [the [NP fx]]

In (7), " x " is an individual variable of type e , and " f " a relation variable of type eet . In (6), " x " is bound by the higher quantifier, and " f " may be assigned by a contextually given assignment function the relation y is a room in x . The meaning of the nounphrase ("NP"), a property, and the meaning of " fx ," also a property, compose via predicate modification, an operation which intersects two property arguments to provide a new property.³

Returning to the unbound case of (5), both covert variables are determined by a contextually determined assignment function. Perhaps " x " is assigned the particular teacher in the class, and " f " the identity relation. Here again the value of the variables combine to form a property, which is intersected via predicate modification with the overt NP property, and the resulting property is the one the use of the definite description presupposes is uniquely satisfied. Hence, the property P that the definite determiner's semantic value takes as argument is the overt NP potentially restricted in a contextually given way. This restriction is what is called the "implicit domain restriction" of definite descriptions.

There is much to be said for the uniqueness view. Let me present a recent argument for it from Blumberg (2020), since working through his argument leads naturally to how the uniqueness view may be extended to demonstratives. Consider first the following contrast.

³See Heim & Kratzer (1998: 65) for more on predicate modification. The implicit domain restriction modeled here may be represented with only the traditional compositional rule of functional application, but doing so involves complicating the type of the variable " f " in a way that is needless for present purposes.

- (8) Context: *It is common knowledge that there was exactly one witch at a Halloween party last night.*
- a. The witch left early.
 - b. ??The tall witch left early.

What is the source of (8-b)'s oddness? Consider the following pragmatic principle (Schlenker 2005).

- (9) MINIMIZE RESTRICTORS!: Relative to a context c , a use of a definite description of the form $\ulcorner \textit{the } \psi \phi \urcorner$ is infelicitous if ψ is redundant, i.e. if:
- a. $\ulcorner \textit{the } \phi \urcorner$ is grammatical and ϕ (plus implicit domain restriction) denotes a singleton set in c ; and,
 - b. ψ does not serve another purpose, i.e. it is pragmatically irrelevant.

According to this principle, if additional predicative material is added to a definite description—which is not required to make the overall predicative component denote a singleton set, and which does not serve any other pragmatic purpose—then the definite description is marked. In (8), “witch” plus the plausible implicit restriction of y was a person at x , where x is the party, contributes a singleton set. So adding the further predicate “tall” is redundant. Hence, the use of “the tall witch” is marked.

In light of Minimize Restrictors!, consider the following contrast between demonstratives and definite descriptions, first noticed by Wolter (2006).

- (10) a. The author of *Blood Meridian* is a genius.
 b. ??That author of *Blood Meridian* is a genius.

Blumberg's insight is that the infelicity of the demonstrative (10-b) may be explained by Minimize Restrictors!, given the uniqueness account of definite descriptions, and the following simple and orthodox account of demonstratives.

Interpreting Kaplan (1989) loosely, as must be done and as many have, we get the following picture of the relationship between demonstratives and definite descriptions.⁴

- (11) That man = The *actual* man *being demonstrated*

Rigidity is not of central concern in this chapter, so I ignore here the rigidifying operator

⁴Prominent views of this type are in King 2001, Schlenker 2005, and Elbourne 2008.

“actual.” Crucially, a demonstrative has additional covert material, provided by a demonstration, which affects the uniqueness presupposition. Accordingly, here is the lexical entry for a demonstrative, as well as the syntactic structure of its phrase.⁵

- (12) a. $\llbracket \textit{that} \rrbracket = \lambda Q_{et} \lambda P_{et} . \exists ! x (Q(x) \wedge P(x)) . \iota x (Q(x) \wedge P(x))$
 b. [That *g* [NP *fx*]]

The demonstrative’s meaning takes in two properties: *Q* and *P*. The second, *P*, is supplied in the same way as the definite description’s argument, by implicit domain restriction of the overt NP. In addition, *Q* is supplied by the contextual assignment to a covert property variable “*g*”, which represents the semantic effect of a demonstration. So, for instance, a use of “that table” may be used to talk about the table *to the right* (demonstration) *in this room* (implicit domain restriction).

Minimize Restrictors! does not mention demonstratives, but, given the uniqueness treatments of definite descriptions and demonstratives, Blumberg derives from it an explanation of the oddity of the demonstratives in Wolter’s contrasts. Given the proposed semantics, a use of a demonstrative $\ulcorner \textit{that } \phi \urcorner$ is equivalent to a use of a definite description $\ulcorner \textit{the } \psi \phi \urcorner$, where ψ is a predicate that denotes a property equivalent to the property supplied by the demonstration. Since, by Minimize Restrictors!, $\ulcorner \textit{the } \psi \phi \urcorner$ is in competition with a use of $\ulcorner \textit{the } \phi \urcorner$, it follows that $\ulcorner \textit{that } \phi \urcorner$ is also in competition with $\ulcorner \textit{the } \phi \urcorner$. Thus, in a situation where ϕ denotes a singleton, both $\ulcorner \textit{that } \phi \urcorner$ and $\ulcorner \textit{the } \psi \phi \urcorner$ will be infelicitous, and $\ulcorner \textit{the } \phi \urcorner$ preferred (unless ψ serves some other pragmatic purpose).

The second condition of Minimize Restrictors!—that the additional restrictor material may felicitously serve some other pragmatic purpose—is important to note. For there are felicitous uses of demonstratives where the overt material is semantically unique, as in the following.

- (13) That US president is at it again!

These uses of demonstratives have been said to partake in the phenomenon of “emotional deixis” (Lakoff 1974, Acton & Potts 2014), since there is an emotive ring to them.⁶ I return to this phenomenon in §4.2.3. For present purposes, note that the current appeal to Minimize

⁵I leave the treatment of deictic features to §4.3.

⁶Note that, regardless of the subtle emotive ring, their mere felicity poses a problem for the accounts of demonstratives from Hawthorne & Manley (2012) and Nowak (2018) that claim demonstratives presuppose that their overt NP component, plus implicit domain restriction, does not generate a singleton set.

Restrictors! should explain emotional deixis as follows: the additional restriction “g” of the demonstrative is supplied an expressive property, such as the overt “stupid” would. So the demonstrative is felicitous, since the redundant restrictor is serving some other pragmatic purpose.

But we can immediately see that something has gone wrong with the uniqueness approach. Consider the following in contrast with (13).

(14) The US president is at it again!

There simply is not a use of (14) in which the definite description has the emotive ring of the demonstrative in (13); the sentence cannot mean something such as the *stupid* US president is at it again. But, if there is an expressive property that can be assigned to the covert restriction “g” of the demonstrative, then in principle something similar could be assigned to the covert “x” and “f” of a definite description’s implicit domain restriction.

It might be possible to draw upon a semantic difference between expressive and non-expressive properties, and allow that expressive properties may be assigned to “g” but may not be the result of combining what is assigned to “x” and “f.” But the uniqueness view of definite descriptions falters for similar reasons on the deictic contrasts (1) and (2), where no expressive properties are involved. In the attempted deictic uses of definite descriptions, why cannot the covert restriction allow one to talk about the painting being pointed at, yet the demonstrative’s further restriction can? Or to talk about the different chairs in turn? If there are general pragmatic principles at work in determining what “x” and “f” are assigned in context, then it is a total mystery why these cannot be determined in such a way as to make the definite descriptions in (1-b) and (2-b) felicitous. So, the defender of uniqueness faces the following difficult question: Given that implicit domain restriction for definite descriptions must be made sense of, why is it that such restriction cannot occur in their attempted deictic uses? I take the difficulty of this question to indicate that we should look elsewhere for an account of the semantics of definite descriptions.⁷

Finally, it is worth mentioning that King (2001: 126–139) hits upon the problem for the uniqueness view that is here raised by the deictic contrasts. In response, he suggests that

⁷There is another way of making sense of implicit domain restriction than the one just given, which employs covert situation variables (Heim 1990, Elbourne 2005, Wolter 2006, Ahn 2019). And Wolter (2006) in fact develops this situation-theoretic account of implicit domain restriction in order to explain the infelicity of definite descriptions in the deictic contrasts, though her account relies upon claims about the pragmatics of situations, a generally underdeveloped topic. I do not mean to suggest here that it is not worth developing. I have simply chosen to take another route.

something about the structural location of demonstratives' additional implicit restriction may explain how they can be restricted in certain ways that definite descriptions cannot. That is, the implicit domain restriction of definite descriptions occurs via their noun phrase complements; whereas, the demonstrative itself combines with the variable that provides its additional restriction (recall the logical forms in (7) and (12-b)). But it is unclear how the mere structural placement of variables should make the relevant difference. On the present view, the contextual resolution of free variables is up to general pragmatic mechanisms, which should presumably act in such a way as to bring about felicity. But, again, it is unclear why the implicit domain restriction cannot then be resolved in such a way as to bring about the felicity of the definite descriptions in the deictic contrasts.

4.1.2 Introducing Familiarity

I turn now to the second main approach to definite descriptions, which foregrounds familiarity.⁸ As glossed in the introduction, the view holds that a definite description presupposes that there is a familiar *F*, and accordingly may be used to say something about that *F*. The notion of familiarity is typically understood in terms of discourse referents, so that there is a familiar *F*, at a given point in a discourse, just in case there is discourse referent for something *F* present at that point in the discourse. Discourse referents are, in turn, typically elaborated within a dynamic semantics. I do not lay out a dynamic semantics in this section, since there is considerable technical sophistication in doing so, and I develop one later for my own proposal in §4.2. A relatively informal discussion will suffice for present purposes.

Talk of discourse referents informally captures how there is semantic machinery that coordinates information at the discourse level with the aid of indices on expressions such as indefinites and definite descriptions. It is crucial to note that specifying this semantic machinery does not require adding entities, the alleged discourse referents, to the model as potential referents. A discourse referent is paradigmatically introduced by an indefinite, so let me say a few words about the dynamic treatment of indefinites. An indefinite semantically introduces a discourse referent associated with a given index, but this is simply to say that an indefinite “an *F_n*” may be coordinated semantically with, say, a definite description “the *F_n*,” despite the fact that the definite description may be outside of the indefinite’s syntactic scope.⁹ Thus, the coordination is at the discourse level in that a definite description in one

⁸Prominent views of this type are from Heim (1982), Szabó (2000), and Ludlow & Segal (2004). This type of view also includes the recent proposals of Coppock & Beaver (2015) and Köpping (2020).

⁹I am following here the convention of using a superscripted index for the introduction of a discourse ref-

sentence may be coordinated with an indefinite in an earlier sentence. More specifically, for example, a sentence “An F^n is G ” may be followed downstream in a discourse by a sentence “The F_n is H ,” and as a matter of semantics the discourse may thereby make the existential claim that there is something that is F , G , and H .

So our informal grasp of the dynamic semantics of indefinites is that they introduce discourse referents associated with indices. Let me now informally elaborate the semantics of definite descriptions. We may retain the basic structure of the uniqueness approach, but with the modifications that uniqueness is assessed among discourse referents, and that implicit domain restriction is provided by an index. Hence, a definite description “the F_n ” presupposes uniqueness among F discourse referents initiated with the index “ n ,” and if defined helps elaborate the information about that discourse referent. On the present treatment of definite descriptions, the uniqueness presupposition is trivialized, since a definite description may be subscripted in one way or another in order pick up one or another discourse referent. Assuming that there cannot be distinct discourse referents associated with the same index, the present way of stating the familiarity view—by appeal to a trivialized uniqueness presupposition—is equivalent to a more traditional version of the view, according to which definite descriptions simply presupposes that there is some F discourse referent under a given index (Heim 1982). But the way of stating it that I have chosen is helpful because it clarifies the familiarity approach’s relationship to the other views of definite descriptions covered in this section.

Furthermore, the present version of the familiarity view vivifies the problem that the deictic and anaphoric contrasts pose, since the contrasts show that a definite description’s uniqueness presupposition is not trivialized by anaphoricity (Roberts 2003, Wolter 2006). Consider first the use of a definite description in the anaphoric contrast, elaborated in (15). The first two sentences introduce discourse referents for women, each associated with distinct indices, and the familiarity view predicts either of the continuations should be acceptable to elaborate the information about one or the other woman.

- (15) A woman¹ entered from stage left. Another woman² entered from stage right.
- a. ...#[The woman]₁ was carrying a bouquet of flowers.
 - b. ...#[The woman]₂ was carrying a bouquet of flowers.

But neither continuation is felicitous; when a definite description is anaphoric to an indefinite,

erent, and a subscripted index to indicate which discourse referent an anaphoric expression picks up.

a uniqueness effect nonetheless remains.

So the anaphoric contrast poses a clear problem for the familiarity view, and consider in that light the attempted uses of definite descriptions in the deictic contrasts, repeated here.

- (16) *Context: John and Mary are in an art gallery with a number of paintings in front of them.*
- a. #The painting [pointing at a painting] is beautiful.
- (17) *Context: John is directing Mary as they are working together rearranging his living room.*
- a. #You will take the chair [pointing to one chair] and I will take the chair [pointing to another chair].

Suppose, as is plausible, that a demonstration is similar to an indefinite in that it introduces a discourse referent. According to the familiarity theory, in (16-a) and (17-a) the definite description should thus be able to be co-indexed with the demonstration. Yet there is infelicity, which also seems to result from a violation of a uniqueness presupposition, for there are multiple paintings/chairs around the interlocutors.

Before moving on, I wish to note that the main traditional motivation for the familiarity approach over the uniqueness one is the phenomenon of donkey anaphora. I leave discussion of that phenomenon to the next section and the conclusion. But this point about donkey anaphora is part of a larger observation about definite descriptions and demonstratives, and saying a few words about the larger observation now will help bring about some terminological clarity.

The larger observation is that both definite descriptions and demonstratives satisfy what is called the “anaphoric paradigm” (Heim 1982, Partee 1984, Roberts 2002). What this means is that they both occur (a) without and (b) with overt linguistic antecedents, where in some of the latter cases (c) the linguistic antecedent has the demonstrative or definite description within its syntactic scope. (Donkey anaphora are cases that vividly illustrate uses of type (b) that are not also of (c), in which an indefinite binds a definite description or demonstrative outside of its syntactic scope.) Cases of each these types are discussed in the chapter, but here is some simple data to illustrate the taxonomy.

- (18) a. *Context: A dog walks into the room.*
The/that dog is beautiful.
- b. A woman walked in. The/that woman was happy.
- c. Every guilty senator testified to congress before the/that senator was harangued.

In this chapter, I have adopted the common terminology of calling type (a) cases “deictic,” type (b) cases “anaphoric,” and type (c) cases “bound.”

The main virtue of dynamic semantics, with regard to referential expressions such as definite descriptions and demonstrations, is in uniformly treating all of these uses. In each case, the referential expression is anaphoric to a discourse referent. So note that the common terminology just highlighted is ultimately misleading in the present theoretical context, since all uses of these referential expressions are anaphoric (which is why the three types of uses together are said to constitute the *anaphoric* paradigm). But I retain the common terminology because it is intuitively satisfying.

In this section, I am focused on how the various extant theories of definite descriptions and demonstrations account for deictic and straightforward anaphoric uses of those expressions. In the next section, where I develop my own view within a dynamic semantics, I discuss how the present considerations are related more complex anaphoric data, such as donkey anaphora and cases where there is syntactic binding. Before turning to my own view, however, let me now discuss the final extant view of definite descriptions and demonstratives.

4.1.3 The Hybrid Approach

Having considered the two main views of definite descriptions, we may turn to a third view, which combines familiarity and uniqueness, but without trivializing uniqueness (Roberts 2003, Barker 2004). On this view, a definite description “the *F*” presupposes that there is a unique *F* discourse referent (regardless of index), and if defined helps elaborate the information about that discourse referent. Recall that the familiarity view of the previous section holds that a definite description “the *F*”’s uniqueness presupposition is that there must be a unique *F* discourse referent *associated with a given index*, hence this trivialized uniqueness presupposition may be satisfied while there are multiple *F* discourse referents. That is, a definite description “the *F*” presupposes that there is a familiar *F*, and so may be used to elaborate the information about this or that familiar *F*. But the hybrid holds that the presupposition is that there is a unique familiar *F*, and so a definite description “the *F*” must be used to elaborate the information about the unique familiar *F*. So the present view draws from the uniqueness view since it does not trivialize the uniqueness presupposition in the way that the familiarity view does. But the present view is also distinct from the uniqueness view, since on that view the definite description presupposes that there is a unique thing, satisfying some property, in all of existence. On the present proposal, in contrast, uniqueness is assessed among what is

familiar in discourse.

The hybrid view explains the infelicity of the definite description in the anaphoric contrast, since the two indefinites introduce two woman discourse referents. In order to make sense of the infelicity of the attempted deictic uses of definite descriptions, we must broaden our understanding of how discourse referents may be initiated. As mentioned above, a use of an indefinite is the paradigm case of initiating a discourse referent. But it is plausible that, when an object is mutually taken to exist in the interlocutors' perceptual environment, there is thereby a discourse referent for it initiated. So in the deictic contrasts there are multiple paintings/chairs, which leads to the failure of the definite descriptions' presuppositions. Thus the hybrid view provides an elegant account of the infelicity of the definite descriptions in the cases at hand. For this reason, I take the view to be correct. But there are two issues that I wish to raise with the main way in which the view has been developed, in the work of Roberts (2002, 2003).

The first issue stems from the fact that the appeal to a nonparadigmatic way in which discourse referents may be introduced raises the question of how exactly the deictic discourse referents—those introduced on the basis of the interlocutors' perceptual environment—may be considered alongside the paradigm salient discourse referents—those introduced by an indefinite. Roberts's (2003) answer relies on the notion of the COMMON GROUND: the information mutually accepted by the interlocutors for the purposes of conversation (Stalnaker [1978] 1999, 2002). For her, a discourse referent may be initiated on the basis of INFORMATIONAL EXISTENCE, when the existence of an object can be inferred on the basis of the common ground. Thus when something is perceptually recognized by the interlocutors, the common ground is appropriately updated, and a corresponding discourse referent is thereby initiated. So the present proposal is that deictic discourse referents are subtype of MERELY INFORMATIONAL ones: those not introduced by an indefinite but instead in virtue of information in the common ground.

The issue, however, is that deictic discourse referents do not behave like other members of the alleged type of merely informational discourse referents. As the deictic contrasts show, deictic discourse referents are relevant for a definite description's uniqueness presupposition in that they can bring about the presupposition's failure. But not so for merely informational discourse, as the following data from Mandelkern & Rothschild (2020) show.¹⁰

¹⁰(19-a) is based on an example given by Heim (1982), where she considers the pronoun "it" instead of a definite description.

- (19) a. Sue bought a sage plant and bought eight others along with the sage plant.
 b. Several couples came in today. There was one, a man and a woman. The man was being so annoying.

Both these sentences are felicitous. Hence the definite description at the end of (19-a) does not have its uniqueness presupposition violated, despite the fact that many sage plants can be inferred to exist on the basis of the mention of eight sage plants earlier in the sentence—which should have updated the common ground before the definite description at the end of the sentence is considered. Similarly, the definite description in the final sentence of (19-b) does not have its uniqueness presupposition violated, despite the fact that multiple men can be inferred to exist on the basis of the first sentence’s update to the common ground. So I conclude that the category of merely informational discourse referents is not a genuine category, in that such alleged discourse referents do not genuinely enter in the context when they can be established on the basis of the common ground. They are not antecedently present, like deictic discourse referents are, in such a way as to lead to the infelicity of a definite description.

