

***Global Expressivism:
Language Agency without Semantics, Reality
without Metaphysics***

Stephen J. Barker

University of Nottingham

October 2007

Preface

0. Language Agency

§0 Foreshadowing: Interpretation, Meaning-Theory and the Home Language

§1 Functional and Interpretative Orientations to the Language Agent

§2 The Interpretative Stance and Representation

§3 Metaphysics, the Metaphysics of Meaning, and Expressivism

§4 Rival Programs: Davidson, Dummett, Brandom, Fodor, Quine, and Others

§5 What to Expect

1. Expressivism and Semantic and Cognitive Categories

§6 Preliminaries

§7 The Essence of Expressivism

§8 Semantic Minimalist Expressivism

§9 Semantic and Cognitive Categories: A problem for Everyone

§10 Inferentialism

§11 Expressivism Returns as a Theory of Language Agency

2. How to be a Global Expressivist

§12 Basic Structures of Language Agency

§13 Signalling Systems and Communicative Exchanges

§14 Signalling Systems for Humans: Assertion and Proto-Assertion

§15 Non-Assertions, Belief, Judgement, and Embedding

§16 Truth-Aptness and Basing

§17 Content-Attribution, Truth-Attribution, and S-Processing

§18 Wissenschaft und Verstehen, and The Home Language

§19 Expressivism Resurrected

3. Reason, Truth, and Logical Complexity

§20 Truth

§21 Grounding and Expressivism about Reason

§22 *Kinds of Reasons, Kinds of C-laws*

§23 *Expressivism goes Negative*

§24 *Disjunction*

§25 *Validity and Rightness*

4. *Expressivism about Everything: The Non-Evaluative*

§26 *II-Properties, Concepts, and Interpretative Stances*

§27 *Formal Predicates*

§28 *Perceptual Predicates: Colour, a Case Study*

§29 *Context-Relativity and Open Texture*

§30 *Predicate Mental Elements: Representationalism and Inferentialism*

§31 *Two Perspectives, and Objectualism about Predication*

§32 *Expressivism about Natural Kinds*

§33 *Expressivism about Cause and Disposition*

§34 *Blueprint for Expressive Analysis*

5. *Expressivism about Everything: Value*

§35 *The Return to Value-Expressivism*

§36 *Value-Predicates*

§37 *Disagreement, Negation, Irrationality, and Rejection*

§38 *Thick Value Terms: Value-Content and Implicature*

§39 *Value Modals*

§40 *Propositional Attitudes, Embedding, and Validity*

§41 *Wishful Thinking, Objectivity, and Value's Reality*

6. *Word and Object*

§42 *Referring Terms and Semantic and Causal Perspectives*

§43 *Inferentialism and Reference*

§44 *Referring, Proto-Referring, and Denoting*

§45 *Objects and Non-Objects*

§46 *Reference to Objects and ι -Properties*

§47 Names

§48 Reference by Descriptive Uniqueness

§49 Events and Property Instances

7. *Word and Non-Object*

§50 Non-Objects and Nominalization

§51 Predication: Objects and Non-Objects

§52 Modes of Reference for Non-Object Terms

§53 Deconstructing Quine's Quantification: Existentials

§54 Deconstructing Quine's Quantification: Universals

§55 Propositions as Non-Entities

§56 Objects, Non-Objects, and Natures

8. *Truth-Making and Facts as Non-Objects*

§57 Truth-Making and its Discontents

§58 Reification and Facts

§59 Facts Freed: Not all Relations relate Entities

§60 Fact-Talk and Nominalization

§61 Making and Analytical Recipes

§62 Truth-Making and Logically Complex Truths

§63 Formal and Non-Formal Truth-Making

§64 Facts and Being without Metaphysics

9. *Objective and Non-Objective Truth*

§65 Objective Truth

§66 Faultless Disagreement

§67 Realism, Non-Cognitivism, Content Relativity, and Relativity of Truth

§68 Two Modes of Assertion and Belief: Objective and Non-Objective

§69 Faultless Disagreement and the Relativity of Assertion and Belief

§70 Putnam's Conceptual Relativity and Epistemic Modals

§71 Objectivity and Representation

10. *Meaning without Nature*

§72 *Meaning-Theory and Language Agency Revisited*

§73 *Talk about Content*

§74 *Normativity and Meaning*

§75 *Meaning without Ontology: Meanings as Non-Objects*

§76 *Folk Semantic and Pragmatic Principles*

§77 *Interpretation*

§78 *Dissolving Meaning Scepticism*

11. *Realism without Metaphysics*

§79 *Meaning-Theory and Metaphysics: Partners in Crime*

§80 *Metaphysics*

§81 *Realism*

§82 *From the Structure of Language to the Structure of Reality?*

§83 *Self-Application and Last Words*

Preface

There is a wide-spread belief amongst theorists of mind and language. This is that in order to understand the relation between language, thought, and reality we need a theory of meaning and content; that is, a normative, formal science of meaning, which is an extension and theoretical deepening of folk ideas about meaning. That formal, normative science may take the form of a theory of truth-conditions, a theory of inferential practice, a theory of assertion, and so on. The idea is that conditions of truth, inferential relations, or assertability conditions fix meaning and illuminate thereby thought's place in a broader reality.

According to this book, the widespread belief is wrong. There is no theory of meaning waiting out there for us that will explain the relation between thought, reality, and language. Rather the form of a theory that illuminates the relation of language, thought, and reality is a theory of *language agency*. In a nutshell, the theory of language agency is a theory of competence, without being a theory of understanding or grasping rules. It is a theory of cognitive structure and language production. This theory distils all there is to say about language, thought, and reality. It does not supplement a theory of truth-conditions or semantic norms. It is not the explanation of how a speaker, qua cognitive system causally embedded in a larger reality, is able to use a language with some pre-existing semantic characterization. There is no pre-existing semantic characterization. The theory of language agency is a solvent through which the theoretical questions about meaning, truth, reference, and the nature of semantic norms disappear. This does not mean that facts about meaning, reference, truth, and norms disappear. There are such facts. They are as solid and real as any other facts. They are not mere fictions or somehow insubstantial.

The dissolution of the theory of meaning is accompanied by another disappearance. That is the disappearance of metaphysical questions in a number of domains. Once we complete the theory of language agency, then just as theoretical questions about meaning disappear, certain theoretical questions about existence disappear. Having provided a theory of the language agency for talk of meaning, fact, property, relation, and proposition, there is no question left over about what meanings, facts, properties, relations, and propositions are.

There is no theory to be given of their natures. This is not because they have primitive, irreducible natures. Rather it is because, in a sense to be clarified in this work, they lack natures. The metaphysical perspective on facts, properties, propositions, and so on, imputes natures to these things, natures demanding investigation. But as we shall see, that perspective is not inevitable, indeed, it is undermined by the theory of language agency to be given.

The disappearance of meaning-theory and the removal of a metaphysical perspective on features of reality have a common cause. It is the denial that representation and related notions have an explanatory role in the theory of language function. This is not to deny that there are representational properties, relations of correspondence between language and a larger world, or that our thoughts are made true or false by reality. All this remains in place as it should. All that changes is the explanatory role of representation.