Before moving on to the second issue, it must be noted that Mandelkern & Rothschild take their data in (19) to in fact help support the familiarity view of definite descriptions. Their argument for the familiarity involves considering minimal pairs such as (19-a) and (20-a), and (19-b) and (20-b).

- (20) a. ??Sue went to the store and bought eight sage plants along with the sage plant she bought.
 b. ??Several couples came in today. The man was being so annoying.

The argument goes as follows. The felicity of (19-a)/(19-b) is supposed to support the claim that definite descriptions do not have any uniqueness presupposition, as maintained by the familiarity view. And then the reason they suggest that there is infelicity in (20-a)/(20-b) is that a discourse referent for the definite description must be accommodated,¹¹ since the definite description has no indefinite as antecedent, and there is something about that process of accommodation that triggers a uniqueness requirement—where, finally, that uniqueness requirement is violated by the merely informational discourse referents.

The main issue with this argument is that it requires that the (putative) merely informational discourse referents both are and are not genuine discourse referents. If, for instance,

¹¹Accommodation is a process whereby a presupposition is not met by the prior context, yet it is smoothly accepted by the audience, and hence there is felicity. See von Stechow (2008) for a recent overview.

it is claimed that there is no discourse referent antecedently present for the definite description in the second sentence of (20-b)—despite the fact that the first sentence mentions several couples—then it cannot also be held that the felicity of (19-b) supports the view that definite descriptions have no uniqueness presupposition—since it cannot also be claimed that the first sentence does initiate a number of discourse referents that *would* violate a uniqueness presupposition.

Again, I take (20-b) and (19-a) to show that merely informational discourse referents are not genuine discourse referents, so those sentences are not problems for the present version of the hybrid view. But there is an open question about the source of infelicity in (20-b)/(20-a), though something like the answer that Mandelkern & Rothschild propose is consistent with the present proposal: since there is no prior discourse referent for the definite description in (20-b)/(20-a), one must be accommodated, and it may be that something about that process of accommodation triggers a uniqueness requirement that concerns merely informational entities (but this triggered uniqueness requirement is not the same as the definite description's uniqueness presupposition). That feature of accommodation—if it is in fact one—is consistent with definite descriptions having their own semantic uniqueness presupposition, and explains the relevant infelicity. I think, however, that the relevant facts about accommodation that are hereby revealed about accommodation are more subtle than Mandelkern & Rothschild suggest, but I relegate a brief discussion of the matter to this footnote.¹²

¹²First, it seems that the infelicity of (20-a) arises because the sentence is claiming that Sue simultaneously bought eight and nine sage plants. Consider the sentence as clarified so that it is clearly saying that Sue bought nine.

(21) Sue went to the store and bought eight other sage plants along with the sage plant she bought.

In (21) there is felicity. And note that there is a definite description that requires accommodation, yet the mention of eight other sage plants does not lead to infelicity.

With regard to (20-b), I suspect that the problem is two-fold: (a) a discourse referent for a man is unable to be accommodated in such a way that it is one of the men of the several couples; and (b) it is unclear how some other man discourse referent could be accommodated. Note that if the example is modified slightly to make it more clear how there is some other man being talked about, there is felicity.

(22) Several couples came in today. The man was being so annoying by asking each of the couples where they met.

The felicity of (22) shows that the mere accommodation of a discourse referent for a definite description does not lead to a uniqueness requirement violated by merely informational discourse referents. The problem with (20-b), I stress, seems to be that a discourse referent cannot be accommodated *that is one of the men in the couples*

The second issue that I wish to raise is with the way that Roberts (2002) extends the approach to treat demonstratives. On her account, a demonstrative “that *F*” presupposes that there is a uniquely demonstrated familiar *F* discourse referent. So demonstratives presuppose demonstrations, which restrict the domain of discourse referents among which the uniqueness presupposition must hold. In this way, the felicity of the demonstratives in the deictic contrasts follows: the demonstration satisfying a demonstrative “that *F*”’s presupposition helps pick out one or another of the *F* discourse referents. But, if demonstratives presuppose the existence of demonstrations, then all of their uses require demonstrations. Intuitively, however, it does not seem that anaphoric or bound demonstratives are accompanied by demonstrations of any kind, or even that a demonstration must be accommodated.

In addition to this intuitive issue, the proposal—packing demonstrations into the presuppositions of demonstratives—makes incorrect predictions. For the parallel between deictic and anaphoric demonstratives is not exact. Consider again the behaviour of demonstratives in those contrasts.

- (23) *Context: John and Mary are in an art gallery with a number of paintings in front of them.*
That painting [pointing at a painting] is beautiful.
- (24) *Context: John is directing Mary as they are working together rearranging his living room.*
You will take that chair [pointing to one chair] and I will take that chair [pointing to another chair]
- (25) A woman entered from stage left. Another woman entered from stage right. This woman was carrying a bouquet of flowers.

In the deictic (23) and (24), a demonstrative, alongside a demonstration, is used to pick out one or the other deictically given painting or chair. In contrast, the anaphoric demonstrative in (25)—even if the distal form “that woman” is used instead—must be anaphoric to the second mentioned woman. Furthermore, in contrast with (24), the setup of the first two sentences of (25) does not allow one to follow with two demonstratives—type identical or not, no matter distal or proximal—to talk about one and then the other woman in turn.

- (26) A woman entered from stage left. Another woman entered from stage right.
a. ...#This woman was taller than this woman.

mentioned in the first sentence. I propose further that a similar explanation holds for the other data Mandelkern & Rothschild provide.

- b. ...#This woman was taller than that woman.

Thus, it is hard to see how anaphoric demonstratives may be accompanied by demonstrations in such a way as to make the anaphoric case parallel with the deictic. Bound cases are similar to anaphoric ones in this respect, and I discuss them as well in the following section.

In sum, the problem is that Roberts treats the anaphoric behaviour of demonstratives as too similar to deictic ones. Despite the fact that those environments are similar with regard to the demonstratives but not definite descriptions being felicitous, it is not the case that demonstrations are present in both. Rather, it is plausible that there are simply no demonstrations present in the anaphoric case, and in the deictic case demonstrations are playing a role similar to indefinites.

4.2 Salience and Deixis

The view I am about to develop retains the semantics of definite descriptions given by the hybrid view: a definite description “the *F*” presupposes that there is a unique familiar *F*, and if defined helps update the information about the unique familiar *F*. Yet, as has emerged from the previous section, the path ahead revises the notion of familiarity and the semantics of demonstratives and demonstrations. I also refine the talk of discourse referents and the semantics of the relevant expressions by drawing upon a dynamic semantics.

4.2.1 Core Proposal

Let me begin with the notion of familiarity. Roughly, I propose that familiarity is constituted by salience and deixis. Recall that talk of familiarity lines up with talk of discourse referents: to say that there is a familiar *F* is to say that there is an *F* discourse referent. So, my proposal about familiarity is that all discourse referents are either salient or deictic. As a first gloss, a salient discourse referent is one established by an overt indefinite, and a deictic one the interlocutors’ mutual perceptual environment.

In order to elaborate my notion of familiarity, I develop some formal machinery to refine the talk of discourse referents, in such a way that respects the distinction between salient and deictic ones. As was hit upon in the previous section, talk of discourse referents recapitulates how chunks of discourse have existentially quantified truth-conditions, in which indefinites may coordinate information with definite expressions outside of their syntactic scope. Making sense of truth-conditions within a dynamic framework involves some subtlety, so I leave

it to §4.3.2. What matters for current purposes is that the introduction of discourse referents occurs via dynamic existential quantifiers and, with the help of a dynamic context, such quantifiers may thereby semantically bind expressions outside of their syntactic scope.

A CONTEXT, or information state, c is a set of information points. I represent formally the distinction between salient and deictic discourse referents with two stacks in the dynamic context. So each INFORMATION POINT i is an ordered pair of sequences. More specifically, each point i is the ordered pair $\langle \sigma_s, \sigma_d \rangle$, where the first sequence σ_s^i helps represent the salient discourse referents and the second sequence σ_d^i the deictic ones. I emphasize though that discourse referents are not elements of any sequence: they are not posited by the formal model as potential referents of definite expressions. Instead, insofar as it is ultimately necessary to talk about discourse referents as existing in the formal model, they correspond to positions among a set of information points that is a context. See Table 4.1 for further illustration.

c	c'	c''	c'''
{ $\langle \langle a \rangle, \emptyset \rangle,$	{ $\langle \langle c, a \rangle, \emptyset \rangle,$	{ $\langle \emptyset, \langle a \rangle \rangle,$	{ $\langle \emptyset, \langle c, a \rangle \rangle,$
	$\langle \langle d, a \rangle, \emptyset \rangle,$		$\langle \emptyset, \langle d, a \rangle \rangle,$
$\langle \langle b \rangle, \emptyset \rangle$	$\langle \langle c, b \rangle, \emptyset \rangle,$	$\langle \emptyset, \langle b \rangle \rangle$	$\langle \emptyset, \langle c, b \rangle \rangle,$
	} $\langle \langle d, b \rangle, \emptyset \rangle$ }	}	$\langle \emptyset, \langle d, b \rangle \rangle$ }

Table 4.1: Discourse Referents. Assume in the background model that only objects a and b are F , and only c and d are G . In context c , there a single discourse referent, for an F object, and it is in the salience stack. In c' that same discourse referent is present, but there is another, for a G object, that is in the top position of the salience stack. The same situation is then illustrated in the deictic stack. In c'' , there is a single discourse referent and it is in the deictic stack. In c''' that same discourse referent is present, but there is another in the top position of the deictic stack.

The bidimensional structure of contexts here is taken from the work of Bittner (2001, 2009). She distinguishes between the center and periphery of discourse attention with two stacks she labels top “ \top ” and bottom “ \perp .” I have however relabeled them the salience stack “ s ” and the deictic one “ d .” Given the two stacks in the context, we may introduce two types of dynamic existential quantifiers. These existential quantifiers are the following update boxes, which appear in logical form, and influence either stack as a matter of their semantics.

- (27) a. $[dx|Fx]$
 b. $[sx|Fx]$

As for their semantic interpretation, both boxes are assigned functions from contexts to contexts. In each case a discourse referent that is indefinite between the individuals that are F is added to the context. For (27-a), the DEICTIC EXISTENTIAL, the discourse referent is added to the top of the stack of deictic discourse referents. For (27-b), the SALIENCE EXISTENTIAL, it is added to the top of the stack of salient discourse referents. See Table 4.2 for further illustration.

c_o	$c_o[[dx Fx]]$	$c_o[[sx Fx]]$
{ $\langle \emptyset, \emptyset \rangle$	{ $\langle \emptyset, \langle a \rangle \rangle$,	{ $\langle \langle a \rangle, \emptyset \rangle$,
	$\langle \emptyset, \langle b \rangle \rangle$,	$\langle \langle b \rangle, \emptyset \rangle$,
	$\langle \emptyset, \langle c \rangle \rangle$,	$\langle \langle c \rangle, \emptyset \rangle$,
	} $\langle \emptyset, \langle d \rangle \rangle$ }	$\langle \langle d \rangle, \emptyset \rangle$ }

Table 4.2: Basic Existential Updates. Here is the effect of each existential updates on the open context $c_o = \{\langle \emptyset, \emptyset \rangle\}$. In the background model, only objects a, b, c , and d are F , so the discourse referent introduced is indefinite between those four objects. I follow the convention of providing the semantic value of for an expression e , $[[e]]$, its context argument on the left hand side, $c[[e]]$, to represent how the semantic value takes in a context and updates it to a new one.

Here is quick note on the syntax that I use throughout this chapter. On the left-hand side of an update box “[$\alpha v, \dots, \alpha v' | C$]” are variables, where “ α ” stands for either “ s ” or “ d ,” and which initiate discourse referents. On the right-hand side there is a condition “ C ,” which may include referential expressions, “ αv_n ,” which pick up the n back value, in the α sequence, at each point in the input context. Note, however, that when these referential expressions pick up values introduced by variables in the left-hand of their own box, I simply write “ x .” There also may appear referential expressions representing definite descriptions and demonstrations, but before getting to those here is a bit more about the existential quantifiers.

The salient existential quantifier (27-b) represents in logical form an indefinite “an F .” For the purposes of this chapter, I assume a simplistic model of salience: a SALIENT discourse referent is one introduced by an overt indefinite, and once a discourse referent is established as salient it remains so. This model is overly simplistic, but the general notion of salience I am appealing to here is fairly uncontroversial, despite the fact that details must be filled in.

In contrast, the introduction of deictic discourse referents and the corresponding deictic existential might seem questionable. A DEICTIC discourse referent is one appropriately related to the interlocutors’ perceptual situation. I do not here dwell on the question of in what

exactly the “appropriate relation” appealed to consists. Roughly, it has to do with an object’s existence being recognized or taken to hold by the interlocutors on the basis of their mutual cognitive grip on their perceptual environment. So, for instance, when John and Mary are facing a wall of painting in an art gallery, those paintings have corresponding deictic discourse referents in the discourse. But not so the paintings in the room adjacent. Moreover, I am suggesting that a deictic discourse is initiated by an element of logical form: the deictic existential quantifier in (27-a). This way of implementing how deictic discourse referents are initiated might seem implausible, but I adopt it here because it provides a simple resolution of the issue, raised in §4.1.3, concerning how it is that deictic discourse referents are genuine aspects of context, and also account for why they must be initiated in a stack distinct from the salient discourse referents. Let me explain.

The upshot of the relevant discussion in §4.1.3 was that deictic discourse referents are not akin to (putative) merely informational discourse referents, since the latter are not genuinely present in context prior to the use of a definite description and instead must be accommodated. In order to intuitively illustrate this difference, compare the following.

(28) *Context: A dog walks into the room and the interlocutors notice.*

The dog is so cute!

(29) The tallest man on earth must have serious back pain.

There is, I claim, an intuitive difference with regard to how the antecedents of each definite description here are established. With regard to (29), it may be antecedently part of the common ground that there is a tallest man on earth, but nonetheless a discourse referent for that man must be accommodated in order to fix an antecedent for the definite description. But, in (28), it seems odd to require that there must be a similar accommodation of a dog when the definite description is interpreted. The dog is clearly present to the interlocutors before the speaker comments on it.

Reflection upon contrasts such as the one between (28) and (29) is admittedly a subtle matter, and very different proposals have been made as to what the difference is. In her seminal work on dynamic semantics, Heim (1982: §2.2) considers a view according to which the discourse referents for the definite descriptions in both (28) and (29) must be accommodated. But I reject that suggestion. The main reason I instead include deictic discourse referents, such as the one in (28), on a par with salient ones—in that they are not accommodated—is that it helps explain the otherwise mysterious infelicity of the attempted deictic uses of definite

descriptions in (1)/(16) and (2)/(17). Given that a definite description “the *F*” presupposes that there is a unique *F* discourse referent, the infelicity in those attempted deictic uses requires that there be multiple *F* discourse referents present. So, despite the fact that deictic discourse referents are not introduced by overt indefinites, they must not be such that they are simply accommodated when needed. For it is unclear why a number of discourse referents would be accommodated in a way that leads to the *infelicity* of a definite description.

Heim (1982: §2.3), moreover, also ultimately endorses a view according to which the discourse referent in (29) is accommodated, but not so the one in (28). She claims, however, that, while the latter discourse referent is not accommodated, it is also not initiated by an element of logical form. But I have chosen here to model in a simple way how deictic discourse referents are introduced into the dynamic context other than via accommodation: initiation by a covert element of logical form. Furthermore, providing an element of logical form, the deictic existential quantifier, also helps represent how such discourse referents must be initiated in a stack that is separate from the salience stack, since the update of the deictic stack is encoded in the formal semantics of the deictic existential. I leave to future work the question of whether the present approach or Heim’s view—that deictic discourse referents are neither accommodated nor initiated by an element of logical form—is ultimately correct.

Let me point out that I have not yet provided motivation for the view that salient and deictic discourse referents must be distinguished by two stacks in the context. The empirical motivation for this aspect of my proposal is the dissimilarity in the behaviour of demonstratives between the deictic and anaphoric contrasts, which was touched upon at the end of §4.1.3. I show how my proposal explains that behaviour in §4.2.2. The parts of my proposal relevant for the explanation are the semantics of demonstrations and demonstratives, which I turn to now.

I treat demonstrations as elements of logical form akin to indefinites; demonstrations are a special type of existential quantifier.¹³ What is special about demonstrations is that they introduce discourse referents to the deictic stack in a relational way: the new discourse referent stands in some relation to an existing deictic discourse referent. Basic demonstrations, which will only concern us for now, employ the identity relation: the new discourse referent

¹³Let me now briefly compare my account of demonstrations to the two other recent and similar proposals mentioned in the previous section. First, my account is very similar to that of Ebert et al. (2020), since they also use relational discourse referents with the identity relation; however, they do not implement a perspectival dimension in the way I do. Second, my account is in starker contrast with that of Stojnić et al. (2013), since for them demonstrations simply introduce discourse referents non-relationally, and also they do not implement different stacks of discourse referents.

is identified with an earlier discourse referent in the deictic stack. So, a demonstration may select one or another deictic referent and raise it to deictic prominence. The form of a basic demonstration is as follows (the squiggly arrow represents abbreviation).

$$(30) \quad \Delta_n \rightsquigarrow [dx|x = dx_n]$$

In (30), the referential expression “ dx_n ” picks out the n th discourse referent back in the deictic stack of an input context (where counting starts at 0). So, a demonstration “ Δ_n ” bumps up the n th back deictic discourse referent to the top of the deictic stack. See Table 4.3 for further illustration. Note that not every gestures that might accompany a demonstrative or definite description counts as a demonstration: the simple demonstrations that are the present focus are contrastive in that they select one or another deictic discourse referent, and establish it as the highest ranked deictic discourse referent.

$c_o[[dx Fx]]$	$c_o[[dx Fx]; [dx Gx]]$	$c_o[[dx Fx]; [dx Gx]; \Delta_1]$
{ $\langle \emptyset, \langle a \rangle \rangle,$	{ $\langle \emptyset, \langle c, a \rangle \rangle,$	{ $\langle \emptyset, \langle a, c, a \rangle \rangle$
$\langle \emptyset, \langle b \rangle \rangle,$	$\langle \emptyset, \langle d, a \rangle \rangle,$	$\langle \emptyset, \langle a, d, a \rangle \rangle,$
	$\langle \emptyset, \langle c, b \rangle \rangle,$	$\langle \emptyset, \langle b, c, b \rangle \rangle,$
	} $\langle \emptyset, \langle d, b \rangle \rangle$ }	$\langle \emptyset, \langle b, d, b \rangle \rangle$ }

Table 4.3: Basic Relational Update. Here is the effect of a demonstration: first there are two deictic discourse referents added to the open context, then a demonstration bumps up the first one introduced. In the background model, a and b are the only objects that are F , and c and d are the only objects that are G . The semi-colon stands for dynamic conjunction.

The dynamic effect of demonstrations works in tandem with the meaning of demonstratives to allow the latter to pick out one or another of the deictic discourse referents. But the same does not hold with the meaning of definite descriptions, since definite descriptions presuppose discourse uniqueness. Let me now present my account of the meanings of those expressions.

I propose that a demonstrative “that F ,” as a matter of its meaning, helps elaborate the information about the highest ranked F discourse referent, considering the deictic stack first. In contrast, definite description “the F ” requires that there is a unique F discourse referent, considering both the salient and deictic stack. More formally, here are the update boxes associated with uses of demonstratives and definite descriptions used in simple sentences.

- (31) The/That F is G .
- a. $[G(\textit{the}_F)]$
 - b. $[G(\textit{that}_F)]$

The update of (31-a) that the definite description contributes to is defined just in case there is a unique F discourse referent. If defined, the update eliminates all points where that unique object is not G . The update of (31-b) that the demonstrative contributes to is defined just in case there is some F discourse referent. If defined, all points are eliminated where the highest ranked F object, considering the deictic stack first, is not G . See Table 4.4 for further illustration.

$$\begin{array}{ccc}
 c_o[[[sx|Fx]]][[sx|Fx]] = c' & c'[[[G(\textit{the}_F)]]] & c'[[[G(\textit{that}_F)]]] \\
 \{ \langle \langle a, a \rangle, \emptyset \rangle, & \# & \{ \langle \langle a, a \rangle, \emptyset \rangle, \\
 \langle \langle b, a \rangle, \emptyset \rangle, & & \langle \langle a, b \rangle, \emptyset \rangle \\
 \langle \langle a, b \rangle, \emptyset \rangle, & & \langle \langle a, b \rangle, \emptyset \rangle \\
 \langle \langle b, b \rangle, \emptyset \rangle \} & & \}
 \end{array}$$

Table 4.4: Demonstrative and Definite Description Updates. Given the introduction of two salient F discourse referents in c' , the definite description's uniqueness presupposition is violated; thus, the update is undefined, which results in the context crashing (represented by "#"). The demonstrative, however, picks up the top ranked salient F discourse referent, since there are no deictic discourse referents (if there were any deictic F discourse referents, the demonstrative would pick up the highest ranked of them, regardless of what is in the salient stack). In the background model, a and b are F , and only a is G .

I have thus presented the core elements of my proposal, which is a formal semantics for definite descriptions, demonstratives, and demonstrations. See the appendix for the nuts and bolts. The system I have just presented is extensional, since intensionality is not required for the applications of central concern. I leave discussion and formal treatment of intensionality to §4.3 (the appendix contains the full system). I now turn to applying the proposal to account for the data discussed in the introduction and §4.1.

4.2.2 Application

I begin with the deictic contrasts, and work through slightly different data for the sake of variation.

- (32) *Context: There are a couple of cars in front of the interlocutors.*
- a. That car [pointing to one car] is beautiful.
 - b. #The car [pointing to one car] is beautiful.

Logical forms for these two sentences are as follows.

- (33) a. $[dx|C(x)]; [dx|C(x)]; \Delta_1; [B(that_C)]$
 b. $[dx|C(x)]; [dx|C(x)]; \Delta_1; [B(the_C)]$

The first two boxes of each logical form represents that two car discourse referents are in the deictic stack. The demonstration in (33-a) bumps one of them up for the demonstrative to pick up. In (33-b), on the other hand, the third box will crash the update because the uniqueness presupposition of the definite description is not satisfied.

Here is a case similar to (2), where there are multiple demonstrations and intrasentential shift of reference.

- (34) *Context: There are a couple of house plants on the table in front of the interlocutors.*
- a. John loves that plant [pointing to one plant] and Bill loves that plant [pointing to another].
 - b. #John loves the plant [pointing to one plant] and Bill loves the plant [pointing to another].

Logical forms for these two sentences are as follows.

- (35) a. $[dx|P(x)]; [dx|P(x)]; \Delta_1; [L(j, that_P)]; \Delta_0; [L(b, that_P)]$
 b. $[dx|P(x)]; [dx|P(x)]; \Delta_1; [L(j, the_P)]; \Delta_0; [L(b, the_P)]$

The first two boxes in (35-a) and (35-b) represent that there are two plant discourse referents are deictically present in the context. In (35-a), the demonstrations bump one and the other of these background discourse referents to the top of its stack for the demonstratives in turn to pick up. In (35-b), however, the update will crash when processing the fourth box because there is more than one plant discourse referent in the context. Note that, for the uses of definite descriptions here and in the case (33-b), the update with the definite description will crash even if there is no demonstration update. That result is welcome, since (34-b) and (32-b) are infelicitous even if no demonstration is used.

We may now consider the anaphoric contrast, which I vary as well.

- (36) A cat toyed with the mouse, and another cat looked on.
- a. ...This cat was ready to pounce.
 - b. ...#The cat was ready to pounce.

In the first two sentences of (36), two salient discourse referents for cats are established by the indefinites. The demonstrative in (36-a) felicitously picks up the second mentioned cat (and must do so), while the definite description in (36-b) has its uniqueness presupposition violated.

Now, recall my claim from the previous section that part of the main virtue of dynamic semantics is that a treatment of bound uses of definite expressions simply falls out from a treatment of their anaphoric uses. The standard extensional treatment of universal quantification in dynamic semantics relies upon indefinites introducing discourse referents within the scope of conditionals, where a conditional is understood in terms of conjunction and negation.

- (37) a. $\forall F : \phi \dashv\vdash [sx|Fx] \rightarrow \phi$
 b. $\phi \rightarrow \psi \dashv\vdash \neg(\phi \wedge \neg\psi)$

According to (37-a), a universally quantified sentence such as “Every man runs” is read as *every man is such that he runs*, with a property as the first argument of the universal quantifier, and a sentence as the second. The translation scheme in (37-a) then tells us that that quantified sentence is equivalent to a conditional with the antecedent as a salience existential opening a discourse referent for an arbitrary man, and the consequent saying of this arbitrary man that he runs. According to (37-b), a conditional is equivalent to the negation of the conjunction of the antecedent and negation of the consequent. Hence, given standard background dynamic notions of negation, conjunction, and existential quantification, the translation key in (37) gives universal quantification the following semantics.

- (38) $c[\forall F : \phi] = \{i \in c \mid \text{For all } j \in \{i\}[[sx|Fx]], \{j\}[\phi] \neq \emptyset\}$

Hence, updating with a universally quantified sentence narrows down the input context in the following way: given an input context c , updating with “ $\forall \mathcal{F} : \phi$ ” returns a context containing every information state i in c such that, for every result j of adding an appropriate discourse referent to i , j can survive the result of updating it with “ ϕ .”¹⁴

On the current dynamic approach, we should therefore expect that the way in which

¹⁴For more detail on how constructions involving universal quantification and conditionals are treated in the current system, see the work of Stone (1999) and Bittner (2009, 2011).

demonstratives are restricted in their anaphoric uses rears in the quantificational environment. And, in fact, we do observe an analogue of the anaphoric contrast in such an environment, by considering an instance of donkey anaphora, where there are multiple indefinites in the antecedent of a conditional.

- (39) a. Every poor farmer who has befriended a rich farmer asked that farmer for a loan.
 b. #Every poor farmer who has befriended a rich farmer asked the farmer for a loan.

Following (37), the universal quantifier contributes a salient discourse referent for a poor farmer in the quantificational context. So the definite description in (39-b) has its uniqueness presupposition violated, since the indefinite “a rich farmer” also contributes a salient discourse referent in the quantificational. But the demonstrative in (39-a) may (and must) pick up the rich farmer, since it is more recently mentioned.

Moreover, there is the following similarity among demonstratives in quantificational and anaphoric environments. As already shown, the exact same demonstration-sensitivity does not rearise in either environment as it does in the deictic cases. But there is a similar way of getting something close in both quantification and anaphora.

- (40) A man [the speaker points to a spatial location on her left] entered from stage left. Another man [the speaker points a spatial location on her right] entered from stage right. This man [pointing again the location on her left] was happier than this man [pointing again to the location on her right] (Charlow p.c.)¹⁵
- (41) *Context: A developer and her assistants are inspecting a house constructed as a model for a subdivision project. They are standing in the living room. (Roberts p.c.)*
- a. Every house will have a light by that door [pointing to one of the doors].
 b. Every living room will have that shelf [pointing to a shelf] closer to that shelf [pointing to another shelf].

Upon some reflection, it should be apparent that in (40) and (41) what is occurring is deferred ostension. In both there is required some kind of deictic basis for the demonstrations: a position in front of the speaker in (40), or a door or shelf that is an instance of the type that many more will be made of in (41). As is more commonly discussed in the literature, deferred ostension occurs as well in deictic uses of demonstratives (Elbourne 2008).

¹⁵Here we have an imitation of the linguistic means of tracking discourse referents in American Sign Language. See Schlenker (2011).

- (42) *A farmer looks out over his land. To his left is a field where a certain donkey, Dim, usually grazes, and to his right is a field where another donkey, Dull, usually grazes.*
- a. That donkey [pointing to the left field] is not too bright.
 - b. But, that donkey [pointing to the left field] is smarter than that donkey [pointing to the right field].

Though I do not extend it in this way here, my view can treat deferred ostension by generalizing the account of demonstrations given above in (30). On this generalization, relations other than that of identity may feature in the demonstration's relational update. So, for instance, a demonstration may introduce a discourse referent based on the relation *x is the inhabit of field y*, where there is an earlier deictic discourse referent for a field.

4.2.3 Pragmatics of the Referential System

Recall that in §4.1.1 I presented an argument from Blumberg (2020) for the uniqueness view of definite descriptions. The argument relied upon the pragmatic principle *Minimize Restrictors!*, which is that redundant restrictive material in a definite description leads to infelicity. Crucially, Blumberg applied this principle to demonstratives, by treating them as definite descriptions with additional covert restrictive material. I concluded discussion of Blumberg's argument there by noting a problem, based upon emotional deixis, for *Minimize Restrictors!* and the background uniqueness view of definite descriptions and demonstratives.

The plan for this subsection is to wrap up the development of my main proposal by first showing how we may appeal to a related pragmatic principle to retain the insight behind Blumberg's argument. Second, I show how this alternative pragmatics fares better with emotional deixis. I conclude by drawing upon similar considerations to suggest how the current approach may be extended to context-sensitive referential expressions other than demonstratives and definite descriptions.

Consider the following contrast (Percus 2006).

- (43) a. Both of John's eyes are open.
 b. ??All of John's eyes are open.

For (43-b), there is the sense that it is odd, despite the fact that it may be strictly speaking true. Given that John is a regular human with two eyes, why not say the stronger (43-a) instead? But (43-a) is only stronger than (43-b) with regard to its presuppositions. So the following

pragmatic principle seems to be at play sustaining the competition between “all” and “both.”

Maximize Presupposition! Take sentences $\lceil \phi[a] \rceil$ and $\lceil \phi[b] \rceil$, where a and b are lexical alternatives and where $\lceil \phi[a] \rceil$ has strictly stronger presuppositions than $\lceil \phi[b] \rceil$. In a (local) context c where $\lceil \phi[a] \rceil$ and $\lceil \phi[b] \rceil$ are equivalent and both have their presuppositions satisfied, $\lceil \phi[b] \rceil$ is infelicitous.

Roughly, with alternative expressions that are equivalent in content, where one has stronger presuppositions that are satisfied, a speaker should use the alternative with stronger presuppositions. If a speaker uses the expression with weaker presuppositions, then that use implicates that the speaker is at least uncertain about whether the stronger presuppositions are satisfied (Sauerland 2008). Hence, if the stronger presuppositions are certainly met, then oddness results. Another way this point may be put is that the use of the weaker expression ANTI-PRESUPPOSES the stronger presuppositions of the alternative expression (Percus 2006).

Maximize Presupposition! was first introduced by Heim (1991) in order to explain contrasts involving indefinites and definite descriptions.

- (44) a. The Sun is shining.
b. ??A Sun is shining.

If both definite and indefinite descriptions have the same existential content, and the latter have no presuppositions while the former presuppose uniqueness, then Maximize Presupposition! may be applied. Accordingly, an indefinite anti-presupposes uniqueness. In (44-b), this anti-presupposition is violated because we are all aware that there is a unique Sun. On my semantics for indefinites and definite descriptions, the conditions apply Maximize Presupposition! to this case are satisfied. Hence, the current proposal follows Heim in explaining the oddness of (44-b).

Now compare definite descriptions and demonstratives. I repeat Wolter’s contrast, the explanation of which was a part of Blumberg’s argument from 4.1.1.

- (45) a. The author of *Blood Meridian* is a genius.
b. ??That author of *Blood Meridian* is a genius.

Again, Blumberg explains this contrast by appeal to Minimize Restrictors! and a uniqueness semantics for both definite descriptions and demonstratives, where the latter have additional covert restriction. Presently, however, we may appeal to Maximize Presupposition to explain

Wolter's contrasts. Demonstratives have the same existential content as definite descriptions. Furthermore, they have strictly weaker presuppositions than definite descriptions: a unique *F* discourse referent entails at least one *F* discourse referent, and not vice-versa. Hence, since both of these presuppositions are satisfied in (45-a) and (45-b), the latter is infelicitous because an alternative with a strictly stronger presupposition is not being used.

So, there are two plausible explanations of Wolter's contrasts on the table. But recall the problem noted earlier with Blumberg's use of Minimize Restrictors! alongside the uniqueness-based semantics for demonstratives and definite descriptions. The problem was that demonstratives but not definite descriptions participate in emotional deixis, yet this fact is mysterious if demonstratives differ from definite descriptions simply in having additional covert restrictive material. Here is the relevant contrast from before.

- (46) a. That US President is at it again!
 b. The US President is at it again!

On the view developed in this section, there is no covert restrictive material in definite descriptions (or demonstratives). So, even given Minimize Restrictors!, (46-b) is not predicted to have an interpretation where covert material plays a special pragmatic purpose.

So the present proposal explains the definite description's lack of emotive flavour in (46-b). But there are two questions that (46-a) raises for the present approach. First, given that (46-b) is felicitous, how is it that (46-a) is also felicitous, and not a violation of Maximize Presupposition!? Second, what is the source of the emotive flavour of the demonstrative in (46-b)? Let me begin with the first question.

The reason that one might think that (46-b) is a violation of Maximize Presupposition! is that there must be a discourse referent accommodated for the demonstrative, but, if there is only one relevant discourse referent accommodated, then the stronger uniqueness presupposition of the definite description in (46-b) is satisfied. In response, I suggest that for the demonstrative's felicity in (46-a) there are two relevant discourse referents accommodated. Accommodating more than one relevant discourse referent creates a context for the demonstrative in which Maximize Presupposition! does not predict its infelicity. This additional accommodation is possible in the case at hand, since there are multiple US Presidents (who held tenure at different times). And note how the use of the demonstrative is degraded if the overt material approximates semantic uniqueness, as in the following.

- (47) ??That current US President is at it again!

It is an odd suggestion that there are multiple *current* US Presidents, but that suggestion arises from the attempt to accommodate here multiple relevant discourse referents in order to avoid violating Maximize Presupposition! (a similar point holds for (45-b) above).¹⁶

It is also worth noting that there are cases where both demonstratives and definite descriptions are felicitous, and where the demonstrative does not give rise to emotional deixis. Here are two such cases.

- (49) a. *Context: A dog struts in front of the interlocutors.* The/that dog has tiny boots on.
 b. A dog chased a cat. The/that cat was very big.

Following the suggestion in the previous paragraph, I suggest that in both these cases the demonstrative is felicitous because an additional relevant discourse referent is accommodated. But what these cases also show is that the additional accommodation that is responsible for the felicity of demonstratives in cases that would otherwise violate Maximize Presupposition! is not itself responsible for the phenomenon of emotional deixis. So let me turn now to the second question, which concerns the emotive flavour the demonstrative in (46-a)—a task that will take us through the next subsection.

It is worth laying out first some more instances of emotional deixis.

- (50) *Context: A mechanic is assessing the state of a customer's car.*
 a. That left front tire is pretty worn.
 b. The left front tire is pretty worn.
- (51) *Context: A nurse is speaking to a patient.*
 a. How's that throat?
 b. How's the throat?

One way we might go is in locating the source of the phenomenon in the demonstration-

¹⁶In this footnote, let me mention an issue with the point just made about how accommodation can help make demonstratives felicitous by altering the context so that Maximize Presupposition! does not apply. Demonstratives may be felicitously used in a double possessive construction, as in (48-a).

- (48) a. That mother of John's is a real piece of work!
 b. ??That mother of John is a real piece of work!

It is hard to see how it may be reasonably accommodated that John has multiple mothers. In response, however, I simply note that the double possessive is a strange construction, and I do not have anything to say about it here. And note how the felicity of (48-a) does seem to depend on that construction, since (48-b) is degraded.

sensitivity of demonstratives. On the recent conception of demonstrations from Roberts (2020), a paradigmatic demonstration locates the referent in physical space, relative to the interlocutors. Yet Roberts argues that that same structure may be abstracted to other types of space. With an anaphoric demonstrative, for instance, one might be locating the relative distance of the antecedent in textual space; and it may be that epistemic modals are accompanied by something like a demonstration: when one says it must be raining one is in part sharing a perspective that locates the proposition that it is raining in epistemic space. So, in the cases of emotional deixis, it may be that the speaker signals the relative location of the referent in emotive space.

It is unclear, however, how my account of the semantics of demonstrations can be related to this conception. I do not represent semantically different types of space and indications of relative position through them. And I do not think that anaphoric demonstratives may be accompanied by demonstrations. On my view, demonstrations simply raise a deictic discourse referent to deictic prominence. Though it is not implausible that my view could make sense of how emotional deixis arises from demonstrative-sensitivity. It may, for instance, be that the speaker, in bringing one deictic discourse referent to prominence, indicates how emotively charged things are.

The proposal that the demonstration-sensitivity of demonstratives is responsible for emotional deixis has some initial plausibility, despite the fact that much more must be done to refine it. But, even at the impressionistic level, it faces the following problem. Pronouns with gender features are demonstration-sensitive in the same way as demonstratives, as the following shows (Maclaran 1982).

- (52) *Context: John and Mary are looking at an old class photo, and John is sharing his memories.*
 He [pointing to one kid] was my best friend, and he [pointing to another kid] was my mortal enemy.

But third-person pronouns with gender features do not give rise to emotional deixis, which suggests that demonstration-sensitivity is not the source of emotional deixis. Compare the following.

- (53) *Context: The interlocutors are waiting in line at a coffee shop, while the customer before them is harassing the barista.*
- a. That guy is a real piece of work!
 - b. He is a real piece of work!

The deictic features of demonstratives play a crucial role in emotional deixis, and pronouns do not have deictic features. So far, however, I have only treated the demonstration-sensitivity of demonstratives. The semantic nature of their deictic features involves introducing intentionality into the core system presented in this section, and I do so in the next section. My central proposal, however, is summarized in Tables 4.5 and 4.6.¹⁷ I now briefly conclude this section by suggesting how the semantics of pronouns and simple demonstratives may be treated. The results of these final suggestions are also represented in Table 4.6.