The removal of representation from the theory of language agency can only occur because the latter provides a specific characterization of the function that lies behind semantic vocabulary that dissolves the need for representation. All talk of truth, content, rightness, semantic norms, and reference is talk in which the language agent takes on a *simulative* orientation to verbal behaviours of others or of itself. Roughly, and deploying an as yet unexplicated vocabulary, to understand or interpret is to simulate. To assign meaning is to *defend* a particular simulation. To assign truth is to defend the state that is simulated. In short, the theory of language agency provides, in functional terms, a characterization of that cognitive style which underpins interpretation and assignment of truth and reference. I call this approach *Global Expressivism*. That is because it generalizes some of the insights brought to the study of value-language by expressivists. However, it removes these insights from the clouding affects of attempting to make expressivism a semantic theory. Expressivism about value fails as a semantic theory of value talk. However, global expressivism can succeed as a theory of all talk because it is not a semantic theory but a theory of language agency, wherein the theory of meaning is replaced by a theory of talk about meaning.

Introduction

Language Agency

§0 Foreshadowings: Interpretation, Meaning-Theory, and the Home Language

Philosophy of language seeks to illuminate the relation between thought and reality. Many people think that the way to provide such illumination is through a theory of meaning, say, a theory of truth-conditions or use-conditions. The main contention of this book is that this is false. Illumination of the relation between thought and reality is not provided by a theory of meaning. (By *theory of meaning* we include theories of reference and representation.) Instead of a theory of meaning, what we need and what I offer below is a theory of *language agency*, a theory that explains the functional nature of the language-using agent and how that agent is causally embedded in a wider world. This theory of agency is not a theory of understanding or grasping meaning. Theories of grasping meaning require that we have a theory of meaning since they are about *grasping meaning*, but that, I urge, is what we cannot have. Nor is it a normative account of use. Normativity is just as problematic as meaning. Rather the relation that the theory of language agency I offer has to meaning, understanding, and normativity is this: the theory of language agency is a purely functional (and non-representational) analysis of the cognitive architecture of an agent and their activities with signs that includes, as a sub-component, a theory of that cognitive structure and activity which underpins talk about meaning, understanding, and norms. In this approach, meaning-theory is replaced by a theory of talk about meaning, norms, and understanding.

I dub the approach to language function *global expressivism*—GE for short—because it generalises some ideas implicit in value expressivism. More about that in a moment. Two general lines of argument are proposed to motivate this approach which are developed in this book. The first appeals to general features about language, meaning, and interpretation,

pointing to the need to transcend meaning-theory. The second is that theories of meaning inevitably bring semantic and metaphysical puzzles with them, which are removed by renouncing meaning-theory. The semantic puzzles centre upon language's relation to reality, model-theoretic paradox, objectivity, meaning scepticism, context relativity. The metaphysical puzzles are puzzles about the nature of truth, facts, properties, relations, truth-making, truth-bearers and propositions. By escaping meaning-theory, GE escapes these puzzles. The solutions it offers are read off the structure of language agency. They are not solutions within semantics. *Semantics* is the problem, not the solution. The rest of this introduction sketches these two arguments and the form of GE.

Semantics as a Formal, Normative Science

The dream of philosophy of language is a theory of meaning. Its paradigm form is called *Semantics*. Semantics comes in many varieties of which truth-conditional semantics is the paradigm. Truth-conditional semantics can be reductive—the truth-conditions of sentences are conceived of as objects like sets of possible worlds or states of affairs—or non-reductive—theories aim to specify compositionally truth-conditions for sentences through T-sentences—biconditionals of the form '*S* is true iff *P*'.¹ Semantic theories can also take the form of so called use-theories, which assert that semantic facts are facts of use grounded in verification or assertability conditions, inferential practice, or dispositional regularities.²

Semantics is, paradigmatically, a kind of formal normative science. It is formal because it removes itself from the empirical workings of minds and practices of language users to describe an abstract structure—for example, sets of worlds and functions thereupon, the workings of a truth-theory, or the activities of ideal conversational participants. It is

¹ Theorists of the first kind include, Lewis (1970), Barwise and Perry (1983), Stalnaker (1983), Montague (1975), Cresswell (1985), Katz (1990). Theorists of the second kind include Davidson (1967, 1984), McDowell (1977, 1978), Lycan (1984), Appiah (1986), LePore and Kappellen (2005), Sainsbury (2005), Taylor (2005).

² Dummett (1976) and Tennant (1997) offer verificationist semantics. Brandom's (1994, 2005) offers a normative pragmatics that equates sentence-meaning with inferential role cashed out in terms of the normative impact that assertions have on conversational score. Horwich (2005) takes a theory of use to be a theory of dispositions, which include inferential dispositions that are naturalistically described. Theorists like Fodor (1987), hold out for reduction of representation, as do Field (1972), Devitt and Sterelny (1987) and Devitt (1996). Theorists like Grice (1957, 1969) propose Lockean theories, which are use-theories in which words and sentences are correlated with concepts or thoughts, that are psychological types that can be tokened in speakers. Other neo-Lockeans include: Bennett (1975), Alston (2000), and Davis (2005). Gauker (2002) offers a theory treating assertability relative to a context generated by conversational purpose as a primitive.

normative because it carries directives about what speakers should ideally do. It is science because it sees itself as empirically constrained and because it uses formal methods akin to mathematics and logic. The term *science* in this context may be closer to *hermeneutic science* in the case of some theorists and more like *cognitive science* in the case of others.³

We can see Semantics as a refinement of a pre-theoretic body of thought—call it *folk semantics*. Just as physics is a development of folk physics, Semantics is meant to be a development of folk semantics. Think then of semantic theories as axiomatizations of an extended folk semantics, with added theoretical posits facilitating this formalisation or allowing the ontology of the folk theory to be theoretically reduced.

The orthodox view is that there is a formal, normative scientific theory of meaning, a refinement of folk semantics, out there waiting for us, which will be central to the illumination of language, thought and reality. This book argues that this is false. If we are looking for illumination about the nature of language in relation to thought and non-mental reality, looking for theories of meaning—doing Semantics—is the wrong way to go. There is no theory of meaning. Philosophers have offered views falling under the *no-theory-of-meaning* label before. The following could be said to fall under that heading: Quine’s view that there is no theory of synonymy or analyticity, and no meanings qua objects; Davidson’s denial of reductive theories of reference; Wittgenstein’s idea that meaning theory should be replaced by talk of practices and forms of life. GE (Global Expressivism), the theory of language agency developed in this book, is both distinct from and more radical than these positions. In GE, there is nothing at all theoretically illuminating to be told about meaning. There is no theory of word-meaning, sentence-meaning, sameness of meaning, possessing meaning; no theory of understanding, no theory of truth, no theory of reference, no theory of meanings. There is no theory of concepts. Yet, at the same time, GE accepts the perfect legitimacy of talk of meaning, reference, truth, synonymy, language, propositional content, concepts, and word-meanings. This talk is not a mere *façon de parler* or a form of fictional discourse. It is as legitimate as any other talk. It is just that there is no theory, as required by

³ Theorists like Davidson (1984) are hermeneutic, taking interpretative stances to be irreducible features of meaning. Theorists like Horwich (2005) are not. Horwich also denies that meaning is intrinsically normative, so falls outside the paradigm in one respect.

Semantics, about the subject matter of this talk. This is not to say that meanings, truth, propositions, and so forth are primitive features of reality. On the contrary, our talk of meaning, truth, propositions and so forth is perfectly intelligible. No residue of mystery remains. It is rather that meaning, truth, propositions, and normative statuses are entirely without nature. It is language agency that has a nature, whose character we can theoretically investigate, not meanings and norms.

The Principle Argument

This book develops two sustained arguments against Semantics and for the no-theory-of-meaning view embodied in GE. The first is a broadly theoretical argument that goes to the essence of meaning in relation to explanation. Here is the argument in outline:

1. Meaning is irreducibly interpretative. To assign meaning is to engage in interpretation. There are no assignments of meanings or concepts outside of an interpretative stance towards an agent.

2. Meaning attributions presuppose an interpretative framework, and consequently so does meaning-theory. Because meaning-theory presupposes an interpretative framework, it cannot supply a theory of the activity that is interpretation and the interpreter's home language.

3. A theory of the home language can be neither a theory of meaning of that language nor a theory of grasping meaning. It will be a theory of something else: a functional characterisation of the language-using and interpreting agent, which includes, as a sub-component, a functional account of talk of meaning and understanding.

4. A theory of thought and its relation to reality is essentially a theory about the home language, which is to say the cognitive and computational achievement that allows there to be any interpretative activity at all.

5. Therefore, a theory illuminating thought, language and reality cannot be a theory of meaning. It will be a theory of language agency for the home language.

The steps in this argument are evidently open to dispute, and the import of the moves made not entirely clear. I now attempt to improve matters, though the full effect of the argument will only become apparent as we proceed to the main exposition.

Premise 1: There are philosophers who would deny the first step, that meaning is irreducibly interpretative—Fodor (1978) and Horwich (1998) are two. They would argue that we can define necessary and sufficient conditions for meaning something in purely functional, causal terms. But two general lines of argument suggest that they are wrong. First, the work of Kripke (1981) indicates that speakers cognitive limitation, the fact that they can make mistakes, bars any functional characterization of the distinction between correct and incorrect use. Second, Davidson (1984) argues that attribution of meaning is holistic. Meaning possession can depend indefinitely on a range of other cognitive features, in ways that cannot be grounded in a theory of non-intentional cognitive facts. I endorse both these lines of argument and reinforce them as we proceed.

Step 2: If premise 1 is correct, meaning attributions presuppose an interpretative stance towards a speaker. If meaning-theory is about meaning and words possessing meanings, then theorizing about meaning occurs within the scope of interpretation. What is interpretation? One answer is that interpretation is the activity of using the meaningful terms of a home language to display the meanings possessed by the interpretee's terms. What is it for this home language to possess meaning? If the interpretative thesis of meaning is right, there is no answer to this question that does not invoke an interpretative perspective. It is useless to say that possessing meaning resides in being interpretable. The explanatory circle here is too tight to provide satisfaction. Meaning-theory, I suggest, cannot provide an account of the nature of interpretation and of the home language when it participates in interpretation, since meaning-theory presupposes an interpretative framework. The nature of home language, seen not as an object for interpretation but as a system that enables there to be any interpretation at all, escapes meaning-theory.

Step 3: Won't a theory of understanding provide the analysis we need of the home language? The analysis of the home language is simply a question of the competence of the speaker. If so, one might think that a complete theory of a language-using agent U comprises: (i) a theory of meaning, a formal normative science, applied to the language used by U; (ii) a theory of competence, a theory of understanding, detailing what U needs to grasp the meanings described in the theory of meaning. But this conception of the task is open to

dispute. It implies a reductive account of what it is to follow a rule, possess a concept, or meaning something by words. Meanings only live within an interpretative perspective, but a theory of grasping suggests a functionalist, cognitive theory outside that perspective. If our claims about the irreducibility of meaning are right, we must repudiate the task so described. Indeed, Quine and Davidson repudiate it: they think there is nothing of any general theoretical value to be said about the home language.

Repudiating a theory of grasping meanings is one thing. Repudiating the need to understand that achievement that is the home language is another. The capacity to use the home language requires a cognitive system that interacts with reality. The general functional character of that system and its links to words and world, and the activity we call interpretation, needs to be and can be described. A functional description of the language agent provides that description. That's what GE, the theory developed in this book, does. But the perspective that GE gives us on the home language is neither a theory of meaning nor understanding.

Step 4: Our aim is to understand the nature of language, thought and its relation to reality. That is partly about seeing meaning in a way that makes its relation to the natural world unproblematic. It seems to many language-theorists that the way to do that is through a theory of meaning. But what can a theory of meaning do if all meaning talk presupposes an interpretative framework, and a theory of the home language is a functional analysis?

There are those who think that reflecting on the activity of interpretation as it is revealed to us in ordinary thought using folk categories will provide illumination. Such reflection on interpretation does give us some general theoretical insights: that an interpretation of a speaker's language must also interpret that speaker's mental states and behaviour; that interpretation must be guided by norms of maximising rationality, charity, etc. But one has doubts that such reflection in itself will uncover the relation of thought to reality. Meaning theories, one might object, relate the sentences used in the object language to the world. That is true. It does do this. But meaning theory does so using the interpreter's language in this way: it attempts to *say* in the interpreter's language what is said by object-language sentences. One provides truth-conditions in clauses like the following:

‘Every good boy deserves fruit’ is true iff every good boy deserves fruit.

In this exercise, the interpreter uses fully-fledged concepts, EVERY, GOOD, BOY, etc., linking them to terms in the interpreted language. This exercise will not illuminate the nature of the activity that underpins deployment of words, *every*, *boy*, *good*, etc, that enables one to say of those words that they convey certain concepts. Rather it presupposes it. Truth-theoretic approaches reveal nothing about concepts—as Davidson (1967) admits. But there is something to be revealed, and it is startling to think that illumination about the relation of thought to reality could go on without revealing it. Here again we have to look to the functional theory of the home language. The revelation is not so much analysis of concepts as functional analysis of the cognitive sub-structure of our use of predicates.

Again, it may be objected that within a theory of meaning for a language, T-sentences like those above are derived from axioms, which apply to each term, *every*, *boy*, *good*, and so on. This process of derivation from axioms reveals the compositional workings of the language. There lies the illumination. That’s a popular idea but I argue that it is false. First, the compositional story just relates simple expressions to complex expressions. Secondly, the compositional structure revealed through truth-theory is but a phenomenal shadow—and a not always faithful shadow—of a deeper structure revealed at the level of language agency.

In sum, what provides deep illumination about language and its relation to a larger world is a theory of the home language, but that is beyond the purview of meaning-theory.

Conclusion 5: The conclusion is that an illuminating theory of language to reality is fundamentally a theory of the home language. But a theory of the home language is not a theory of meaning or of understanding. Rather it must be something else: a functional analysis of language agency. But how then do meaning, interpretation, and understanding get into the picture? They do so in this way: part of the functional description of language agency will include a functional characterisation of meaning-talk. Providing a functional analysis of meaning talk will uncover the structures that underpin that activity we call *interpretation*.

The lines of argument sketched in steps 1 to 5 will be fleshed out in the course of this book. Before we get to that I now try and give a preliminary idea of what a theory of language agency does, and how it relates to questions of meaning and interpretation.

§1 Functional and Interpretative Orientations to the Language Agent

The principle argument invoked implicitly the idea of two perspectives one can take up towards the language using subject: *the functional* and *the interpretative*. These orientations correspond to a familiar distinction from the 19th century, that between *Wissenschaft* and *Verstehen*. *Wissenschaft* is natural science couched in terms of law, causation, and dispositions. *Verstehen* is human science, couched in terms of interpretation. The divide is mirrored in the contemporary distinction of a purely functionalist stance towards an agent and an *intentional stance*—Dennett (1989). How do we reconcile these two perspectives on language activity and mental life? My principle argument points to how we should do that. Reconciliation comes not by providing a theory of content, doing Semantics, but by providing a functional theory of the language-agent whose activities, including that activity underpinning talk of content, become functionally explicable. It is not the reduction of the subject matter of *Verstehen* to *Wissenschaft*, but the reduction of the activity that lies behind practising *Verstehen* that is studied by *Wissenschaft*. With this functional analysis, no substantial question about meaning remains. How can the substantial questions about meaning just disappear like that? This section and the next sketch an answer to that question.