I do not address pronouns in any detail, but I briefly suggest, following Roberts (2005b), that, for instance, “he” has essentially the same semantics as “that man”—except that demonstratives have deictic features, which I discuss below. This proposal is motivated by the simple observation that third-person pronouns with gender features may be used alongside demonstrations in the same way as demonstratives, as observed already in (52). Note, however, that the third-person pronoun “it,” which lacks gender features, does not seem to be demonstration-sensitive.

- (54) *Context: John is directing Mary as they are working together rearranging his living room.*
 #You will take it [pointing to one chair] and I will take it [pointing to another chair].

Thus, I propose, following again Roberts (2005b), that “it” must be pick up the most salient inanimate discourse referent. Since demonstrations do not raise their discourse referent to salience—instead, to the top of the deictic stack—the non-demonstration-sensitivity of “it” is explained. Finally, I have not mentioned SIMPLE demonstratives, which are ones without any nounphrase complement. In this chapter, I focus on COMPLEX demonstratives, which do have a complement nounphrase. But I suggest that simple demonstratives simply pick up the top-ranked discourse referent, with no conditions on what kind of thing it represents, considering the deictic stack first.

4.3 Intensional Extensions

With the importance of treating the deictic features of demonstratives emerging from the previous section, I turn now to my account of them as contributing not-at-issue, perspective-sensitive updates. Doing so involves incorporating intensionality into the core system that is already on the table. I discuss afterwards the alleged difference between demonstratives and

¹⁷Note that on my view demonstratives have additional presuppositions and dynamic effects based upon their deictic feature that are not presented in the Table 4.6. I cover these in the next section.

<i>Type</i>	<i>How is it Introduced?</i>
Deictic	In virtue of the mutual perceptual environment of the interlocutors.
Salient	In virtue of an overt linguistic indefinite.

Table 4.5: Discourse Referents

<i>Type</i>	<i>Presupposition</i>	<i>Dynamic Effect</i>
Complex Demonstrative (“this man,” “that man”)	There is some discourse referent whose value is (e.g.) male throughout the context.	Contributes to an atomic update concerning the top ranked discourse referent whose value is (e.g.) male throughout the context, considering the deictic stack first.
Simple Demonstrative (“this,” “that”)	There is some discourse referent.	Contributes to an atomic update concerning the top ranked discourse referent, considering the deictic stack first.
Demonstration (a pointing gesture)	There is some discourse referent in the deictic stack.	Bumps some selected deictic discourse referent to the top of the deictic stack.
Definite Description (“the man”)	There is unique discourse referent whose value is (e.g.) male throughout the context.	Contributes to an atomic update concerning the unique discourse referent whose value is (e.g.) male throughout context.
Demonstrative Third-Person Pronoun (“he”)	There is some discourse referent whose value is (e.g.) male throughout the context.	Contributes to an atomic update concerning the top ranked discourse referent whose value is (e.g.) male throughout the context, considering the deictic stack first.
Non-Demonstrative Third-Person Pronoun (“it”)	There is some salient discourse referent whose value is inanimate throughout the context.	Contributes to an atomic update concerning the top ranked salient discourse referent whose value is inanimate throughout the context.

Table 4.6: Semantic Entries

definite descriptions with regard to rigidity. The formal system in the appendix includes the intensional resources for the applications here as well as that of the previous section.

4.3.1 Deictic Features

Consider again a simple contrast between demonstratives and definite descriptions mentioned at the end of the previous section.

- (55) *Context: A dog struts in front of the interlocutors.*
- a. The dog has tiny boots on.
 - b. That dog has tiny boots on.

The present reason for mentioning this contrast is that illustrates that demonstratives but not definite descriptions have deictic features. The demonstrative “that” encodes distality, in contrast to the proximal demonstrative “this.”

Here are some empirical details concerning how demonstratives and definite descriptions differ with regard to deictic features (Diessel & Coventry 2020, Rubio-Fernandez forthcoming). Demonstratives appear in all known languages and are both ontogenetically basic and phylogenetically old: they are ancient and among the first words children learn. Furthermore, cross-linguistically definite descriptions arise from demonstratives through GRAMMATICALIZATION, a process of language change whereby a syntactic form changes or loses some of its meaning. Compare again the demonstrative and definite description in (55). Both have as a part of their meaning a QUALITATIVE FEATURE, contributed by the complement noun-phrase “dog,” which indicates what kind of thing the referent is. The semantic effect of this qualitative feature has been modeled in the system of the previous section as constraining the anaphoricity of each expression. In addition, however, the demonstrative in (55-b) has a DEICTIC FEATURE, which indicates the referent’s location relative typically to the speaker. All known languages have a deictic system. In some languages, the system merely encodes the proximal/distal contrast. English is one such language, though we use to have an intermediate “yon.” More sophisticated deictic features are syntactically encoded in other languages: for instance, the relative elevation of the referent or whether it is within view of or hidden from the interlocutors (e.g. Khasi, West Greenlandic, Lahu, and Tauya).

My analysis of deictic features holds that they contribute dynamic updates as a part of the semantics of demonstratives. And their dynamic effect has two special aspects: perspective-sensitivity and not-at-issueness. Demonstratives, in virtue of having deictic features distinct

from qualitative features, are perspective-sensitive expressions in they have content that specifies something's position relative to some agent, typically the speaker. I draw here upon Roberts's (2020) notion of perspective-sensitivity, according to which a perspective-sensitive expression is one that presupposes a *DISCOURSE CENTER*: an agent whose (doxastic) perspective is salient. Typically the discourse center is the speaker, and for the purposes of this chapter I only consider the typical case.

Roberts's insight about perspectival presuppositions, however, only goes so far in accounting for the dynamic effect of demonstratives. A demonstrative, in virtue of its deictic feature, adds its own particular content to the context. But the special behaviour of this content does not seem to be fully a presuppositional matter, since the content is typically novel. The distal demonstrative "that," for instance, is typically used to help the audience identify an object as the one the speaker is talking about in virtue of its relative distance from the speaker. In this paradigm case, new information is being provided to the audience. Hence, it is unnatural to treat the overall content as presuppositional.

So, the content contributed by deictic features is dissimilar to presuppositional content in that it is typically novel. But, though typically adding information to the context, the content deictic features provide is *NOT-AT-ISSUE*: backgrounded relative to the main point of an uttered sentence containing them. In response to (55-b), for instance, it is much easier to deny that the dog in question has tiny boots on than that it is not relatively distant from the speaker. Thus the deictic content of demonstratives seems similar to that of Potts's (2005) conventional implicatures, which suggests that deictic features are similar to non-restrictive relative clauses, or appositives, which are paradigm cases of expressions that generate conventional implicatures. Compare the following (Potts 2005).

- (56) a. Chuck, who killed a coworker, is in prison.
 b. Chuck killed a coworker and is in prison.

The non-restrictive relative clause in (56-a) contributes the same information as the first conjunct of (56-b), but in a way that is backgrounded from the main point of the utterance.

Murray's (2014) recent treatment of appositives within dynamic semantics treats them as contributing novel, not-at-issue content. I follow Murray in modeling the distinction between at-issue and not-at-issue content in terms of different ways of updating the *CONTEXT SET*: the information in the common ground. Hence, in order to present my account of deictic features, the dynamic system I have presented so far must be enriched to incorporate intensionality. Before turning to the formal details, however, here are a couple initial points about how I

am drawing on Murray's treatment of appositives. First, there is another recent and similar dynamic account of appositives, from Anderbois et al. (2015). But Murray develops her account in the same background semantics from Bittner that I have used in the previous section, so that is why I take her account as my starting point. Second, and more pressingly, there are criticisms of the views of both Murray (2014) and Anderbois et al. (2015) on the basis that appositives may sometimes be at-issue, while both views predict that they cannot be at-issue. Syrett & Koev (2015) present compelling empirical evidence that relative clause appositives may in some cases be at-issue. The alternative view they outline, as well as the account of appositives in Martin (2016), allows for at-issueness. There are, therefore, genuine issues with the proposal I am drawing upon as it relates to relative clause appositives. But, I suggest, the deictic content of demonstratives is a kind of supplemental content that cannot be made at-issue. Hence, the application of Murray's proposal in this chapter is not susceptible to those issues. And note that here is another point of contrast between the content of deictic features and presuppositions: presuppositional content may vary in its at-issueness (Simons et al. 2010), but not so the content of deictic features.

Moving on, I follow Bittner and Murray in incorporating intensionality by introducing worlds and propositions into the dynamic economy of discourse referents. Accordingly, a context is an "infotention" state that represents rankings of discourse referents, where some discourse referents track intensional information (Bittner 2009: 2). Hence, I introduce discourse referents for worlds and propositions, and with corresponding variables make updates sensitive to these discourse referents in various ways. For the at-issue content of a declarative sentence—the only type of sentence considered in this chapter—a propositional discourse referent is initiated in the deictic stack, and it is then proposed to update the context set with the information in that initiated propositional discourse referent. Thus, it must be ensured that a context for processing a sentence is set up with a discourse referent representing the context set. Following Bittner (2011), what ensures this setup is Stalnaker's ([1978] 1999) COMMONPLACE EFFECT: when a speaker begins speaking, the context relevant for processing the main content of the sentence will be guaranteed to include information and discourse referents that can be mutually inferred on the basis of recognition that a particular a speech act is taking place. So, when a speaker begins speaking, the context is updated with a discourse referent representing the context set, which may then be narrowed down in various ways by the uttered sentence.¹⁸

¹⁸I set aside the question of how the present appeal to the commonplace effect and the context set are related to my appeal to those notions in the previous chapter.

An example of the addition of a context set and the simple effect of at-issue content is represented in the following series of updates, where the sentence providing the at-issue content is “John is happy.”¹⁹

$$(57) \quad [s\omega|\omega \in CS]; [d\omega|H_\omega(j)]; [d\Omega|\Omega = d\omega_0||]; [s\omega_0 \in d\Omega_0]; [s\Omega|\Omega = s\omega_0||]; [s\omega|\omega \in s\Omega_0]$$

The first update box represents the commonplace effect: a world discourse referent is introduced that is indefinite between the worlds in the context set, where the latter is for simplicity referred to with the constant “CS.” In the second two boxes, world and propositional discourse referents are introduced, which work together to represent the main content of the uttered sentence. What happens is that a world discourse referent is introduced that is indefinite between all the worlds where John is happy, then in the next box a definite propositional discourse referent is introduced that stands for that entire set of worlds. Note that the way that definite propositional discourse referent is constructed is with a global update, which considers an entire input context. Crucially, a global update relies on a new kind of referential expression, “ $\alpha\omega_n||$,” that denotes all the different world values in the n th position back of the input context. In the final three boxes, which represent the speech act force associated with a declarative sentence, the main content of the sentence is added to the context set. The way it is added is that the world discourse referent representing the information in the context set, which was initiated in the first update box of the entire logical form, is narrowed down in such a way as to eliminate any worlds that are not consistent with the main content of the sentence, which again is represented by the definite propositional discourse referent initiated by the third box of the entire logical form. Then, in the final two boxes, a new definite propositional discourse referent is initiated that includes all the worlds in the original context set in which John is happy, and then a new world discourse referent is initiated so that a similar update process may be undertaken by a subsequent sentence.

So what the logical form in (57) does is construct a piece of information, the worlds in which John is happy, and then the last three boxes, again representing the speech act force, executes the proposal to update the context set with that piece of information. Since there is a proposal to add that information to the context set, it is possible for a hearer to reject that content in a way distinctive of at-issueness. In contrast, the way that not-at-issue content is modeled is in directly updating the context, and that is why it cannot so easily be challenged. In particular, demonstratives supply a not-at-issue update in virtue of their deictic

¹⁹See Murray (2014: §3) for more detailed discussion of similar examples.

features. And given perspective-sensitivity, in order for that update to be defined there must be a discourse referent associated with the speaker—so that the perspectival presupposition is satisfied. So, with the background commonplace effect in place, before the main content of a sentence is processed, the context must be additionally updated with the appropriate deictic discourse referents, including at least the speaker. The speaker has a corresponding deictic discourse referent, since, at least in ordinary conversation, the speaker is an aspect of the interlocutors' mutual perceptual environment. So, the initial context for the processing of the main content of an utterance can be represented as follows, where $c_o = \{\langle \emptyset, \emptyset \rangle\}$ is the open context.

$$(58) \quad c_o \llbracket [s\omega | \omega \in CS]; [dx | speaker_{s\omega_0}(x)]; [dx \dots] \rrbracket$$

The final update represents schematically how there may be deictic discourse referents in addition to the one for the speaker.

Let me now return to the simple sentence (55-b) containing a demonstrative with which this subsection began: “That dog has tiny boots on.” First, for this sentence, the commonplace effect must be as follows.

$$(59) \quad c_o \llbracket [s\omega | \omega \in CS]; [dx | speaker_{s\omega_0}(x)]; [dx | D_{s\omega_0}(x)] [dx \dots] \rrbracket$$

Then, as the following shows, the logical form of the sentence constructs an at-issue-proposal, but, before updating the context set with that proposal, the context set is directly updated by the demonstrative's deictic feature.

$$(60) \quad [d\omega | TB_\omega(that_D)]; [distal_{s\omega_0}(that_D, the_{speaker})]; [d\Omega | \Omega = d\omega_o]; \dots$$

What is crucial here is that the second update box represents how the demonstrative “that” has the deictic feature of distality, which directly adds not-at-issue information to the context. I use the referential expressions “*that_D*” and “*the_{speaker}*” in the specifying the deictic feature to simply represent how the deictic update is perspective-sensitive in that it presupposes there is a speaker, and if defined contributes the content that the relevant thing is distal to the speaker. I do not hold that deictic features have demonstratives and definite descriptions as a part of their syntax. Note finally that the material I have elided in (60) is simply the addition of the at-issue context to the context set, which was shown in (57) in full.

In addition to being difficult to reject, not-at-issue content is also projective, as the following illustrates with regard to the content of deictic features.

- (61) *Context: The interlocutors are arguing about the breed of the dog they see.*
That dog isn't a poodle.

The sentence (61) cannot be used to deny that the relevant dog is distal to the speaker, which is to say that deictic content projects past operators such as negation. The way this projection is accounted for, in the case of negation, is that in negation the context set is updated to include worlds that are not a part of the at-issue information (Stone & Hardt 1999, Bittner 2011, Anderbois et al. 2015). This process is illustrated in the following logical form for (61) (which I assume has a background context with the commonplace effect and the introduction of speaker and dog discourse referents).

- (62) $[d\omega|P_\omega(\text{that}_D)]; [distal_{s\omega_0}(\text{that}_D, \text{the}_{speaker})]; [d\Omega|\Omega = d\omega_0||]; [s\omega_0 \notin d\Omega_0]; [s\Omega|\Omega = s\omega_0||]$

In the first box, a world discourse referent is initiated that is indefinite between the worlds in which the dog discourse referent is a poodle. Then, in the second box, before a definite propositional discourse referent is constructed for that set of worlds in the third box, the dynamic effect of the demonstrative's deictic features takes place, directly updating the context set. In the penultimate box is the effect of negation: context set worlds are eliminated that are not ones in which the prejacent holds. In the final box, a new context set discourse referent is initiated.

Finally, with an account of deictic features in place, we may return to the phenomenon of emotional deixis. Crucially, observe that a deictic property is very general. What it means, for instance, for something to be distal is not semantically specified as involving, say, relative physical distance; rather, the property of distality only semantically specifies a very general property of distality in some type of space. This loose semantic content is enriched in more particular ways in particular contexts as a broadly pragmatic affair. That deictic features are semantically nebulous is supported by various empirical studies that illustrate how factors such as relative control, visibility, and ownership are relevant to how language users interpret distal and proximal demonstratives (Coventry et al. 2008, 2014, Peeters & Özyürek 2016, Caldano & Coventry 2019). Hence, what explains emotional deixis is that certain uses of demonstratives have their deictic updates pragmatically interpreted as concerning emotional distance. This type of enrichment occurs when the deictic content is not required to be interpreted in a way that helps identify the referent, for instance as involving physical distance, since the overt qualitative feature of the demonstrative phrase provides suffi-

cient identifying information. This explanation of emotional deixis is similar to that given by Acton & Potts (2014), who provide the most thorough discussion of this generally underappreciated phenomenon. They hold that because demonstratives, as opposed to definite descriptions, are semantically perspective-sensitive, the use of a demonstrative signals that the addressee should consider the speaker's perspective, which in certain cases may give rise to a signal that the speaker is attempting to establish solidarity with the addressee. So, my view is a refinement of their suggestion: I hold that emotional deixis is a pragmatic phenomenon based on the semantic perspective-sensitivity of demonstratives. Note, however, that Acton & Potts' discussion is potentially misleading because they do not distinguish deictic features and their perspective-sensitivity from the demonstration-sensitivity of demonstratives. But, as we noted in §4.2.3, some third-person pronouns are demonstration-sensitive, yet they do not give rise to emotional deixis.

4.3.2 Rigidity and Existentialism

Following Kaplan's (1989) seminal work, one might think that a discussion of demonstratives and definite descriptions should place, front and center, their differences with regard to rigidity. Kaplan's influence supports a picture on which (i) demonstratives are always rigid, (ii) definite descriptions never (unless special material is added), and (iii) definite descriptions can interact with intensional operators in ways that demonstratives cannot. But this picture is oversimplistic, and I have left this topic to the end because these alleged differences are not as robust as the traditional wisdom would have it.

Kaplan's master argument for (i) and (ii) relies upon intuitions about the truth-conditions of simple sentences in context. Compare the following.

- (63) *Context: There is a man, m , wearing a hat in front of the interlocutors.*
- a. That man [pointing at m] is bald.
 - b. The man wearing a hat is bald.

Consider a possible world w (distinct from the world in which the interlocutors are interlocuting) in which (a) instead of m , there is another man m' in front of the interlocutors wearing a hat, (b) m' is bald in w , and (c) m is present somewhere else in w (far away and not wearing a hat), and is not bald in w . The intuitive contrast is that what is said by (63-a), as uttered in its actual context, is false in w ; whereas, what is said by (63-b), as used in its actual context, is true in w .

For a proposition expressed, or what is said, by a simple sentence to be RIGID, where the sentence contains a demonstrative or definite description, and is considered in a context c , is for the proposition's truth at any possible world w to depend only on the properties at w of the individual picked out in c . So the proposition expressed by (63-a) is rigid, but not so the one expressed by (63-b). Kaplan proposes that in general the proposition expressed by a simple sentence in context containing a definite description is never rigid—except when special material is added, as with the actuality operator in the following.

(64) The *actual* man wearing a hat is bald.

Recall the Kaplanian “equation” (11) in §4.1.1 of the relationship between demonstratives and definite descriptions. On that view, demonstratives are rigidified versions of corresponding definite descriptions. Accordingly, on the view, the proposition expressed by a simple sentence in context containing demonstratives is always rigid.