By *functional analysis* I mean the specification of causal and dispositional structures governing some domain of phenomena. This description specification may be at a level of abstraction; we describe a causal-dispositional pattern that may be realized in different ways by underlying patterns of properties. GE (global expressivism) provides a functional specification of the language agent. The latter is a functional system basically comprising:

(i) A set of dispositions to *signalling activity*. This activity is the production and reception of signs. Signalling has a causal structure, which I describe in chapter 2. In a signalling exchange between two agents, one produces signs s as the result of an inner state Ω

and a receiver perceives s and mirrors or simulates the production of s through structural *homologues* Ψ in her system. I call this *S-processing*. Language users spontaneously S-process others' utterances and gestures in terms of *homologues*, states that are potential causes of their own signalling productions. The functional structure of understanding and interpretation is *simulation*. Meaning-statements are unconscious *defences* of dispositions to S-process symbols. In a sense to be made precise, they are expressions of simulative mental states.

(ii) An underlying system of mental states which form a network, called the *NC* or *natural cognitive system*, in an individual and their relation to a broader reality. This idea of a network is a familiar one. Philosophers talk of networks of beliefs, networks of sentences linked by inter-animation or inference, and networks of cognitive states.⁴ What GE does is something new with the network idea. The NC system is a network of *pre-doxastic* states linked by cognitive-laws, which are, in some cases, the precursors of activities that we call *reason* and *inference*. Concepts of content play no role in characterizing this system.

A language agent, a signalling system, is a rich and variegated thing. The productions of the signalling system and its NC system have many forms. Most of GE's time is spent describing these. Content is not functionally reducible, but a lot of activities related to it are. GE provides a functional specification of semantic, pragmatic, and cognitive categories. It provides functional specifications of speech-act-types, and of belief and desire as kinds of states a system can be in. GE provides functional specification of logical syntax. That means a specification of the structure of the kinds of mental elements that underpin production of different categories of symbols, and what constitutes such categories. It provides a functional specification of logical complex sentences, and a functional specification of the mental elements that underpin use of logical constants, negation, conjunction, disjunction, and so on. It provides a functional theory of predicates and of simple sentences. That includes a theory of divisions between kinds of predicates: evaluative predicates like *is good*, *is tasty*, as distinguished from non-evaluative, formal predicates, like *exists*, and empirical predicates,

⁴ See Quine and Ullian (1970) for the web of belief; Quine (1960), Wittgenstein (1975) for networks of sentences, Brandom (1994) networks of sentences inter-animated by inference, and Clark (1993) for networks of cognitive states.

such as colour terms, *red*, natural-kind sortal terms, like *dog*, or *water*, or causal-dispositional terms, like *causes* or *impenetrable*.

GE assigns cognitive structures in a compositional way to words and sentences. So GE explains the productive, creative aspects of language. But these compositional structures are not based in semantic properties, such as truth or reference. Languages are not truth-conditionally or referentially compositional. No normative concept is deployed in characterising such functions, except that of *normal function* as used in natural science. It certainly makes no appeal to *semantic rules*.

Mention of semantic rules brings us to the question of semantic and normative vocabulary. What is the functional specification of such talk, that is, talk of truth, meaning, entailment, belief, reference, correctness, rightness, rules, and what semantically ought to be done? Answering this question takes us through a reflexive twist in relation to the two perspectives on the language agent that we have introduced: *the functional stance* and *the interpretative stance*. Because talk of meaning is intimately connected with interpretation, in characterising the functional underpinning of semantic and normative vocabulary, we must construct the distinction between the interpretative stance and the functional stance within a functional characterisation of language agency. In other words, the two stances reflect distinct functionally characterisable orientations that we may have to the behaviour and mentality of others.

The mode of *Verstehen*, the interpretative stance, is present when we make interpretative claims. Assertions about meaning, reference, representation, truth, predicate application, entailment, rightness, correctness, and so on, are literally activities in which we are simulating another's language production activity. In a technical sense to be defined below, in such assertions we *defend* dispositions to S-process. Comprehending, understanding, and interpretation are exercises in simulation. Let me emphasise that simulation is not simply an epistemic route through which we arrive at hypotheses about what a speaker means. Rather it is literally what the mode of speech, *meaning-attribution*, is. So in asserting *The word 'E' means n*, one is literally defending deployment of a mental

element underpinning *n*, one's homologue, to simulate productions of 'E'.⁵ Simulation is not conscious. We have no grasp of the structures that do the simulation. It is what we do, not what we grasp.

Contrast now the functional orientation: *Wissenschaft*. It is a view from *no interpretative perspective* in this sense. When one takes up the position of *Wissenschaft* to an agent, no assertions have the simulative nature characteristic of those within the interpretative orientation. One does not describe a speaker's activity through a simulative orientation. In the mode of *Wissenschaft*, we ignore the deliverances of those natural systems in ourselves that are the simulating, or S-processing systems. No explanatory statements by the theorist in the functional stance are statements in which they defend a simulative (S-processing) stance to others. Someone in the orientation of *Wissenschaft*, or the functional stance, may spontaneously understand their subjects' utterances, but they leave these invitations aside. It is my contention that the pure stance of the cognitive scientist in which we examine the intrinsic functional structure of signalling systems is, or should be, this stance.

The interpretative stance, *Verstehen*, is simulative. The functional stance, *Wissenschaft*, is not. This is the functional basis of the claim that meaning is irreducibly interpretative. Meanings are irreducibly interpretative not simply because talk of them involves a different vocabulary—the way that talk of clouds might be irreducible to talk of water-molecules. But rather because that vocabulary involves a different cognitive stance: a simulative stance, quite literally, in one case, and a non-simulative stance in the other.

This book spends nearly all its time at the level of functional specification of language agency. That's why it says nothing about meaning. That's why it never uses the terms *refers*, *means*, *represents*, *true*, *has the content*, *applies to*, or *entails* as part of its explanations of the home language. That's why it does not describe the mental elements of the signalling system in terms of the idioms of content, such as representational states. All its concepts are in the pure realm of functional analysis.

⁵ There is a familiar distinction in philosophy of mind between *theory theory* and *simulation theory*. See Currie and Ravenscroft (2002). That distinction is largely an epistemic one about how thinkers, say children, learn to apply mental attributes. The simulation theory view is that they simulate others' states as a basis for attitude ascription. The present view is pushing this idea further. It is saying that an assertion of *By 'E' O means n*, is itself structured by simulation. The claim is not merely an epistemic one.

The radical thought of GE is this. Once we have described, from the functional perspective, the general capacities of a language agent and in particular its ability to take up an intentional stance—based in S-processing—then we have described all there is to the substantial problem of language, reality and thought, and thus about meaning. There is no room left for a science of meaning, or substantial questions about meaning. We are now ready to see how this can be so.

§2 The Interpretative Stance and Representation

In GE, semantic notions like reference, satisfaction, truth-conditions, or representation have no explanatory role in the analysis of language function. They cannot because their language function is based in simulation, and simulation is excluded in the functional orientation. Thus GE does not explain how a simple sentence *T is F* functions in terms of its conditions of truth, fixed by what *T* picks out, and by what *F* applies to. GE does not explain how a predicate functions by assigning a worldly correlate to it, like a universal, an extension, or a mereological fusion, or function, in terms of which truth can be defined. The fixing of an extension is not required, because we are not defining conditions of truth. Mental elements are assigned to predicates, and the natures of these mental elements are defined by their interactions with the world. It's just not an interaction that is semantic: the embedding in the world is causal. Semantic properties do not appear at the level of functional structure. The function of the truth-predicate is not given in terms of the conditions under which it holds of a sentence or a thought. Rather the function of the truth-predicate is given by a description of the activity and cognitive underpinnings of the language-using agent in terms of an intersubjective, S-processing orientation. (We describe that in chapters 2 and 3.)

Having described the language agent from the functional perspective, we can return to an interpretative perspective, which is to say we activate that mode of engaging with language through simulation—S-processing—that lies dormant in the functional perspective. We can now use our semantic vocabulary, and talk of meanings and reference, but now we realize that the analysis of language activity given from the functional perspective has left us

with nothing to say about the nature of meaning. The question *What is the meaning of a predicate?* has no interesting answer. We might ask: what is the meaning of *red*? In the language agency analysis of *red* we assign no extension, property, or mereological fusion of red things to *red* or the mental elements underpinning its deployment. To interpret someone as meaning red is not to assign an extension to their term. It is rather to defend a simulative stance to their activity with that term. In short, when we look back and view the language agency foundation of our activities with predicates and semantic vocabulary, nothing reveals itself as a ground for theoretical claims about meaning. Talk of meaning is perfectly legitimate talk, but meaning lacks any theoretically describable nature.

This concept of predicate function stands deeply opposed to the orthodox idea that the function of a predicate is explained in terms of what it contributes to fixing truth-conditions of sentences in which it appears. Given a predicate's function is to perform this role, its function requires some worldly thing to be correlated with it: an extension or a property. So we need to know what the correlate is and how the correlation is achieved. We are back with serious semantic questions, where meaning has a nature to be theoretically investigated. Call the view of predication that issues in this result the *objectual view of predication*.

We shall recur to the objectual view of predication and GE's opposition to it frequently in what follows. One of my concerns will be to undermine reasons for thinking it must be right. The main reason for thinking so is summed up in two theses: (a) Realism is correct at least to this extent: reality fixes truth, truth does not fix reality. (b) A necessary condition for maintaining realism is that the function of predication is to determine conditions of truth. GE's position is perfectly realist: it accepts (a). What we reject is (b). The dependency of truth on being does not require a representational analysis of predicate function. Appeal to representation does no work at all in illuminating realism and related concepts, such as objectivity.

We are now in a position to sum up the basic idea in GE. The recognition of the interpretative nature of meaning demands a functional analysis of the home language, which includes a functional analysis of the interpretative stance. That analysis is given in terms of a simulative orientation. But the simulative stance enables us, when coupled with the analysis

of the language agency underpinning the truth-predicate, to dispense with representationalist models of predicate and language function. But the dissolution of representationalism means that talk of meaning cannot issue in interesting theoretical view on what meaning is. Meaning is there, and we happily talk of it, but there is no theory of meaning beyond platitudes of folk semantics. Moreover, it is not these platitudes that do theoretical work, as proposed in some forms of semantic minimalism. Nothing does theoretical work.

In relation to meaning, reference, and semantic facts, GE is neither a form of reductionism, nor intentional primitivism, nor instrumentalism, nor fictionalism. Rather its view is that there is no deep explanatory viewpoint upon which we reveal the nature of language qua system of words and meanings. There is a sense in which—to be made quite precise below—language is not an object, and meaning has no nature. It is certainly false to say that a language is a set, i.e., a set of sentences and a set-theoretical object corresponding to an interpretation. Semantics is an attempt somehow to find an objective standpoint on language qua meaningful system. There is not one. For that reason, Semantics is a doomed project. With the passing of Semantics its paradoxes pass as well. We dissolve of a range of semantic puzzles including: open-texture, faultless disagreement, model-theoretic paradox, and meaning scepticism.⁶

GE does not leave our talk about meaning high and dry. As well as explaining the functional underpinnings of kinds of meaning-attribution sentences, GE explains meaning intuitions, that is, our dispositions to assert a range of sentences related to meaning: the T-schema, truth-making schemas, claims of synonymy, kinds of meaning attribution. GE provides a theory of intuitions about truth, meaning, and validity, but does not give us a theory of truth, meaning, or validity. GE thereby explains our inclinations to affirm a range of theses about meaning that we can call *folk semantics*. Folk semantics is true and far from trivial, in which many familiar theses are affirmed. For example, that the content of what is said in a sentence corresponds to its conditions of truth. GE does not entail these claims.

⁶ GE also provides solutions to the semantic paradoxes and vagueness. These are the topics examined in Barker (200+).

Rather, it provides a theory of their intuitive plausibility, that is, an analysis of their cognitive underpinnings which issues in an explanation of why we are disposed to assert them.

§3 Metaphysics, the Metaphysics of Meaning, and Expressivism

One might not be entirely satisfied with this sketch of GE's anti-reductism since it might appear to have left out an important issue. We talk of meanings, we talk of truth, we talk of propositions and relations of reference. Many theorists hope to supply theories of these things. What drives the imperative for a theory is an apparently minimal and uncontroversial assumption. We quantify over meanings, propositions, and properties of reference. Quine's (1948) principle of ontological commitment tells us that the values of variables of bound quantifiers are entities to whose existence we are committed. With Quine's principle, our quantification over properties of meanings, propositions, and properties of reference leaves us with a swag of ontological commitments to semantic entities whose natures we must then ponder and seek to analyse.

How does GE extricate itself from this difficulty? The answer is that Quine's ontological principle is false. Not all kinds of quantification carry ontological commitment. This is not to embrace some kind of fictionalism or substitutional quantification. It is an initially more illusive claim: Quine's idea is based in implicit acceptance of objectualism about predication, in the sense defined above. Let me clarify.

GE does not explain the function of a referring term as its role in introducing a referent. Reference does not have an explanatory role in relation to language function. That idea just follows from the fact that semantic vocabulary cannot be used with the functional orientation to the language agent. What GE provides is a theory of how the production of referring terms operates and the kinds of mental elements underpinning referring terms. (These mental elements do not function representationally.) Referring terms come in two forms: There are those whose potential referents are things that can appear at the level of functional analysis, and play a role in explaining the nature of the mental elements underpinning these referring terms, and there are those whose potential referents never appear

at the level of functional analysis. In this second case, we see the language agent embedded causally in the world, with mental elements underpinning their referring terms causally interacting and sensitive to reality, but nothing in the world at this causal level enters the picture that is later called, from the interpretative stance, the referent of the term.

The first kinds of terms I call *object terms*, which, when they refer, refer to things like people, material objects, events. The second are *non-object* terms, which, when they refer, refer to properties, relations, facts, and propositions; the kinds of things philosophers often call *abstract objects*. The non-object terms are, in the simplest cases, *nominalizations*, that is, terms derived by grammatical transformation from predicates, sentences, relational terms.

The referents of object terms can be said to have *natures*. We can legitimately ask what they are, and natural science or common sense can potentially answer. In contrast, the referents of non-object terms cannot be said to have natures. It is because these referents are without natures that we cannot really say they are objects. Objects are things that have natures, a place in the world, and knowing them requires some causal contact with them. None of these concerns apply to the referents of non-object terms. That is because, as we shall see, at the level of language function, there is literally nothing that features as referent. That might look like a kind of fictionalism. But it is not at all.