It is worth mentioning that these intuitions about the different truth-conditions of (63-a) and (63-b) are subtle. Another perhaps more intuitive way of putting the difference is in terms of aboutness, or ways of referring. The demonstrative in (63-a) seems to be about the particular man actually picked out in its context of use: it is *that* man's baldness (or lack thereof) relevant to the truth-value of what is said, even when assessed at other possible ways things might have been. In contrast, the definite description in (63-b) is about whoever is the man wearing a hat, which picks out m in the interlocutors' actual context, but which may vary from world to world. So with demonstratives there is direct reference; definite descriptions, reference via description.

But even granting that these intuitions are legitimate, similar ones can be harnessed to create problems for Kaplan's generalization. First, there are non-rigid uses of demonstratives, as noticed by King (2001).

(65) *Context: A group of corrupt senators are together at a bar. They learn that exactly one senator, though they do not know which, is going to testify the next day to an investigative committee and likely reveal their illegal activity. One of the senators says the following.*

I am going to kill that senator.

Intuitively, what is said by (65) is non-rigid. The speaker is talking about whichever senator is the one that is going to testify. Second, there are rigid uses of definite descriptions (without the aid of special material). Here is case noticed by Fara (2015).

- (66) *Context: The conversational participants are discussing a party they all attended last night.*
 Olga enjoyed the party.

Intuitively, the speaker of (66) is saying of the particular party that it was enjoyed by Olga. We should thus conclude that both definite descriptions and demonstratives partake in VARIABLE RIGIDITY: depending on the context, and without the aid of special material, both sometimes lead to rigidity and sometimes not.

The subtlety of the intuitions about what is said by simple sentences has also lead some to focus instead on comparing embedded cases of demonstratives, which is related to point (iii) of the Kaplanian picture limned at the beginning of this subsection. More thorough and recent work, however, has shown that demonstratives may interact with operators in ways similar to definite descriptions (Heim 1985, Nunberg 1993, Elbourne 2008). I set aside discussion of embedded cases, since what has been set forth is already enough to discuss the more radical way in which my proposal differs from the Kaplanian approach to the semantics of referential expressions.

The considerations just given with regard to rigidity involved appeal to the proposition expressed, or what is said, by a sentence in context. These notions are supposed to be equivalent for Kaplan, and both in line with the truth-conditions of a sentence in context. So, in his two-dimensional system, expressions are assigned CHARACTERS: functions from contexts to content. Given a context, CONTENTS compose to determine truth-conditions. A reader who has been deeply immersed in Kaplan's work may realize that—contrary to Kaplan's insistence otherwise—the content compositionally generated, on the one hand, and what is said in context, on the other, in fact cannot be equated. Partly for this reason, a more modern version of Kaplan's framework has it that characters compose to determine a character for a sentence, which then may be combined with a context to determine a proposition.²⁰

Thus, given that in Kaplan's considered system sentences are provided functions from contexts to propositions, there is in fact some formal similarity between it and the dynamic view developed in this chapter. In both frameworks, sentences are semantically provided values more complex than propositions, and, in particular, these semantic values of sentences are functions from contexts to propositions. But there are crucial differences in the formal and conceptual nature of the input contexts. Setting aside differences in the elements of the information points, here is a high-level analysis of the formal differences (entities of type *s*

²⁰See Ninan (2010), Rabern (2012, 2013), Stalnaker (2014), and Yalcin (2014) for the motivation for modernizing Kaplan's view, though the core observation goes back to Lewis (1980). See also Glanzberg & King (2020) for a defence of classic Kaplan, but also Pickel & Rabern's (forthcoming) rejoinder.

are information points, which in both systems are ordered n -tuples; type t is that of the two truth-values).

- (67) a. Kaplan: $s \rightarrow (s \rightarrow t)$
 b. Dynamic: $(s \rightarrow t) \rightarrow (s \rightarrow t)$

In Kaplan's system, sentence meanings are functions from points to sets of points. In the dynamic system here, sentence meanings are functions from sets of points to sets of points.²¹ So the formal difference between the input contexts allows that, for Kaplan, context supplies a fully determinate way things are as well as particular values for context-sensitive referential expressions. For dynamic meanings, in contrast, the input context is a set of points, so there may be many candidates for actuality, and it may be indeterminate which value a referential expression receives. The resulting conceptual difference is that a Kaplanian input context may be taken to be the actual determinate milieu of the interlocutors, so that the proposition determined may be reasonably taken to be a truth-condition. In contrast, the input context for a dynamic meaning, given its indeterminacy, is naturally taken to be a state of information had by the interlocutors: the context set. And it is strange for a truth-condition to be something relative to a state of information, since whether what is said is true or false does not depend on the information interlocutors happen to currently share (Stalnaker 2018). So the set of points output by a dynamic meaning conceptually should not be regarded as the truth-condition of the sentence in context.

Despite these differences, it is nonetheless possible to find truth-conditions of sentences in a dynamic semantic framework. Consider first how in static two-dimensional system such as Kaplan's, where sentence meanings are of type $s \rightarrow (s \rightarrow t)$, there is the following notion of a diagonal proposition (Stalnaker [1978] 1999). Given that one element of the n -tuples of type s is a world, we may retrieve a non-context-sensitive proposition from a sentence's meaning by considering first the set S of all n -tuples such that, when that n -tuple is fed in *twice* to the sentence's meaning, the output is truth. The set of all worlds that appear in some element of S is the diagonal proposition.

We may employ a formally analogous tool with dynamic meanings.²² Take a sentence S , which corresponds syntactically to a sequence of update boxes. Consider now the set TC_S

²¹In Kaplan's terms, a sentence meaning is a character, which is a function from a context—an n -tuple containing a world, time, place, and agent—to a proposition—a function from such n -tuples to truth-values. For that reason, in the terms of the paragraph above, a character is a function of type $s \rightarrow (s \rightarrow t)$.

²²This method is adapted from the work of Heim (1982), Starr (2010), Bittner (2011), Murray (2014), and Stokke (2014).

consisting of worlds w which satisfy the following: if the singleton $\{w\}$ is taken as the initial context set in the commplace effect, the result of updating that context with the update boxes does not result in a crash or the absurd state \emptyset . The set TC_S is the truth-condition of S . Informally, this method works by feeding in fully determinate context sets, and seeing which ones survive. Since the context sets are fully determinate—singleton sets of worlds—it is appropriate to speak of truth-conditions. This method of extracting a truth-condition provides a non-context-sensitive, existentialist propositional content for sentences containing referential expressions: for instance, a sentence involving a deictic demonstratives has the propositional content that there exists something with such-and-such properties. I call *EXISTENTIALISM* the view that referential expressions such as demonstratives and definite descriptions merely contribute, as a matter of their semantics, to claims that such-and-such properties are co-instantiated in the world.

Thus, returning to the considerations of rigidity, the existentialism of my current framework means that those Kaplanian considerations of reference are not in fact an aspect of semantics. This result should be welcome in light of how Kaplan's traditional generalizations (i) and (ii) do not in fact hold: there is not a robust difference between demonstratives and definite descriptions with regard to rigidity. But existentialism also has troubling consequences, since we do use expressions such as demonstratives and definite descriptions to talk about particular objects in our world. So, the issue arises of how to account for the notion of reference that seems to outstrip what is represented in dynamic semantics. I now critically discuss two extant ways of avoiding existentialism before discussing how to deal with its consequences.

Following a natural suggestion, recently made explicit in Rothschild & Yalcin (2017), one might avoid existentialism about deictic demonstratives (and other context-sensitive expressions) by enriching the semantic types even further than the dynamic ones in (67-b). On this view, a sentence's meaning takes as input both a determinate Kaplanian context (type s) and an indeterminate dynamic one (type $s \rightarrow t$), and given that input supplies a proposition (another indeterminate context). Accordingly, a deictic demonstrative may receive a determinate value from the determinate Kaplanian context, and thus be about some particular object. The problem with this proposal is that it must treat deictic and anaphoric demonstratives as ambiguous, since anaphoric demonstratives pick up an indeterminate value from the dynamic context. Consider the following.

- (68) a. That man [pointing to a man] was happy.

- b. A man came in. That man was happy.

Positing that the demonstratives in (68-a) and (68-b) are of different expression-types, however, is undesirable.

The second way of avoiding existentialism is from Stojnić et al. (2013, 2017), and does not require positing ambiguity or enriching dynamic meanings. They exploit a technical point to give a unified treatment of deixis and anaphora in which all of demonstratives are anaphoric to dynamic existentials, yet in which deictic ones contribute to singularist truth-conditions. The technical point is that existentially quantified sentences can have the same truth-conditions as ones only containing singular terms. Compare the following equivalent sentences of simple predicate logic.

- (69) a. $\exists x(x = n \wedge Fx)$
 b. Fn

In these sentences, “*n*” is a non-context-sensitive name. Hence, with the use of singular restrictors—or definite discourse referents, as they are called in dynamic semantics—Stojnić et al. succeed in treating deictic demonstratives as anaphoric, but at the same time contributing to singularist truth-conditions. But this way of avoiding existentialism is underexplanatory. The proposal seems to be that the source of context-sensitivity is in fact different covert constants (where there must be enough of such constants in order to refer to any object we could refer to with demonstratives). One might attempt to make this account explanatory by claiming that a demonstration really involves the coining of a new proper name. While such a proposal cannot be ruled out a priori, it is empirically implausible.

Given the issues with the two extant ways of avoiding existentialism, I suggest we learn to live with it. And, fortunately, the difficulty is dispelled by the distinction between sentence and utterance understanding: there is a difference between what is involved in understanding an uttered *sentence* and understanding an *utterance* of a sentence.²³ Now note that Bittner (2001, 2009), whose formal system I have relied on this chapter, argues that there are semantic universals of discourse attention covering not only the individual domain, but also the temporal and modal ones. I thus take it that the instantiation of these semantic universals throughout the languages of the world arises from a semantic component of linguistic competence, in Chomsky’s (1965) sense. In this chapter, I am only concerned with the individual domain. But

²³I discuss in more detail the present point in the introduction and conclusion, and there note its connection to other’s work.

in investigating the deictic and anaphoric contrasts we have seen how one manifestation of the semantic universals of discourse attention is in a distinction between deictic and salient discourse referents. With this distinction, we have hit upon what may be called a GRAMMAR OF REFERENCE, since there are linguistic constraints governing deictic uses of demonstratives, which are analogous to the ones governing overt linguistic anaphora. Accordingly, I have modeled how various elements of logical form contribute semantic updates to context, as well as how a truth-conditional sentential content may be retrieved from such updates. So I suggest that what is required for understanding an uttered *sentence* is represented by the sentence's semantic content, which corresponds to an aspect of our linguistic competence. But these constraints do not determine, for instance, what particular object a demonstrative refers to. Accordingly, it is in utterance understanding that one must think of this or that object as referent. So, for utterance understanding one must go beyond semantic content and into pragmatics, broadly speaking. Here is not the place to begin explaining how reference may find a home in pragmatics, but there is no reason to think such a project is hopeless. One may happily look to Strawson's (1950) insight that reference is something people—not words—do.

4.4 Looking Ahead

This chapter developed a treatment of demonstratives and definite descriptions within a dynamic semantics. The backbone of my account is a distinction between two types of discourse referents: deictic and salient. Given that the context keeps track of the different stacks of deictic and salient discourse referents, the semantic presuppositions and dynamic updates of context-sensitive referential expressions can be specified in a way that accounts for the similarities and differences in their behaviour between their (pre-theoretically speaking) deictic, anaphoric, and bound uses.

The existentialism defended at the end of the previous section connects the work here to much of the dissertation that came before. Deictic demonstratives themselves do not refer, so understanding demonstrative reference is not about understanding a linguistic expression, or even a token of one. It is part of understanding one another's actions, which is a part of pragmatics, broadly construed. In the conclusion to this dissertation, I discuss the connection between the existentialist conclusion of this final chapter and the Gricean considerations from the earlier in the dissertation.

Looking ahead, the main upshot of this chapter's treatment of demonstratives and definite

descriptions is in formally elaborating the connection between the semantics of those expressions and discourse structure. But I have assumed in this chapter an overly simplistic model of salience, according to which a discourse referent is added to the top of the salience stack by an overt indefinite, and never leaves the stack. Moreover, there can be no ties in ordering, and the ordering never changes unless a new salient discourse referent is introduced. This crude model of salience must ultimately be revised (and for similar reasons the dynamics of the deictic stack may end up being more complicated than I have modeled in this chapter), and I conclude by discussing the importance of doing so.

Looking ahead, the main upshot of this chapter's treatment of demonstratives and definite descriptions is in formally elaborating the connection between the semantics of those expressions and discourse structure. But I have assumed in this chapter an overly simplistic model of salience, according to which a discourse referent is added to the top of the salience stack by an overt indefinite, and never leaves the stack. Moreover, there can be no ties in ordering, and the ordering never changes unless a new salient discourse referent is introduced. This crude model of salience must ultimately be revised (and for similar reasons the dynamics of the deictic stack may end up being more complicated than I have modeled in this chapter), and I conclude by discussing the importance of doing so with regard to the behaviour of definite descriptions.

To begin, consider the following data.

- (70) The cat is in the carton. The cat will never meet our other cat, because our other cat lives in New Zealand. Our New Zealand cat lives with the Cresswells. And there he'll stay, because Miriam would be sad if the cat went away. (Lewis 1979)
- (71) A man in a dark suit was talking to a man holding a large envelope. After talking for about a minute they left the station going in opposite directions. Thirty seconds after they left, the man showed up and nervously looked around. (Szabó 2000)

In (71), one is able to accommodate a third man in order to felicitously utter the definite description in the final sentence, despite the fact that there seems to be two discourse referents for men initiated by the first sentence. Thus, (71) seems to be a counterexample to the view that definite descriptions presuppose discourse uniqueness. In response, however, I suggest that the complexities of salience make it so that the context relevant the definite description in the final sentence does not include those initial two discourse referents. And I suggest that a similar point holds for (70): the discourse referent that the definite description's in the first

two sentences are anaphoric to is no longer relevant in the final sentence.

Now, in light of these suggestions about how there is felicity in (70) and (71), I wish to consider the following problematic data.

(72) A woman₁ met another woman₂. The woman_{1,??2} was wearing a red coat.

The problem that (72) raises for the account definite descriptions I have here defended, on which they presuppose discourse uniqueness, is that the definite description in the second sentence can be felicitously used to elaborate the information about the first mentioned women in the first sentence, despite the fact that the second indefinite in the first sentence also introduces a discourse referent, which remains present for the context of the second sentence. The fact that the second indefinite in the first sentence introduces a discourse referent that remains present is supported by the felicity of the following.

(73) A woman₁ met another woman₂. That woman_{#1,2} was wearing a red coat.

In (73), the demonstrative may pick up the discourse referent introduced by the second indefinite of the first sentence, which means that it should still be present in the context.

Note furthermore that the following “bishop sentence” (Heim 1990), with definite descriptions, poses a similar problem to that raised by (72).

(74) If a bishop meets a bishop, the bishop blesses the bishop.

Here both definite descriptions are felicitous, despite the fact that there should be multiple bishop discourse referents in the local context of the second clause.

In response to the issues raised by data such as (72) and (74), I wish to say first and foremost that it is genuinely puzzling why the definite descriptions in those cases are felicitous, but not so the one in the anaphoric contrast, repeated here, that I have focused on in this chapter.

(75) A woman entered from the stage left. Another woman entered from stage right. #The woman was carrying a bouquet of flowers.

I have followed Roberts (2003) in appealing to (75) in supporting the view of definite descriptions according to which they presuppose discourse uniqueness; however, data such as (72) and (74) is typically used to support the competing familiarity theory, covered in §4.1.2, according to which definite descriptions simply presuppose there is some relevant discourse referent, and so can be co-indexed with this or that earlier indefinite. To my knowledge, how-

ever, there is existing account of the infelicity of the definite description in (75) on behalf of the familiarity theory, thus that approach has not resolved the present puzzle. Finally, I wish to suggest, on behalf the view that definite descriptions presuppose discourse uniqueness, that the felicity of (72) and (75) can be ultimately accounted for by appeal to the complexities of how the relative salience of discourse referents is introduced and maintained, along the similar lines as the suggestions above concerning (70) and (71).

.1 Formal System

.1.1 Syntax

Singular Terms There are three types: variables (existential quantifiers) $v \in V$, referential expressions $r \in R$, and global referential expressions $g \in G$. The metalinguistic variable n is any natural number, and α is one of s or d ; so, for instance, the set of variables V has a total of six members, as I have made explicit.

1. $V = \{\alpha x, \alpha \omega, \alpha \Omega\} = \{sx, dx, s\omega, d\omega, s\Omega, d\Omega\}$
2. $R = \{\alpha x_n, \alpha \omega_n, \alpha \Omega_n, the_F, that_F, CS\}$
3. $G = \{\alpha \omega_n ||\}$

Predicates A basic predicate $b \in B$ is an n -place predicate. There are also the logical predicates in $l \in L$.

1. $B = \{F, G, R, \dots, distal, speaker\}$
2. $L = \{=, \in, \subseteq\}$

Atomic Sentences Any basic n -place predicate with n referential expressions attached, the first one being $\alpha \omega_n$, is an atomic sentence: $F_{\alpha \omega_n}(r, \dots, r')$. Where l is a logical predicate, r, r' referential expressions, and g, g' global referential expressions: rlr', rlg , and glg' are atomic sentences.

Well-Formed Formulas (WFFs) Any variable or atomic sentence is a WFF. The conjunction of any WFFs is a WFF: $\phi; \psi$.

DRT-Style Abbreviations I use DRT-style boxes as abbreviations for conjunctions of introducing discourse referents and imposing conditions on them. The general translation scheme is as follows, where v is a variable and C, \dots, C' are atomic sentences.

$$(76) \quad [\alpha v, \dots, \alpha' v' | C, \dots, C'] \rightsquigarrow \alpha v; \dots; \alpha' v'; C; \dots; C'$$

A box may introduce no variables, so only consist of C, \dots, C' . The referential expressions in C do not have any subscript or α if it is obvious that they correspond to ones that are introduced by variables in the left hand side of their box.

.1.2 Semantics

Models A model M is a tuple $\langle D, W, I \rangle$. The first two members represent the basic entities of different types: D is the set of individuals; W , worlds. I is an interpretation function from n -place predicates to functions from worlds $w \in W$ to n -tuples of individuals in D .

Contexts A CONTEXT, or information state, c is a set of information points. An INFORMATION POINT i is a tuple $\langle \sigma_s, \sigma_d \rangle$, where the first sequence σ_s^i helps represent the salient discourse referents and the second sequence σ_d^i the deictic ones.

Operations on Sequences For a given sequence σ and natural number n : $n(\sigma)$ is the n th member of σ , where the counting starts at 0 and from the beginning of σ . For a given σ and set S , σ^S is the sequence containing elements of σ that are in S , and appearing in the same order as they appear in σ . Also, for a given σ and S , S^σ is the set of things occurring in σ that are S . Given sequences σ, σ' : their amalgamation $\sigma \uplus \sigma'$ is the sequence formed by combining σ and σ' in order.