Quine is wrong. Ontological commitment is not the necessary concomitant of quantification. Contemporary ontology does rest upon a mistake.⁷ GE is deeply anti-metaphysical. Metaphysics is that discipline that implicitly assumes that having provided an analysis of how language functions, there are questions left over about the nature of things talked about in the language, questions that cannot be answered by natural science or common sense. But this assumption just reflects our implicit acceptance of a representationalist conception of language function.

Expressivism

What I have done so far is try to show, in general terms, how GE is a framework in which theoretical questions about meaning can disappear, without elimination of meanings,

⁷ Yablo (1998) argued for the view the Quine's criterion of ontological commitment assumes an account of when we speak literally as opposed to metaphorically.

and how metaphysical questions about the nature of a range of things can disappear, without those things being eliminated. But why have I introduced the curious term *Global expressivism* to describe the theory that does this? GE earns its name *expressivism* because it extrapolates certain features of value-expressivism to all areas of discourse. According to value-expressivism, a speaker makes an assertion about value with a sentence such as *Lying is bad* by uttering that sentence and *expressing* a motivational state. The speaker does not describe her state in making the utterance. Her utterance is not true if and only if she has the state. Rather, expressing is *voicing* an affective state. What is interesting about value-expressivism as a metaethical theory is that it is not a theory of value or norms. It is rather a theory of the activity of normative language. At no stage does value-expressivism tell you what norms or values are. The theorist hungry for a theory of value must be turned away, indeed, must be shown to have a misplaced ambition.

As such value-expressivism, seen properly, is a perfect model for the orientation given above. I say *seen properly*. By that I mean that we strip expressivism of the first order claims that value-sentences are not truth-apt, never form real assertions, or manifest real beliefs or really correspond to facts. Once we clear away this confusion, seeing expressivism purely as an explanatory doctrine, then we can see it as doing exactly what I have proposed: it describes from the perspective of Wissenschaft, the functional character of value-language—the agency or signalling activity of language-agents who make value-statements. In short, value-expressivism contains the seed of a correct apprehension about how all language functions, and a correct explanatory orientation. Hence the title of this book.

§4 Rival Programs: Davidson, Dummett, Brandom, Quine, Fodor, and Others

The framework I have described will be fleshed out in the course of this work. But before we start doing that it is worthwhile indicating the location of GE relative to rival positions.

Davidson

Davidson (1984) argues that: (i) the interpretative stance is central to the concept of meaning. All content attributions, attributions of reference and propositional content and

concepts only take place within an interpretative stance. *(ii)* Interpretative theories are Tarskian truth-theories, in which sentence-content is given by truth-theories and truth-theory reveals the compositional structure of the language. *(iii)* Studying interpretation—truth-theories empirically fitted to languages—illuminates the relation between thought and reality.

GE's agreement with Davidson is with *(i)*. It disagrees with his claims *(ii)* and *(iii)*. As noted, Davidson denies that analyzing the home language illuminates thought in relation to reality, but for GE that is where all the action is. For GE, meaning-attributions are not claims about truth-conditions. Truth-theories distort our appreciation of the logical form of sentences and are not the way to reveal compositional structure.

Dummettian Anti-Realism

Dummett's work can be seen as a kind of antithesis to Davidson's work—a deep demand for a theory of the home language—which contends that a theory of meaning must be a theory of understanding. For Dummett, that means a theory of use, and use for Dummett boils down to conditions of verification. Dummett is inspired by the intuitionist logic. Logical connectives are explained by introduction and elimination rules, and no-logical terms by limited sets of inference roles and correlation with canonical grounds in perception or proof. Although sympathetic to Dummett's demand for a theory of the home language, GE eschews Dummettian verificationism. Dummett thinks a theory of meaning has to be a theory of understanding. But GE rejects that on two counts: there is no theory of meaning and no theory of understanding. GE's theory of language agency is not a theory of grasping meaning. Nothing in GE takes us towards the anti-realist conclusion that flow from Dummett's verificationism. Verificationism is an artefact of thinking that meaning-theory must give a theory of grasping meaning, that will involve public and manifestable practice, and that human capacities to manifest grasp of truth-conditions are limited. Dummettian verificationism is a reductive thesis, a reduction of meaning to forms of understanding. GE is not reductive about meaning at all.

Brandom's Inferential Normative Pragmatism

Brandom thinks that the inevitable alternative to representationalism is inferentialism; the meaning of a sentence is its inferential role; inference is prior to assertion and

propositional content. So the meaning of connectives is given by their introduction and/or elimination rules. The meanings of predicates are given in terms of the material inferences in which those predicates can feature. Referring terms are also characterised by inferential pattern. Brandom theorises inferential role in terms of commitments and score keeping. Sentence-meaning is the impact that utterance of the sentence has on conversational score, where conversational score is a list of entitlements and obligations that a speaker has at a stage in the conversation in virtue of her and others' assertions. Brandom's approach is couched in terms of public utterances, shared language games, and public sanctions. It says little about cognitive and intentional states of speakers.

GE is not a form of inferentialism or normative pragmatism. It is not a theory of meaning or concepts. GE invokes no normative concepts, incompatibility or commitment, to explain the nature of language activity. Moreover, unlike inferentialism, GE is all about cognitive structure. The critique to be given of inferentialism in the course of this book is that inference cannot be the basic explanatory tool to provide illumination about meaning. I shall argue that inferentialism: (a) provides no adequate analysis of assertion, truth-apt sentences, and belief. In short, it fails to provide an explanation of the basic cognitive and semantic categories. (b) It treats logical incompatibility as primitive. (c) It has unacceptable holistic commitments in relation to concepts. (d) It cannot provide an adequate analysis of referring terms; and (e) requires a battery of unexplained normative notions.

Minimalism

GE is a form of expressivism in a sense that still needs clarification. Contemporary expressivism is associated with the movement in *minimalist semantics*.⁸ Semantic minimalism is the view that truth, truth-aptness, assertion, reference and so on, are illuminated by certain platitudes or minimal linguistic practices, such as our espousal of the T-schema, our disposition to replace *the referent of 'O'* with *O*, or the embeddability of declarative sentences in logical compounds. Beneath the surface of this insubstantial semantic exterior, we find a substantial cognitive and affective reality. Declarative sentences

⁸ See Divers and Miller (1994), Price (1990, 1994), Blackburn (1998), Horgan and Timmons (2000).

may be correlated with affective states, like desire, with substantiate belief states. In both cases, the sentences are truth-apt, assertable, believable, and so on, because they meet the minimal conditions to have these non-robust semantic and pragmatic properties. GE is not a form of minimalism. Minimalism about truth and syntactic and semantic categories like truth-aptness and assertion does not work. In this respect, minimalism is like inferentialism, and has similar failings. It attempts to explicate syntactic, semantic, and pragmatic categories at the surface level of activity with sentences without delving below the surface. My contention is that we have to delve beneath the surface to analyze a cognitive noumenon if the syntactic, semantic, and pragmatic phenomena are to be illuminated.

Locke, Grice and Company

GE spends a lot of time explaining the structure of speech-acts. But it would be a grave error to think that GE is a speech-act analysis of meaning. Speech-act analysis of meaning falls within the Lockean tradition, whose best known modern is Grice (1957, 1969, 1971). Grice sought to analyze speaker-meaning, in terms of produce of utterances with certain complex communicative intentions, and word and sentence meaning in terms of practices, regularities or rules relating utterances to speaker-meanings. This program assumes a vast realm of content and contentful states preceding language. Grice uses intention as the glue that relates dead words to living content. But the living content is already there beneath the surface, waiting to be explained somehow. GE is totally opposed to this forcing content into another realm. It isn't part of the neo-Lockean movement found in the work of Schiffer (1972), Bennett (1975), Altson (2000), and Davis (2005).