Static Interpretation Given a background model, there is a static interpretation function $\llbracket \cdot \rrbracket_s$ that assigns meaning, given an information point i , to referential expressions and atomic sentences. There is also a global static interpretation $\llbracket \cdot \rrbracket_{sg}$ that assigns meaning, given a context c , to global referential expressions.

$$1. \quad \llbracket the_F \rrbracket_s^i = \lambda w. \begin{cases} a & \text{If } a = \iota x (x \in I(F, w)^{\sigma_d^i \uplus \sigma_s^i}) \\ \text{Undefined} & \end{cases}$$

2. $\llbracket \text{that}_F \rrbracket_s^i = \lambda w. \begin{cases} a & \text{If } a = \iota x(x = 0((\sigma_d^i \uplus \sigma_s^i)^{I(F,w)})) \\ \text{Undefined} & \end{cases}$
3. $\llbracket \alpha x_n \rrbracket_s^i = \lambda w. \begin{cases} d & \text{If } d = \iota x(x = n((\sigma_\alpha^i)^D)) \\ \text{Undefined} & \end{cases}$
4. $\llbracket \alpha \omega_n \rrbracket_s^i = \lambda w. \begin{cases} w' & \text{If } w' = \iota x(x = n((\sigma_\alpha^i)^W)) \\ \text{Undefined} & \end{cases}$
5. $\llbracket \alpha \Omega_n \rrbracket_s^i = \lambda w. \begin{cases} P & \text{If } P = \iota x(x = n((\sigma_\alpha^i)^{\wp(P)})) \\ \text{Undefined} & \end{cases}$
6. $\llbracket F_{\alpha w_n}(r, \dots, r') \rrbracket_s^i = \langle \llbracket r \rrbracket_s^i(\llbracket \alpha \omega_n \rrbracket_s^i), \dots, \llbracket r' \rrbracket_s^i(\llbracket \alpha \omega_n \rrbracket_s^i) \rangle \in I(F, \llbracket \alpha \omega_n \rrbracket_s^i)$
7. $\llbracket \alpha \omega_n \rrbracket_{sg}^c = \{ \llbracket \alpha \omega_n \rrbracket_s^i \mid i \in c \}$

Operations on Points Here is a type of operation underlying dynamic existential quantification.

1. $i \bowtie_\alpha o =$ the i' such that:
 - (a) $0(\sigma_\alpha^{i'}) = o$,
 - (b) for some n , σ_α^i has n members and $\sigma_\alpha^{i'}$ has $n + 1$ members,
 - (c) for all k s.t. $0 \leq k < n$, $k(\sigma_\alpha^i) = k + 1(\sigma_\alpha^{i'})$, and
 - (d) For all $\beta \neq \alpha$, $\sigma_\beta^i = \sigma_\beta^{i'}$

Compositional Dynamics Given a background model and static interpretation functions, there is the following dynamic interpretation function $\llbracket \cdot \rrbracket$, which assigns WFFs functions from contexts to contexts. I assume that the metalanguage has weak kleene logic for dealing with undefinedness; hence, in the update for atomics in 1, if any $\llbracket F_{\alpha w_n}(r, \dots, r') \rrbracket_s^i$ is undefined, then the entire update is undefined.

1. $c \llbracket F_{\alpha w_n}(r, \dots, r') \rrbracket = \{ i \in c \mid \llbracket F_{\alpha w_n}(r, \dots, r') \rrbracket_s^i \}$
2. $c \llbracket \alpha x_\beta \rrbracket = \{ i \bowtie_\alpha d \mid i \in c, d \in D \}$

3. $c[\alpha\omega] = \{i \bowtie_{\alpha} w | i \in c, w \in W\}$
4. $c[\alpha\Omega] = \{i \bowtie_{\alpha} P | i \in c, P \in \wp(W)\}$
5. $c[rlr'] = \{i \in c | [r]_s^i l [r']_s^c\}$
6. $c[rlg] = \{i \in c | [r]_s^i l [g]_{sg}^c\}$
7. $c[glg'] = \{i \in c | [g]_{sg}^c l [g']_{sg}^c\}$
8. $c[\phi; \psi] = c[\phi][\psi]$

Truth-Conditions Given a background model, we can take a proposition $P \in \wp(W)$ as the context set. The resulting initial context is as follows.

1. $c_P = \{\langle w, P \rangle, \emptyset | w \in P\}$

Now, we can extract truth-conditions from a series of updates $U = u_1; \dots; u_n$, which represents the logical form of a sentence, as follows.

1. $TC(U) = \{w | c_{\{w\}}[u_1; \dots; u_n] \neq \emptyset\}$

Conclusion

Imagine old friends catching up on a park bench, or a traveler in a foreign city getting directions from a stranger. At least in its paradigm case—that of ordinary conversation—linguistic communication is a joint activity: it is something interlocutors do together. Now imagine friends waving and smiling goodbye, or soldiers making silent hand signals to one another. These practices make sense in light of their guiding purposes, and linguistic communication is similar to them in this regard. So linguistic communication is like many other forms of human communication in that it involves genuine action and is accordingly sensitive to general principles of rationality.

It should not be controversial that linguistic communication is a rational activity, and perhaps not even that it is a joint one. But what is special about the Gricean approach is that it holds that conversation's rational nature constitutes the most basic or minimal achievement of successful linguistic communication: linguistic understanding of an utterance between speaker and addressee. To illustrate, consider a simple version of Griceanism, according to which the addressee must come to recognize that the speaker is intending to say something. On this simple view, the most basic understanding of an utterance involves the speaker doing enough to make their action intelligible, and the addressee thereby making sense of what the speaker is doing by attributing to them the appropriate intention.

The Gricean approach to linguistic communication, also known as INTENTIONALISM, contrasts with CONVENTIONALISM, according to which in linguistic communication a speaker utters a sentence that itself has a meaning—in virtue of the conventional meanings of the constituent words (and how meanings compose)—where that meaning constitutes a minimal contribution to conversation. An addressee is thus able to figure out what an uttered sentence contributes without making sense of the speaker's action by attributing a Gricean intention. An addressee may instead rely upon knowledge of the conventional meanings of the constituent words (and rules of meaning composition).

Between intentionalism and conventionalism has been a “Homeric struggle,” as noted by

Strawson (1971a: 172), which continues to this day (Lepore & Stone 2015, Harris 2016). In this conclusion, I wish to outline how this dissertation contributes to this debate, on behalf of intentionalism. I do so by connecting the formal semantics of the final chapter with the Gricean approach from the earlier chapters, in order to show how an important general challenge to Griceanism can be thereby met. The challenge is that of explaining what is distinctively linguistic about communication with language.

In the first three chapters, the way that I developed the Gricean approach and applied it to the case of demonstratives strayed in various ways from the simple Gricean view given a few paragraphs above. But all my proposals from those chapters are open to the present challenge. For simplicity and clarity, let me explain how the challenge arises for the simple view just given. Nothing in what was said there about how understanding is achieved appealed to the use of language. The “speaker” might act with the intention to say something, but without using language at all; the “speaker” might, instead, be using impromptu non-verbal gestures in order to say something to an addressee. So, in the case where language is used in communication, what role is it playing? What features do words have such that they play a role in linguistic communication?

Note that the challenge raised by these questions is naturally met by conventionalism. Again, according to conventionalism, minimal understanding does not involve the addressee recognizing that the speaker has a Gricean intention. Instead, understanding consists of the speaker and addressee exploiting a conventional system that pairs words with meanings (as well as provides rules for how those meanings compose). The distinctiveness of linguistic communication is here explained by the distinctiveness of the social institution that enshrines the conventional associations of words with meanings.

But there is one popular response to the challenge, on behalf of Griceanism, called “constraint semantics,” which holds that the properties of a sentence, even as tokened on a particular occasion, merely place a constraint on what a speaker may say with it.²⁴ The constraint itself does not typically constitute a minimal contribution to conversation. So what I wish to do now is show how the formal semantics from the final chapter provides a version of constraint semantics. As already mentioned, doing so will serve to connect that final chapter with the rest of the dissertation, and thus consolidate the overall view of linguistic communication that emerges from this dissertation. In what follows I also provide an accessible outline of the type of formal system developed in the final chapter.

²⁴Early versions of constraint semantics are in Strawson (1950), Grice ([1969] 1989), and Bach (1987). For more recent developments, see Schiffer (2003), Neale (2005), Pietroski (2006, 2018), Heck (2014), and Harris (2020).

In order to state my overall view, the distinction must be drawn between a linguistic expression—including a token linguistic expression, as uttered on a particular occasion—and an utterance of a linguistic expression. An utterance is an action, not a token linguistic expression. My proposal is thus that, when a speaker utters a sentence, full linguistic understanding consists of understanding the speaker's action, where that full understanding requires, but is not exhausted by, understanding of the token sentence, which itself is generated by the operation of linguistic competence. As emphasized by Chomsky (1957, 1965) and followers, the properly linguistic component of linguistic communication is distinctively arational. Linguistic competence, in generating understanding of token linguistic expressions, is encapsulated from any normativity involved in the overarching rational activity of making sense of one another in communication. Hence, linguistic competence is a modular or autonomous component of linguistic communication, which makes it distinctively linguistic.

Given that linguistic competence must generate understanding of a token expression, it contains semantics as a part. So now I wish to elaborate what it means to say that semantics is a part of linguistic competence, and how dynamic semantics provides a conception of semantics that accounts for how that is so. I begin by saying more about linguistic competence.

Syntactic rules—to pick the paradigm aspects of linguistic competence—simply do not make sense from the perspective of general reason, even as enshrined in more or less arbitrary convention. Consider, for instance, the requirement in some languages that every clause has a grammatical subject, as illustrated by the following contrast.

- (1) a. *Is raining.
b. It's raining.

This requirement, known as the Extended Projection Principle, explains the defectiveness of (1-a). The dummy, expletive use of "it" must be added, as in (1-b). Yet there is no good *reason* that (1-a) cannot easily be taken to mean that rain is happening. In addition, here are examples famously given by Chomsky (1957), which motivate the autonomy of syntax. First, here is an example where a sentence should have a perfectly intelligible meaning, given principles of general reasoning, but where syntactic facts make it defective.

- (2) *The child seems sleeping.

It should be easy to take (2) to mean that the child seems *to be* sleeping, but its ungrammaticality makes that difficult, and is the source of the ineliminable defectiveness. On the flipside,

the fact that (3) is syntactically well-formed forces us to give it a meaning, even though the meanings of its constituents should not fit together in a reasonable way.

(3) Colourless green ideas sleep furiously.

These cases show that syntactic rules operate autonomously of general reasoning.²⁵

Let me dwell for a moment on (3)—probably the most well-known example from Chomsky’s work—since its upshot is frequently misunderstood.²⁶ Part of the misunderstanding is due to Chomsky himself. In Chomsky 1957, he argued that (3) is syntactically well-formed, yet lacks meaning. But, as clarified in Chomsky 1965, the point of the example is in fact that there is a notion of linguistic meaning that does not line up with our intuitive notion of meaning—where the latter, as an intuitive notion, is connected with what is reasonable to say or think. So the point of the example is that the sentence is forced to have a meaning, by syntactic rules, where that meaning is a part of autonomous linguistic competence, in addition to syntax.

Thus linguistic competence is not an undifferentiated ability, consisting only of syntax. It is instead a repertoire with several components, making up a grammar. In addition to the syntactic, there are the phonological and semantic components. The semantic component is my focus in what follows.

Here are two further groups of evidence for taking semantics to be a part of autonomous linguistic competence. First, it has been received wisdom that certain syntactic environments place constraints on when quantifiers can covertly take scope over others. These syntactic environments are tensed and relative clauses.

(4) Guinevere has a bone that is in every corner of the house. (Rodman 1976)

(5) John said that everyone had left. (Chomsky 1975)

In (4), the universal quantifier phrase “every corner of the house” cannot take scope over the indefinite “a bone,” so the reading that is forced is that Guinevere has one single bone that is (somehow at once) in every corner of the house. But an alternative reading, which is not available—but which makes more sense—would be that Guinevere has, in each corner of her house, a (potentially different) bone. Similarly, the quantifier “everyone” in the embedded clause in (5) cannot covertly take scope over the attitude verb: the sentence cannot be taken to mean that John said of each person that they had left. Recently, however, Barker (2021) has

²⁵For more examples along these lines, see Collins (2004).

²⁶Thanks to Paul Pietroski for discussion on this point.

amassed evidence that these environments are not scope islands *as a matter of their syntax*. Consider the following.

- (6) The picture of himself that everyone sent in annoyed the teacher. (Hulsey & Sauerland 2006)
- (7) A student made sure that every invited speaker had a ride. (Farkas & Giannakidou 1996)

The sentence (6) can be taken to mean that everyone is such that the picture of themselves they sent in annoyed the teacher. So here is a case where a universal quantifier covertly scopes out of a relative clause. And we see in (7) a universal quantifier scoping out of the embedded tensed clause, and over the indefinite “a student”: there is a reading available where every invited speaker had a potentially different student ensure their ride.

Barker proposes a semantic theory of scope islands, but its details are not of present concern. What matters is that semantics is here operating autonomously in that judgments are forced that do not make sense in light of general rationality. Again, the reading of (4) that is forced barely makes sense; whereas, the inaccessible one does.

The second source of evidence for the autonomy of semantics comes from anaphora resolution. Here is a minimal pair that played a crucial role in motivating Heim’s (1982) seminal work in dynamic semantics.

- (8) I dropped ten marbles and found nine of them. #It is blue.
- (9) I dropped ten marbles and found all but one. It is blue.

In order for the pronoun “it” to be licensed, there must be an entity introduced by an explicit indefinite. In (9), that indefinite is “one.” In (8), the second sentence should easily be taken to be about the missing marble whose existence can be inferred from the first sentence. However, since there is no explicit indefinite, “it” is defective. But, since the first sentences of both (8) and (9) imply that there is a single missing marble, from the perspective of general rationality it should be perfectly easy to give a meaning to “it” in both cases.

If—as I am now suggesting—anaphora resolution is an aspect of linguistic competence, then the following picture emerges. Anaphoric expressions pick up their meaning from discourse referents, which is to say, informally, that they help bundle existentially quantified information at the discourse level. So, in treating the dynamic economy of discourse referents as a part of autonomous linguistic competence, the picture is that the principles of discourse attention are based on semantic universals. And this picture is empirically sup-

ported in the work of Bittner (2009, 2011), who argues that there are semantic universals of discourse attention—the economy of discourse referents regulating anaphoricity—that manifest in various ways throughout the languages of the world. Note, in addition, that pronouns are not the only anaphoric expressions. There are also temporal and modal anaphora (Partee 1984, Stone & Hardt 1999). So, the dynamics of the relative prominence of discourse referents involve hard-coded rules that play a core role of coordinating the meanings of sentences making up a discourse.

The standard way of formally refining talk of discourse referents and the semantic rules governing them is with dynamic semantics. In the final chapter, a dynamic semantics was developed, with a focus on capturing the anaphoricity of demonstratives. Here I give a relatively accessible gloss on the dynamic treatment of demonstratives, in order to illustrate how it generates a conception of semantic content of sentences that is in line with my overall proposal that understanding a sentence, even as tokened, is only part of what is involved in full linguistic understanding of a speaker’s utterance of that sentence: the part generated by autonomous linguistic competence.

To begin, consider the following anaphoric use of a demonstrative.

- (10) A woman entered from stage left. Another woman entered from stage right. This woman was carrying a bouquet of flowers. (Roberts 2002)

The demonstrative “this woman” in the final sentence of (10) helps elaborate the information about the woman introduced by the indefinite “another woman” in the second sentence. The information established by the second two sentences of this discourse is thus roughly that there was some (other) woman who entered from stage right and was carrying a bouquet of flowers.

Here is how the tradition of dynamic semantics provides a straightforward way of capturing the anaphoricity of demonstratives. For present purposes, suppose that demonstratives are variables (for simplicity, I do not model their deictic features nor the noun-phrase complement of complex demonstratives). Suppose also that indefinites are simple existentially quantified sentences, but which semantically update a context in a way that may bind variables outside of their syntactic scope. So here a context is not understood along Kaplanian lines, on which it determines a particular world and assignment function.²⁷ (Worlds are max-

²⁷On each of Kaplan’s (1989) proposed “indexical” and “corrected Fregean” theories of demonstratives, a context determines a value for a demonstrative. The difference is in whether the demonstrative itself gets assigned a value directly by context (on this indexical view, think of the demonstrative simply as a variable), or whether

imally determinate ways things could be; more formally, they are points relative to which predicates are assigned sets of individuals, the ones which satisfy the predicate in that world. Assignment functions determinately value context-sensitive expressions; more formally, they are functions from variables to individuals.) In contrast to the Kaplanian notion, a context in dynamic semantics contains indeterminate information about the way the world is and indeterminately values context-sensitive expressions. More formally, a dynamic context is a set of pairs of worlds and assignment functions. Accordingly, an indefinite “ $\exists xF(x)$ ” introduces a discourse referent by adding, for each world in the input context, a valuation of “ x ” for every object that satisfies “ F ” in that world. An ensuing demonstrative “ x ” helps elaborate the information about that discourse referent by helping to rule out, alongside a predicate, some of those assignment function and world pairs: the ones in which the value assigned to “ x ” by the assignment function does not satisfy the predicate in the world.

So, again, informal talk of discourse referents simply captures how there is semantic machinery that coordinates information at the discourse level with the aid of variables. And the present point is simply that the dynamic approach accounts for how the demonstrative in (10) helps semantically elaborate the information about the (other) woman. But this point raises the question of how to make sense of truth-conditions of sentences, or even entire discourses, within a dynamic framework, since the point makes reference to the information built up by a sentence or discourse. So let me say a few words on truth-conditions in a dynamic setting.

On the present approach, a discourse is treated as a conjunction of simple sentences, and the dynamic semantics of conjunction is successive update, or, equivalently, function composition. Hence, successive update is how complex discourses are assigned update functions: the update function of a discourse is compositionally determined from the conjuncts that make it up. Now, puzzlement over how to associate a piece of information—equivalently: a set of worlds, or truth-condition—with a sentence or discourse disappears with the realization that update functions are more complex (type-theoretic) entities than (characteristic functions of) sets of worlds. Accordingly, a set of worlds may be retrieved from an update update, which is a function from sets of world and assignment function pairs to other such sets. Remaining relatively informal, this extraction of a piece of worldly information from a dynamic meaning goes as follows (Heim 1982: 186–187). Consider all the singleton sets containing a single ordered pair of a world and assignment function, and keep only such singletons that the update is de-

the demonstrative picks up a value from a demonstration in logical form, where this demonstration is assigned a value by context (on this corrected Fregean view, think of the demonstration as a variable and a demonstrative as having a constant character and the identity function as its content). See sections XV and XVI of Kaplan 1989.

fined on and which do not result in the empty set. Finally, determine the set of worlds that appear in one or another of the remaining singletons.

Now here is a simple logical form for the discourse in (10).

$$(11) \quad \exists x Woman(x) \wedge FromLeft(x) \wedge \exists y Woman(y) \wedge FromRight(y) \wedge BouquetCarrying(y)$$

Applying the method of truth-condition extraction just limned, the truth-conditional content of (11) is the set of worlds in which there exists some woman who entered from stage left, as well as some woman who entered from stage right and was carrying a bouquet of flowers. Thus, we get some more clarity on the informal talk of discourse referents: coordinating information at the discourse level with the aid of variables involves a discourse building up an existentially quantified content that such-and-such properties are co-instantiated in the world.

In light of this discussion of anaphoric demonstratives such as in (10), consider a deictic demonstrative.

$$(12) \quad \text{That [pointing to a car] is red.}$$

A flatfooted way of extending the dynamic approach to deixis is to treat the demonstration as an indefinite to which the deictic demonstrative is in fact anaphoric—where here “anaphoric” is meant in the theoretical sense that the expression picks up a discourse referent from its linguistic context. Hence, the proposed simple logical form of (12) is as follows.

$$(13) \quad \exists x Demonstrated(x) \wedge Red(x)$$

Semantically, (13) builds up the following information: there exists something demonstrated that is red.

Thus, in general, from a dynamic semantics perspective, context-sensitive referential expressions in all of their uses are ultimately anaphoric expressions. So, an evolving discourse semantically builds up a content that is existentially quantified by means of discourse referents. I just outlined how this perspective may be applied to deictic demonstratives, and, accordingly, I embrace EXISTENTIALISM: the semantic content of a simple sentence containing a demonstrative is a non-context-sensitive proposition that such-and-such properties are co-instantiated in the world. Thus, setting aside semantically encoded discourse effects, what is required for understanding a sentence (even as tokened on a particular occasion) is determined by the sentence type’s standing semantic content, which itself is a technical notion in

cognitive science (in particular, formal semantics) and corresponds to an aspect of our modular linguistic competence. Hence, I do not mean “content” here as Kaplan’s (1989) does: what is said by a sentence token in Kaplanian context. Rather, the content of a sentence and all of its tokens is simply the truth-condition had at the type level, which is determined by the sentence type’s dynamic meaning given the extraction method limned above.

In §4.3.2 of the final chapter, I compared my existentialism to some alternatives. (Recall, for instance, that the problem with introducing Kaplanian contexts to determine referents for token deictic demonstratives is that it would require that deictic and anaphoric demonstratives are ambiguous, since each would be semantically sensitive to different types of context.) Here I summarize the view that has emerged by answering the following puzzle about (12): How can existentialism be reconciled with the intuition that, in order to understand the utterance of (12), the addressee must think of the particular car being pointed at, and not simply that certain properties are co-instantiated in the world?

The answer is that understanding the utterance of (12) requires more than grasping the uttered sentence’s semantic content. Understanding the *utterance* of (12) is what requires the addressee to think of the particular car, for that must be done in, for instance, recognizing the speaker’s referential intentions. Thus, I hold that utterance understanding involves going beyond semantic content and into pragmatics, broadly construed. And it is open to the Gricean to hold that, in addition to the broadly pragmatic considerations surrounding utterance understanding, linguistic communication involves the operation of linguistic competence, itself encapsulated from any normativity involved in the overarching rational activity. This modular or autonomous component is what makes linguistic communication distinctively linguistic. Thus, the puzzle raised by (12) concerning the shortcomings of dynamic semantics in accounting for reference is resolved in the same way as the issue for Griceanism, raised at the beginning of this conclusion, concerning the distinctiveness of linguistic communication: with the distinction between sentence and utterance understanding, and the claim that the former is not sufficient for the latter.

The crucial claim here may be approximated, in more familiar terminology, with the claim that literal utterance or speech act content—what is said by the speaker’s utterance—is not determined by semantic content. But it must be stressed that grasping what is said by a speaker’s utterance is not essential to utterance understanding, since, as covered in chapter 3, there can be understanding without there being anything that is said. Recall that, in cases of felicitous underspecification, there is understanding despite the interlocutors not settling on any object as the referent of an uttered deictic demonstrative. And note how my existentialist account

of semantic content allows for the possibility of felicitous underspecification. Successfully initiating a discourse referent does not require that there is any particular object to which it corresponds. There is therefore no semantic requirement of specificity, even in deixis—in contrast to the mainstream Neo-Russellian and Neo-Fregean accounts, discussed in chapter 2, which rely on Kaplanian contexts to hold that token demonstratives are associated with particular objects or ways of thinking of objects, respectively, and which thus struggle to account for felicitous underspecification.

In conclusion, the Gricean considerations regarding making sense of one another in communication, which culminated in chapter 3, may have seemed to be in tension with the hard-coded semantic rules governing demonstratives laid down in the dynamic semantics of the final chapter 4. But I hope to have explained how the two approaches are complementary. On the one hand, the Gricean approach must say something about what is distinctive about linguistic communication, and dynamic semantics provides a conceptual and formal elaboration of constraint semantics that does so. On the other hand, the hard-coded rules of anaphora resolution merely constrain reference; thus, for deictic demonstratives the pragmatic considerations of Griceanism fill in how it is that reference is important—at least sometimes—for full understanding.

Bibliography

- Acton, Eric K. & Christopher Potts. 2014. That straight talk: Sarah palin and the sociolinguistics of demonstratives. *Journal of Sociolinguistics* 18(1). 3–31.
- Adamson, Lauren B. & Connie L. Russell. 1999. Emotion regulation and the emergence of joint attention. In Philippe Rochat (ed.), *Early social cognition*, Psychology Press.
- Ahn, Dorothy. 2019. *That thesis: A competition mechanism for anaphoric expressions*: Harvard University dissertation.
- Almotahari, Mahrad & Ephraim Glick. 2010. Context, content, and epistemic transparency. *Mind* 119(476). 1067–1086.
- Anderbois, Scott, Adrian Brosoveanu & Robert Henderson. 2015. At-issue proposals and appositives in discourse. *Journal of Semantics* 32(1).
- Aristotle. 1984. Politics. In Jonathan Barnes (ed.), *Complete works of aristotle: The revised oxford translation*, vol. 2, Princeton University Press.
- Bach, Kent. 1987. *Thought and reference*. Oxford University Press.
- Bach, Kent. 1992. Intentions and demonstrations. *Analysis* 52(3). 140–146.
- Bach, Kent & Robert M. Harnish. 1979. *Linguistic communication and speech acts*. MIT Press.
- Barker, Chris. 2004. Possessive weak definites. In Ji-Yung Kim, Yury A. Lander & Barbara H. Partee (eds.), *Possessives and beyond*, 89–113. GLSA Publications.
- Barker, Chris. 2021. Rethinking scope islands. *Linguistic Inquiry* Manuscript.
- Battich, Lucas & Bart Geurts. forthcoming. Joint attention and perceptual experience. *Synthese*.
- Berridge, Kent C. 2004. Motivation concepts in behavioral neuroscience. *Physiology & Behaviour* 81. 179–209.
- Bezuidenhout, Anne. 1997. The communication of de re thoughts. *Noûs* 31(2). 197–225.
- Bittner, Maria. 2001. Surface composition as bridging. *Journal of Semantics* 18. 127–177.
- Bittner, Maria. 2009. Tense, mood, and centering. Manuscript.
- Bittner, Maria. 2011. Time and modality without tenses or modals. In Renate Musan & Monika

- Rathert (eds.), *Tenses across languages*, de Gruyter.
- Blackburn, William K. 1988. Wettstein on definite descriptions. *Philosophical Studies* 53.
- Blumberg, Kyle. 2020. Demonstratives, definites and non-redundancy. *Philosophical Studies* 177. 39–64.
- Bowker, Mark. 2019. Saying a bundle: Meaning, intention, and underdetermination. *Synthese* 196. 4229–4252.
- Bratman, Michael. 1993. Shared intention. *Ethics* 104(1). 97–113.
- Bratman, Michael. 2014. *Shared agency: A planning theory of acting together*. Oxford University Press.
- Bruner, Jerome. 1975a. From communication to language—a psychological perspective. *Cognition* 3(3). 255–287.
- Bruner, Jerome. 1975b. The ontogenesis of speech acts. *J. Child Lang* 2. 1–19.
- Bruner, Jerome. 1977. Early social interaction and language acquisition. In H. R. Schaffer (ed.), *Studies in mother-infant interaction*, Academic Press.
- Bruner, Jerome. 1983. *Child's talk: Learning to use language*. Oxford University Press.
- Buchanan, Ray. 2010. A puzzle about meaning and communication. *Noûs* 44(2). 340–371.
- Buchanan, Ray. 2014. Reference, understanding, and communication. *Australasian Journal of Philosophy* 92(1). 55–70.
- Caldano, Michela & Kenny R. Coventry. 2019. Spatial demonstratives and perceptual space: To reach or not to reach? *Cognition* 191. 1–6.
- Camp, Elisabeth. 2018. Insinuation, common ground, and the conversational record. In Daniel Fogal, Daniel W. Harris & Matt Moss (eds.), *New work on speech acts*, Oxford University Press.
- Campbell, John. 1987. Is sense transparent? *Proceedings of the Aristotelian Society* 88. 273–292.
- Campbell, John. 1997. Sense, reference and selective attention. *Proceedings of the Aristotelian Society Supplementary* 71.
- Campbell, John. 2002. *Reference and consciousness*. Oxford University Press.
- Campbell, John. 2005. Joint attention and common knowledge. In Naomi Eilan, Christoph Hoerl, Teresa McCormack & Johannes Roessler (eds.), *Joint attention: Communication and other minds: Issues in philosophy and psychology*, Oxford University Press.
- Campbell, John. 2011. An object-dependent perspective on joint attention. In Axel Seemann (ed.), *Joint attention: New developments in psychology, philosophy of mind, and neuroscience*, The MIT Press.
- Campbell, John. 2017. Comments on Imogen Dickie's *Fixing Reference*. Presented at the 91st

- Meeting of the APA Pacific Division.
- Campbell, John. 2018. Joint attention. In Kirk Ludwig (ed.), *The routledge handbook on collective intentionality*, Routledge.
- Carpenter, Malinda & Kristin Liebal. 2011. Joint attention, communication, and knowing together. In Axel Seemann (ed.), *Joint attention: New developments in psychology, philosophy of mind, and neuroscience*, The MIT Press.
- Carpenter, Malinda, Katherine Nagell & Michael Tomasello. 1998. Social cognition, joint attention, and communicative competence from 9 to 15 months of age. *Monographs of the Society for Research in Child Development* 63(4).
- Charlow, Nate. forthcoming. Metasemantic quandaries. In Bill Dunaway & David Plunkett (eds.), *Meaning, decision, and norms: Themes from the work of allan gibbard*, Maize Books.
- Chevallier, Coralie, Gregor Kohls, Vanessa Troiani, Edward S. Brodtkin & Robert T. Shultz. 2012. The social motivation theory of autism. *Trends in Cognitive Sciences* 16(4).
- Chomsky, Noam. 1957. *Syntactic structures*. Mouton Publishers.
- Chomsky, Noam. 1965. *Aspects of the theory of syntax*. MIT Press.
- Chomsky, Noam. 1975. *Questions on form and interpretation*. Peter de Ridder Press.
- Chomsky, Noam. 1995. Language and nature. *Mind* 104. 1–61.
- Clark, Herbert H. 1996. *Using language*. Cambridge University Press.
- Collins, John. 2004. Faculty disputes. *Mind & Language* 19(5). 503–533.
- Coppock, Elizabeth & David Beaver. 2015. Novelty and familiarity for free. In *Proceedings of the 20th amsterdam colloquium*, .
- Coventry, Kenny R., Debra Griffiths & Colin J. Hamilton. 2014. Spatial demonstratives and perceptual space: Describing and remembering object location. *Cognitive Psychology* 69. 46–70.
- Coventry, Kenny R., Berenice Valdés, Alejandro Castillo & Pedro Guijarro-Fuentes. 2008. Language within your reach: Near-far perceptual space and spatial demonstratives. *Cognition* 108. 889–895.
- Dennett, Daniel Clement. 1969. *Content and consciousness*. Routledge.
- Dickie, Imogen. 2011. Visual attention fixes demonstrative reference by eliminating referential luck. In Christopher Mole, Declan Smithies & Wayne Wu (eds.), *Attention: Philosophical and psychological essays*, Oxford University Press.
- Dickie, Imogen. 2014. A practical solution to the problem of empty singular thought. In Manuel García-Carpintero & Martí Genoveva (eds.), *Empty representations*, Oxford University Press.

- Dickie, Imogen. 2015. *Fixing reference*. Oxford University Press.
- Dickie, Imogen. 2020. Understanding singular terms. *Aristotelian Society Supplementary Volume* 94(1). 19–55.
- Dickie, Imogen & Gurpreet Rattan. 2010. Sense, communication, and rational engagement. *Dialectica* 64(2). 131–151.
- Diessel, Holger. 1999. *Demonstratives: Form, function, and grammaticalization*, vol. 42 Typological Studies in Language. John Benjamins Publishing Company.
- Diessel, Holger. 2006. Demonstratives, joint attention, and the emergence of grammar. *Cognitive Linguistics* 17(4). 463–489.
- Diessel, Holger & Kenny R. Coventry. 2020. Demonstratives in spatial language and social interaction: An interdisciplinary review. *Frontiers in Psychology* 11. 1–14.
- Doherty, Martin & Josef Perner. 1998. Metalinguistic awareness and theory of mind: Just two words for the same thing? *Cognitive Development* 13. 279–305.
- Dretske, Fred I. 1981. *Knowledge & the flow of information*. MIT Press.
- Dummett, Michael. 1973. *Frege: Philosophy of language*. Gerald Duckworth.
- Ebert, Christian, Cornelia Ebert & Robin Hörnig. 2020. Demonstratives as dimension shifters. Manuscript.
- Egan, Frances. 2014. How to think about mental content. *Philosophical Studies* 170(1). 115–135.
- Egan, Frances. forthcoming. A deflationary account of mental representation. In Joulia Smortchkova, Krzysztof Dolega & Tobias Schlicht (eds.), *Mental representations*, Oxford University Press.
- Eilan, Naomi. 2005. Joint attention, communication, and mind. In Naomi Eilan, Christoph Hoerl, Teresa McCormack & Johannes Roessler (eds.), *Joint attention: Communication and other minds: Issues in philosophy and psychology*, Clarendon Press.
- Eilan, Naomi. 2015. Joint attention and the second person. Draft.
- Eilan, Naomi. 2018. Knowing and understanding other minds: On the role of communication. Manuscript.
- Eilan, Naomi. 2020. Other i's, communication, and the second person. *Inquiry*.
- Elbourne, Paul. 2005. *Situations and individuals*. MIT Press.
- Elbourne, Paul. 2008. Demonstratives as individual concepts. *Linguistics and Philosophy* 31. 409–466.
- Elbourne, Paul. 2013. *Definite descriptions*. Oxford University Press.
- Evans, Gareth. 1977. Pronouns, quantifiers, and relative clauses (i). *Canadian Journal of Philosophy* 467–536.

- Evans, Gareth. 1982. *The varieties of reference*. Clarendon Press.
- Evans, Gareth. 1985. Understanding demonstratives. In *Collected papers*, Clarendon Press.
- Evans, Gareth & John McDowell. 1976. Introduction. In Gareth Evans & John McDowell (eds.), *Truth & meaning: Essays in semantics*, Oxford University Press.
- Fagin, Ronald, Joseph Y. Halpern, Yoram Moses & Moshe Vardi. 1995. *Reasoning about knowledge*. The MIT Press.
- Fara, Delia Graff. 2015. Names are predicates. *Philosophical Review* 124(1).
- Farkas, Donka & Anastasia Giannakidou. 1996. How clause-bounded is the scope of universals? *Semantics and Linguistic Theory* 6. 35–52.
- Fine, Kit. 2007. *Semantic relationism*. Blackwell Publishing Ltd.
- Fine, Kit. 2012. The structure of joint intention. Unpublished Manuscript, available on academia.edu.
- von Fintel, Kai. 2008. What is presupposition accommodation, again? *Philosophical Perspectives* 22(1). 137–170.
- von Fintel, Kai & Anthony S. Gillies. 2011. ‘Might’ made right. In Andy Egan & Brian Weatherston (eds.), *Epistemic modality*, 108–130. Oxford University Press.
- Flavell, John H. 1977. The development of knowledge about visual perception. In C. B. Keasey (ed.), *The nebraska symposium on motivation: Social cognitive development*, vol. 25, 43–76.
- Flavell, John H. 1992. Perspectives on perspective taking. In H. Beilin & P. B. Pufall (eds.), *Piaget’s theory: Prospects and possibilities*, vol. 14 The Jean Piaget Symposium Series, 107–139.
- Flavell, John H., Barbara A. Everett, Karene Croft & Eleanor R. Flavell. 1981. Young children’s knowledge about visual perception: Further evidence for the level 1–level 2 distinction. *Developmental Psychology* 17(1). 99–103.
- Frankfurt, Harry G. 1978. The problem of action. *American Philosophical Quarterly* 15(2). 157–162.
- Frege, Gottlob. 1960a. On sense and reference. In Peter Geach & Max Black (eds.), *Translations from the philosophical writings of gottlob frege*, Basil Blackwell.
- Frege, Gottlob. 1960b. The thought. In Peter Geach & Max Black (eds.), *Translations from the philosophical writings of gottlob frege*, Basil Blackwell.
- Gilbert, Margaret. 1990. Walking together: A paradigmatic social phenomenon. *Midwest Studies in Philosophy* 15.
- Gilbert, Margaret. 2013. *Joint commitment: How we make the social world*. Oxford University Press.

- Glanzberg, Michael. 2009. Not all contextual parameters are alike. Manuscript.
- Glanzberg, Michael & Jeffrey C. King. 2020. Binding, compositionality, and semantic values. *Philosophers' Imprint* 20.
- Grice, Paul. [1957] 1989. Meaning. In *Studies in the way of words*, Harvard University Press.
- Grice, Paul. [1969] 1989. Utterer's meaning and intentions. In *Studies in the way of words*, Harvard University Press.
- Grice, Paul. [1975] 1989. Logic and conversation. In *Studies in the way of words*, Harvard University Press.
- Groenendijk, Jeroen & Martin Stokhof. 1991. Dynamic predicate logic. *Linguistics and Philosophy* 14(1). 39–100.
- Harman, Gilbert. 1990. The intrinsic quality of experience. *Philosophical Perspectives* 4. 31–52.
- Harris, Daniel. 2019. We talk to people not contexts. *Philosophical Studies* .
- Harris, Daniel W. 2016. Intentionalism versus the new conventionalism. *Croatian Journal of Philosophy* 16(47). 173–201.
- Harris, Daniel W. 2020. Semantics without semantic content. *Mind & Language* 1–25.
- Hawthorne, John & Ofra Magidor. 2009. Assertion, context, and epistemic accessibility. *Mind* 118(470). 377–397.
- Hawthorne, John & David Manley. 2012. *The reference book*. Oxford University Press.
- Heck, Richard G., Jr. 1995. The sense of communication. *Mind* 104(413). 79–106.
- Heck, Richard G., Jr. 2000. Nonconceptual content and the "space of reasons". *The Philosophical Review* 109(4). 483–523.
- Heck, Richard G., Jr. 2002. Do demonstratives have senses? *Philosophers' Imprint* 2(2). 1–33.
- Heck, Richard G., Jr. 2007. Are there different kinds of content. In Brian P. McLaughlin & Jonathan D. Cohen (eds.), *Contemporary debates in philosophy of mind*, Blackwell.
- Heck, Richard G., Jr. 2014. Semantics and context-dependence: Towards a strawsonian account. In Alexis Burgess & Brett Sherman (eds.), *Metasemantics: New essays on the foundations of meaning*, Oxford University Press.
- Heim, Irene. 1982. *The semantics of definite and indefinite noun phrases*: University of Massachusetts dissertation.
- Heim, Irene. 1983. On the projection problem for presuppositions. In *Formal semantics - the essential readings*, Blackwell.
- Heim, Irene. 1985. Direct reference explained away. Manuscript.
- Heim, Irene. 1990. E-type pronouns and donkey anaphora. *Linguistics and Philosophy* 13. 137–177.