Fodorian Cognitive Science

Fodor is a theorist who is focussed on providing an account of the home language. He sees the content animating natural language sentences as inherited from the content that animates the inner states of language-using agents. Those inner states are explained by his postulation of a language of thought, which is a syntactically ordered system with representational content by virtue of its causal embedding in a surrounding world. GE is again sympathetic to the thrust of the Fodorian project, but radically differs from it. GE denies that a theory of cognition underpinning surface language is a theory of content, and

denies the idea that surface natural language sentences inherit their content from that possessed by an underlying *mentalese* language. Fodor is an anti-holist: he does not think concepts are constituted by inferential relations. With that GE agrees. But Fodor embraces a representationalist approach to concepts, that is, elements of mentalese have representational relations to external reality. With this GE disagrees, not because it offers another theory of concepts but because it rejects the very idea of a theory of concepts. Cognitive science should shed itself of intentional folk notions including that of concept. That is not because—*pace* Stich (1983) and Churchland (1986)—there really are no intentions or beliefs. There are such things. It is that they play no role in functional analysis of the language agent.

Quinean Naturalism

Quine is perhaps the most important naturalist of the 20th century: he is a logical empiricist tempered by holism and the American pragmatism. Quine gives us two views of language. One is a behaviouristic concept, of language-agents deploying a web of sentences, inter-animating each other within the interior, causally reactive to sensory input at the periphery (1960). The other is interpretative: the radical translator devising translation manuals (1960). One might say the first corresponds to a functional perspective of *Wissenschaft*, and the second to the interpretative stance. Quine thinks that the behaviouristic orientation gives us little by way of illumination, all the work is with an extensional referential treatment of the activities of the radical translator. GE disagrees with Quine's hermeneutic orientation, denying that the home language, studied from the point of view of *Wissenschaft*, is the product of the idiosyncratic, conditioned-dispositions of behaviourism. On the contrary, it is highly structured, with much intersubjective uniformity. Here GE has much in common with the position of Chomsky (1969) on Quine.

GE is not in accord with Quine's attack on analyticity and synonymy, which are perfectly acceptable folk notions, though with little theoretic interest. GE eschews the referential orientation to predicates that generates Quine's problems of inscrutability of reference and indeterminacy of translation.

Quinean naturalism is difficult, as is all naturalism within his shadow. The obstacle is Quine's famous criterion of ontological commitment. This is the principle that where we

seriously quantify over things, we are committed to those things existing as objects. GE, as we have already noted, denies Quine's quantificational principle.

§5 *What Follows*

This is how the rest of the book progresses. The first two chapters are directed towards getting expressivism up and running as a theory of language agency. Chapter 1 outlines the basic explanatory doctrine of expressivism, and addresses an issue that every theory of language function needs to address: the nature of semantic and cognitive categories. These are the categories of assertion, belief, truth-aptness, reason, inference. I show that expressivism, couched in its currently popular form of semantic minimalism, is unable to provide a theory of these categories. It fails at a foundational level. The great irony is that no theory working within the umbrella of orthodox semantics can explain these categories either. This is the invitation to enter below the surface phenomenon of syntax and familiar folk notions of content to explore a speculative realm of cognitive structure. This is how GE begins. I outline the basic structures of GE in chapter 2. In chapter 3, I apply this approach to truth-, reason-, entailment-, and validity-talk, and logical complexity. In chapters 4 and 5, I develop expressivism for predicates like identity, empirical predicates like colour and natural kind terms, cause and disposition, and then value-talk. The approach to value-talk entails that value-sentences can be issued in genuine assertions and believed. There are facts of value, and properties of value as good as any other facts and properties. Yet we embrace value expressivism. In chapters 6 and 7, I provide a functional analysis of referring terms and generality, and explicate the crucial ideas of *object* and *non-object*, and deconstruct Quinean, ontologically loaded, quantification. In chapter 8, I provide a theory of the language agency of talk of facts and truth-making, which simultaneously dissolves any possibility of the metaphysics of fact or being. In chapter 9, I examine issues to do with objectivity through the lens of faultless disagreement. Here we find a distinction between modes of assertion. Assertion can come in an objective and a non-objective form, a distinction that is functionally specifiable. It is in terms of this distinction that we can analyse the language agency

underpinning discourse about *objectivity*. In chapter 10, I offer reflections on meaning, meaning-talk, folk semantics, and meaning scepticism. Finally, chapter 11 provides discussion of GE in relation to realism and metaphysics. GE offers the prospect of realism without metaphysical realism.

Chapter 1

Expressivism and Semantic and Cognitive Categories

§6 Preliminaries

The core thesis of this book is that expressivists about value have dimly seen a deep truth not just about how value-talk functions, but how all domains of talk function. But this insight has been obscured and distorted—at bottom because the insight has been implanted in foreign soil, the soil of Semantics. When sown in its native soil, the theory of language agency, it produces a vigorous plant. To uncover its potential, we have to know exactly what expressivism is and how it works. This chapter focuses on these matters. Below I attempt to articulate what the essential thesis of expressivism is. I then focus on what I take to be the current orthodox vehicle for expressivism. I call it *semantic minimalist expressivism*. It uses minimalist theories about truth, truth-aptness, assertion, belief, etc, as a way of maintaining that expressive discourse can really be truth-apt, involve assertion and belief, and so on. It does not work. Examining its failure is deeply instructive. That failure of minimalist expressivism is just the failure of so-called minimalist semantics. Truth-aptness, assertion, and so on cannot be given a minimalist analysis. To those opposed to expressivism, this will appear to be just one more confirmation that their scepticism about expressivism is well-founded. But here is the punch line: I argue that orthodox semantic theories are in no better position than the semantic minimalist expressivists. They likewise have no adequate theory of truth-aptness, assertion, belief, and so on.

That is the negative part of this chapter. The positive part involves going back to the drawing board to provide a theory of these basic categories. This is centred on a new approach to understanding assertion. The key is what turns out to be a radical twist in our

understanding of the relation of assertion, judgement, belief, and reason. I suggest that all these phenomena of cognitive life are on the same explanatory level. What needs to be uncovered is a cognitive order that underpins all these achievements, a realm of pre-doxastic, pre-truth-apt states. The resulting theory of assertion issues in a theory of belief, reason, and the truth-predicate. It is a functional account of basic cognitive categories that is the backbone of a theory of the language agent. The semantic category of truth-aptness, however, has no functional reduction. We can, however, functionally explicate talk of truth-aptness. The surprise is that the theory that carries out these tasks meets the conditions laid out below for being a form of expressivism, that is, expressivism in its proper place as the centrepiece of a theory of language agency, not as semantics.