- Heim, Irene. 1991. Artikel und definitheit. In A. von Stechow & D Wunderlich (eds.), *Semantik: Ein internationales handbuch der zeitgenössischen forschung*, 487–534. de Gruyter.
- Heim, Irene & Angelika Kratzer. 1998. *Semantics in generative grammar*. Blackwell Publishers.
- Hobson, R. Peter. 2002. *The cradle of thought*. Pan Macmillan.
- Hobson, R. Peter. 2005. What puts the jointness in joint attention? In Naomi Eilan, Christoph Hoerl, Teresa McCormack & Johannes Roessler (eds.), *Joint attention: Communication and other minds: Issues in philosophy and psychology*, Oxford University Press.
- Hulsey, Sarah & Uli Sauerland. 2006. Sorting out relative clauses. *Natural Language Semantics* 14. 111–137.
- Jankovic, Marija. 2014. Communication and shared knowledge. *Philosophical Studies* 169. 489–508.
- Kadmon, Nirit. 1990. Uniqueness. *Linguistics and Philosophy* 13(3). 273–324.
- Kaplan, David. 1989. Demonstratives: An Essay on the Semantics, Logic, Metaphysics and Epistemology of Demonstratives and other Indexicals. In J. Almog, J. Perry & H. Wettstein (eds.), *Themes from kaplan*, Oxford University Press.
- King, Jeffrey C. 2001. *Complex demonstratives: A quantificational account*. The MIT Press.
- King, Jeffrey C. 2014a. The metasemantics of context sensitivity. In Alexis Burgess & Brett Sherman (eds.), *Metasemantics: New essays on the foundations of meaning*, 97–118. Oxford University Press.
- King, Jeffrey C. 2014b. Speaker intentions in context. *Noûs* 48(2). 219–237.
- King, Jeffrey C. 2017. Strong contextual felicity and felicitous underspecification. *Philosophy and Phenomenological Research* 1–27.
- Kirk-Giannini, Cameron Domenico. 2018. Uniformity motivated. *Linguistics and Philosophy* 41. 665–684.
- Köpping, Jan. 2020. How to trivialize uniqueness. In *Proceedings of sinn und bedeutung* 24, 429–446.
- Lakoff, Robin. 1974. Remarks on ‘this’ and ‘that’. In *Proceedings of the chicago linguistics society*, 345–356.
- Lederman, Harvey. 2017. Two paradoxes of common knowledge: Coordinated attack and electronic mail. *Noûs* 1–25.
- Lederman, Harvey. 2018. Uncommon knowledge. *Mind* 127(508). 1069–1105.
- León, Felipe. 2021. Joint attention without recursive mindreading: On the role of second-person engagement. *Philosophical Psychology* 34(4). 550–580.
- Lepore, Ernie & Matthew Stone. 2015. *Imagination and convention*. Oxford University Press.

- Lewis, David. 1969. *Convention: A philosophical study*. Harvard University Press.
- Lewis, David. 1979. Scorekeeping in a language game. *Journal of Philosophical Logic* 8(1). 339–359.
- Lewis, David. 1980. Index, context, and content. In Stig Kanger & Sven Öhman (eds.), *Philosophy and grammar*, 79–100. Reidel.
- Loar, Brian. 1976. The semantics of singular terms. *Philosophical Studies* 30(6). 353–377.
- Ludlow, Peter & Gabriel Segal. 2004. On a unitary semantical analysis for definite and indefinite descriptions. In Marga Reimer & Anne Bezuidenhout (eds.), *Descriptions and beyond*, Oxford University Press.
- MacFarlane, John. 2016. Vagueness as indecision. *Proceedings of the Aristotelian Society* 90. 255–283.
- Maclaran, Rose. 1982. *The semantics and pragmatics of english demonstratives*: Cornell University dissertation.
- Mandelkern, Matthew & Daniel Rothschild. 2020. Definiteness projection. *Natural Language Semantics* 28. 77–109.
- Marr, David. 1982. *Vision: A computational investigation into the human representation and processing of visual information*. W. H. Freeman and Company.
- Martin, Scott. 2016. Supplemental update. *Semantics and Pragmatics* 9. 1–61.
- McDowell, John. 1994. *Mind and world*. Harvard University Press.
- McDowell, John. 1998a. De re senses. In *Meaning, knowledge, and reality*, Harvard University Press.
- McDowell, John. 1998b. Meaning, communication, and knowledge. In *Meaning, knowledge, and reality*, Harvard University Press.
- McDowell, John. 1998c. On the sense and reference of a proper name. In *Meaning, knowledge, and reality*, Harvard University Press.
- Mesman, Judi, Marinus H. van Ijzendoorn & Marian J. Bakermans-Kranenburg. 2009. The many faces of the still-face paradigm: A review and meta-analysis. *Developmental Review* 29(2). 120–162.
- Milward, Sophie J. & Malinda Carpenter. 2018. Joint action and joint attention: Drawing parallels between the literatures. *Soc Personal Psychol Compass* 12. 1–11.
- Moll, Henrike & Andrew N. Meltzoff. 2011. Joint attention as the fundamental basis of understanding perspectives. In Axel Seemann (ed.), *Joint attention: New developments in psychology, philosophy of mind, and neuroscience*, The MIT Press.
- Moll, Henrike, Andrew N. Meltzoff, Katharina Merzsch & Michael Tomasello. 2013. Taking versus confronting visual perspectives in preschool children. *Developmental Psychology*

- 49(4). 646–654.
- Moore, George Edward. 1962. *Commonplace book: 1919–1953*. Routledge.
- Moore, Richard. 2017. Gricean communication and cognitive development. *The Philosophical Quarterly* 67(267).
- Murray, Sarah E. 2014. Varieties of update. *Semantics and Pragmatics* 7(2). 1–53.
- Neale, Stephen. 1992. Paul grice and the philosophy of language. *Linguistics and Philosophy* 15. 509–559.
- Neale, Stephen. 2005. Pragmatism and binding. In Zoltan Szabó (ed.), *Semantics versus pragmatics*, Oxford University Press.
- Ninan, Dilip. 2010. Semantics and the objects of assertion. *Linguistics and Philosophy* 33. 355–380.
- Nowak, Ethan. 2018. Saying 'that F' is saying which F: Complex demonstratives, hidden arguments, and presuppositions. Unpublished Manuscript.
- Nunberg, Geoffrey. 1993. Indexicality and deixis. *Linguistics and Philosophy* 16(1–43).
- Nyström, Pär, Emilia Thorup, Sven Bölte & Terje Falck-Ytter. 2019. Joint attention in infancy and the emergence of autism. *Biological Psychiatry* 86. 631–638.
- O'Madagain, Cathal & Michael Tomasello. 2019. Joint attention to mental content and the social origin of reasoning. *Synthese* .
- Partee, Barbara. 1984. Nominal and temporal anaphora. *Linguistics and Philosophy* 7(3). 243–286.
- Paul, Matthias. 1999. *Success in referential communication*, vol. 80 Philosophical Studies. Kluwer Academic Publishers.
- Peacocke, Christopher. 1992. *A study of concepts*. MIT Press.
- Peacocke, Christopher. 2005. Joint attention: Its nature, reflexivity, and relation to common knowledge. In Naomi Eilan, Christoph Hoerl, Teresa McCormack & Johannes Roessler (eds.), *Joint attention: Communication and other minds: Issues in philosophy and psychology*, Oxford University Press.
- Peet, Andrew. 2017. Referential intentions and communicative luck. *Australasian Journal of Philosophy* 95(2). 379–384.
- Peeters, David & Aslı Özyürek. 2016. *This and That* revisited: A social and multimodal approach to spatial demonstratives. *Frontiers in Psychology* 7. 1–4.
- Percus, Orin. 2006. Antipresuppositions. In Ayumi Uyema (ed.), *Theoretical and empirical studies of reference and anaphora: Toward the establishment of generative grammar as an empirical science*, Japan Society for the Promotion of Science.
- Pickel, Bryan & Brian Rabern. forthcoming. The myth of occurrence-based semantics. *Lin-*

- guistics and Philosophy* .
- Pietroski, Paul. 2006. Character before content. In Judith Jarvis Thomson & Alex Byrne (eds.), *Content and modality: Themes from the philosophy of robert stalnaker*, Oxford University Press.
- Pietroski, Paul. 2018. *Conjoining meanings: Semantics without truth-values*. Oxford University Press.
- Potts, Christopher. 2005. *The logic of conventional implicatures*. Oxford University Press.
- Rabern, Brian. 2012. Against the identification of assertoric content with compositional value. *Synthese* .
- Rabern, Brian. 2013. Monsters in kaplan's logic of demonstratives. *Philosophical Studies* 164. 393–404.
- Rakoczy, Hannes, Delia Bergfeld, Ina Schwarz & Ella Fizke. 2015. Explicit theory of mind is even more unified than previously assumed: Belief ascription and understanding aspectuality emerge together in development. *Child Development* 86(2). 486–502.
- Récanati, François. 1993. *Direct reference: From language to thought*. Blackwell Publishers.
- Richard, Mark. 2004. Contextualism and relativism. *Philosophical Studies* 119. 215–242.
- Roberts, Craige. 2002. Demonstratives as definites. In Kees van Deemter & Roger Kibble (eds.), *Information sharing: Reference and presupposition in language generation and interpretation*, 89–196. CSLI Press.
- Roberts, Craige. 2003. Uniqueness in definite noun phrases. *Linguistics and Philosophy* 26. 287–350.
- Roberts, Craige. 2005a. Context in dynamic interpretation. In Laurence R. Horn & Gregory Ward (eds.), *The handbook of pragmatics*, Blackwell.
- Roberts, Craige. 2005b. Pronouns as definites. In Marga Reimer & Anne Bezuidenhout (eds.), *Descriptions and beyond*, Oxford University Press.
- Roberts, Craige. 2012a. Information Structure in Discourse: Towards an Integrated Formal Theory of Pragmatics. *Semantics and Pragmatics* 5(7).
- Roberts, Craige. 2012b. Information Structure: Afterword. *Semantics and Pragmatics* 5(7). 1–19.
- Roberts, Craige. 2020. Character study: a *de se* semantics for indexicals. Manuscript.
- Rochat, Philippe & Tricia Striano. 1999. Social-cognitive development in the first year. In Philippe Rochat (ed.), *Early social cognition*, Psychology Press.
- Rodman, Robert. 1976. Scope phenomena, movement transformations, and relative clauses. In Barbara Partee (ed.), *Montague grammar*, Academic Press.
- Roessler, Johannes. 2005. Joint attention and the problem of other minds. In Naomi Eilan,

- Christoph Hoerl, Teresa McCormack & Johannes Roessler (eds.), *Joint attention: Communication and other minds: Issues in philosophy and psychology*, Clarendon Press.
- Roessler, Johannes. 2020. Plural practical knowledge. *Inquiry* .
- Rothschild, Daniel & Seth Yalcin. 2017. On the dynamics of conversation. *Noûs* 51(2). 24–48.
- Rubio-Fernandez, Paula. forthcoming. Pragmatics markers: The missing link between language and theory of mind. *Synthese* .
- Russell, Bertrand. 1905. On denoting. *Mind* 14(56). 479–493.
- Rüther, Johanna & Ulf Liszkowski. 2020. Ontogenetic emergence of cognitive reference comprehension. *Cognitive Science* 44. 1–18.
- Sauerland, Uli. 2008. Implicated presuppositions. In A. Steube (ed.), *Sentence and context: Language, context, and cognition*, de Gruyter.
- Scaife, Michael & Jerome Bruner. 1975. The capacity for joint visual attention in the infant. *Nature* 253. 265–266.
- Schiffer, Stephen. 1972. *Meaning*. Oxford University Press.
- Schiffer, Stephen. 2003. *The things we mean*. Oxford University Press.
- Schiller, Henry Ian. 2019. Acquaintance and first-person attitude reports. *Analysis* 79(2). 251–259.
- Schlenker, Philippe. 2005. Minimize restrictors! (notes on definite descriptions, condition c and epithets). In *Proceedings of sinn und bedeutung* 9, 385–416.
- Schlenker, Philippe. 2011. Donkey anaphora: The view of sign language (asl and lsf). *Linguistics and Philosophy* 34. 341–395.
- Seemann, Axel. 2010. The other person in joint attention: A relational approach. *Journal of Consciousness Studies* 17(5–6). 161–182.
- Seemann, Axel. 2019. *The shared world: Perceptual common knowledge, demonstrative communication, and social space*. The MIT Press.
- Simons, Mandy, Judith Tonhauser, David Beaver & Craige Roberts. 2010. What projects and why. *Proceedings of SALT* 20. 309–327.
- Siposova, Barbora & Malinda Carpenter. 2019. A new look at joint attention and common knowledge. *Cognition* 189. 260–274.
- Smithies, Declan. 2011. What is the role of consciousness in demonstrative thought? *Journal of Philosophy* 108(1). 5–34.
- Smithies, Declan. 2019. *The epistemic role of consciousness*. Oxford University Press.
- Sperber, Dan & Deirdre Wilson. 1986. *Relevance: Communication and cognition*. Blackwell Publishers.

- Stalnaker, Robert. [1978] 1999. Assertion. In *Context and content*, 78–95. Oxford University Press.
- Stalnaker, Robert. 2002. Common ground. *Linguistics and Philosophy* 25. 701–721.
- Stalnaker, Robert. 2009. On hawthorne and magidor on assertion, context, and epistemic accessibility. *Mind* 118(470). 399–470.
- Stalnaker, Robert. 2014. *Context*. Oxford University Press.
- Stalnaker, Robert. 2018. Dynamic pragmatics, static semantics. In Daniel Fogal, Daniel W. Harris & Matt Moss (eds.), *New work on speech acts*, Oxford University Press.
- Stanley, Jason & Zoltán Gendler Szabó. 2000. On quantifier domain restriction. *Mind & Language* 15(2). 219–261.
- Starr, William B. 2010. *Conditionals, meaning and mood*: Rutgers University dissertation.
- Stern, Daniel N. 1985. *The interpersonal world of the infant: A view from psychoanalysis and developmental psychology*. Harvard University Press.
- Stojnić, Una, Matthew Stone & Ernie Lepore. 2013. Deixis (even without pointing). *Philosophical Perspectives* 27. 502–525.
- Stojnić, Una, Matthew Stone & Ernie Lepore. 2017. Discourse and logical form: Pronouns, attention and coherence. *Linguistics and Philosophy* 40. 519–547.
- Stokhof, Martin, Jeroen Groenendijk & Frank Veltman. 1996. Coreference and modality. In Shalom Lappin (ed.), *Handbook of contemporary semantic theory*, 179–216. Blackwell.
- Stokke, Andreas. 2014. Truth and context change. *Journal of Philosophical Logic* 43. 33–51.
- Stone, Matthew. 1999. Reference to possible worlds. Tech. rep. Rutgers University Center for Cognitive Science.
- Stone, Matthew & Daniel Hardt. 1999. Dynamic discourse referents for tense and modals. In Harry Bunt & Reinhard Muskens (eds.), *Computing meaning*, vol. 2, 302–321. Dordrecht.
- Strawson, Peter F. 1950. On referring. *Mind* 59(235). 320–344.
- Strawson, Peter F. 1964. Intention and convention in speech acts. *Philosophical Review* 73(4). 439–460.
- Strawson, Peter F. 1971a. Meaning and truth. In *Logico-linguistic papers*, Methuen & Co Ltd.
- Strawson, Peter F. 1971b. Propositions, concepts, and logical truths. In *Logico-linguistic papers*, Methuen & Co Ltd.
- Strawson, Peter F. 1988. Perception and its objects. In Johnathan Dancy (ed.), *Perceptual knowledge*, Oxford University Press.
- Striano, Tricia & Evelin Bertin. 2005. Coordinated affect with mothers and strangers: A longitudinal analysis of joint engagement between 5 and 9 months of age. *Cognition and*

- Emotion* 19(5). 781–790.
- Syrett, Kristen & Todor Koev. 2015. Experimental evidence for the truth conditional contribution and shifting information status of appositives. *Journal of Semantics* 32(3). 525–577.
- Szabó, Zoltan. 2020. The goal of conversation. *Aristotelian Society Supplementary Volume* 94(1). 57–86.
- Szabó, Zoltán Gendler. 2000. Descriptions and uniqueness. *Philosophical Studies* 101. 29–57.
- Tenenbaum, Sergio. 2015. Representing collective agency. *Philosophical Studies* 172. 3379–3386.
- Toates, Frederick. 1986. *Motivational systems*. Cambridge University Press.
- Tomasello, Michael. 1999. *The cultural origins of human cognition*. Harvard University Press.
- Tomasello, Michael. 2009. *Why we cooperate*. The MIT Press.
- Tomasello, Michael. 2014. The ultra-social animal. *European Journal of Social Psychology* 44. 187–194.
- Tomasello, Michael. 2018. How children come to understand false beliefs: A shared intentionality account. In *Proceedings of the national academy of sciences of the united states of america*, .
- Tomasello, Michael. 2019. *Becoming human: A theory of ontogeny*. Harvard University Press.
- Tomasello, Michael, Malinda Carpenter, Josep Call, Tanya Behne & Henrike Moll. 2005. Understanding and sharing intentions: The origins of cultural cognition. *Behavioral and Brain Sciences* 28. 675–735.
- Trevarthen, Colwyn. 1979. Communication and cooperation in early infancy: a description of primary intersubjectivity. In M. Bullowa (ed.), *Before speech: The beginning of human communication*, 321–347. Cambridge University Press.
- Tronick, Edward, Heidelise Als, Lauren Adamson, Susan Wise & T. Berry Brazelton. 1978. The infant's response to entrapment between contradictory messages in face-to-face interaction. *Journal of the American Academy of Child Psychology* 17(1). 1–13.
- Williamson, Timothy. 2000. *Knowledge and its limits*. Oxford University Press.
- Wittgenstein, Ludwig. [1953] 2009. *Philosophical investigations*. Wiley-Blackwell.
- Wolter, Lynsey Kay. 2006. *That's that: The semantics and pragmatics of demonstrative noun phrases*: University of California, Santa Cruz dissertation.
- Yalcin, Seth. 2014. Semantics and metasemantics in the context of generative grammar. In Alexis Burgess & Brett Sherman (eds.), *Metasemantics: New essays on the foundations of meaning*, Oxford University Press.