§7 The Essence of Expressivism

What is the essence of expressivism? People have largely taken expressivism to be restricted to value claims. Typically, value-expressivism is characterised as the doctrine that holds that in uttering value-sentences speakers do not report facts of value, or refer to properties of value in the world. Nor do they express belief states and produce sentences with propositional content. Rather what they do is *express*, in a sense to be made clear, motivational or effective states. People are attracted to such value-expressivism because it promises to demystify value. It provides an epistemology of value. It seems to explain the relation between value and motivation simply. It has the metaphysical payoff that there is no special metaphysical question about the nature of value.⁹

Expressivism defined in these terms, however, is a revisionist doctrine: it implies that in making value-statements speakers do not report facts, that value-sentences lack propositional content, that there are no beliefs about value, that value-sentences lack truth-value. All this contravenes appearances. But expressivism does not have to be defined in these terms. The doctrine that has the metaphysical and epistemological payoffs noted above does not have to imply these counterintuitive views. Rather the essence of expressivism,

⁹ See Blackburn (1993, 1998), Smith (1995).

according to me, is an explanatory doctrine. It is a claim about what role notions like truth-conditions, reporting, and representation, play in explaining how, say, value-sentences function. We do not analyse the activity of value-statement making in terms of speakers being directed towards representing structured features of reality—facts of value as components of being—but rather in terms of expressing a motivational state of some kind. Motivational states are not truth-apt states—they do not represent in the way that beliefs are meant to represent. Thus utterance of *Killing is wrong* involves expression of my hatred of killing, where hatred is a non-doxastic, non-truth-apt state. This does not imply—as we shall see—that value-sentences are not truth-apt, that they are not manifestations of belief, that value-sentences lack propositional content, that in uttering a value-sentence a speaker reports nothing.

Generalising, I define expressivism for sentences in any domain *D* of discourse as the explanatory doctrine that claims:

Expressivism: A domain *D* of discourse is expressive just in case an account of what speakers do in asserting sentences in *D* does not view speakers as being directed towards relations of representation between their mental states and reality, but rather views speakers as expressing—in some sense of *expression* to be made clear—states of mind, where these states of mind are not representational or truth-apt states.

Such a view only implies that there are no moral facts, value-sentences lack truth-value, and so on, if we have certain assumptions about assertion, belief, fact, truth-aptness, and *expressing*. We shall examine these assumptions later and see how they can be discarded.

Expressivism is a thesis about *expressing*. Clearly the term *expressing* can mean different things. Here are some instances of expressing: (a) expressing thanks by saying *thanks*; (b) expressing disapproval in saying *boo!*; (c) expressing desire in making a request; (d) expressing belief by making an assertion; or (e) expressing a mental state in performing any speech-act. These are all cases of expressing, but are they the kind of expressing that the expressivist postulates to be going on when a speaker asserts a value-sentence and expresses an affective state?

This question may depend on what expectations we have about expressivism. Classical expressivism—a doctrine close to our initial characterisation of expressivism—proposes that sentences in the expressive domain D are not open to truth-assessment.¹⁰ Classical expressivism is a radical, indeed revisionist doctrine, since typical candidates for D , value-statements, really appear to be truth-apt. On the other hand, it is a conservative doctrine in that its argument for non-truth-aptness is a conservative argument: if a sentence expresses a non-cognitive, affective state, it cannot be truth-apt.

I will not concern myself with this orientation further here, for several reasons.¹¹ First, this form of expressivism has been superseded by forms that admit that value-sentences are truth-apt. Furthermore, if expressivism has any hope of going global it cannot accept this form.¹² Let us then take on board the idea that sentences in the domain D are truth-apt, can enter into assertions and manifest beliefs. What conditions of adequacy then must be met by a theory of expressing? Here are three requirements.

(a) A theory of expressing needs to explain why sentences in the expressive domain are open to truth-evaluation. Orders, questions, exclamations are not truth-apt, but utterances of value-sentences are. The question of truth-aptness is related to the question of assertion. It seems utterances of value-sentences can be real assertions, so what is assertion such that expressing affective states can be assertions?

(b) A theory of expressing needs to provide a theory that distinguishes between *reporting* and *expressing* a mental state. U's utterance of *Killing is wrong* is correlated with an affective state, say hatred of killing, and its content depends on that correlation. Yet the sentence *Killing is wrong* is not *about* the disapproval state, it does not mean, *I hate killing*. Its truth-conditions, if it has any, are not given in these terms: '*Killing is wrong*' as uttered by *U* is true iff *U* hates killing. Expressivism about a domain wants to distinguish itself

¹⁰ A paradigm classical expressivist is Ayer (1946). Blackburn (1984) is a more recent proponent.

¹¹ Classical expressivism has been criticised heavily in the literature. First classical expressivism needs an alternative semantic account of logical connectives and validity, but it is far from obvious that it has one. For attempts to provide an account see Blackburn (1984, 1988) and Gibbard (1990). For critiques see Geach (1960), Hale (1993, 2002), Dreier (1996), van Roojen (1996), Kölbel (2002).

¹² Value-expressivism has also taken a dual content form. There are writers who hold that value-sentences are dual-content. They have said-content and value-content, which is the expression of an attitude. Barker (2000) proposes that the expression of the attitude is done through conventional implicature. Versions of dual-content theories include Ridge (2006). These are not theories that have much hope of being globalised.

absolutely from *subjectivism* about that domain according to which sentences in *D* are just reports about subjective non-doxastic states, and are true if and only if speakers have these states.

(c) The theory of expressing needs to explain not only self-standing illocutionary acts, like assertions or questions, but also the use of sentences in *D* in embedded contexts. Value-sentences embed, and enter into logical argument and are subject to truth-evaluation and can enter into belief claims. For example we have the argument in *I*:

I: You ought not to lie. If you ought not to lie, getting your friend to lie is bad.

So, getting your friend to lie is bad.

Embedding presents us with several explanatory issues:

(i) Expressing must be able to function both in self-standing assertions of sentences in *D* and in embedded contexts. It follows that expressing content cannot always be understood as a kind of illocutionary force, since illocutionary forces do not embed.

(ii) How do we understand logical validity? How do we analyse logical compounds themselves? It is not obvious that the expressivist can or should accept the standard treatments in terms of truth-functions.

(iii) If sentences in *D* can enter into arguments as premises or conclusions, it follows that one can have reasons for value-claims, and they in turn can provide reasons. If they can be reason providers, they can also be objects of belief.

With that thought in mind we can now outline the form of a theory expressive discourse, which is a theory of expressing meeting the following conditions:

Expressing: For any sentence *S* in *D*, if production of *S* involves expression of Σ then: (i) *S*'s production is correlated with Σ ; (ii) *S* does not possess truth-conditions specifiable in terms of Σ —*S* does not mean that the speaker has Σ ; (iii) the state Σ itself lacks truth-conditions. (iv) *S* is truth-apt and can embed in logical compounds and enter into relations of reason and be objects of belief.

We can now see that it is far from obvious that the models of expressing given above are what the expressivist is looking for. The expressing of giving *thanks*, of uttering ‘boo!’, of making requests, of expressing a belief in performing an assertion, or of expressing a mental state when we perform any speech-act, all fail to obviously address the subtlety of the problem summed up in **Expressing**. We have a serious explanatory challenge on our hands. There are theories around that make a stab at rising to the challenge. I now look at the approach that currently is receiving most favour.

§8 *Semantic Minimalist Expressivism*

The approach to expressivism that many think fits the bill is what we might call *semantic minimalist expressivism*. The arrival of minimalism about truth was first greeted as a kind of challenge to any form of expressivism. According to minimalists, truth-apt discourse is discourse meeting certain minimal conditions of a syntactic and use-based kind. *Snow is white* is truth-apt because it is grammatically declarative, embeds in logical compounds such as *either snow is white or off-white*, has a meaningful negation, and is used with stable conditions of utterance. Furthermore, truth is just the minimal property defined by our supposed disposition to assert all instances of the T-schema: ‘*Snow is white*’ is true iff *snow is white*. But if this is all that is required for a sentence to be truth-apt and true, value-sentences would appear to be truth-apt and frequently true. *Lying is bad* is embeddable, declarative, has a meaningful negation, and has a disciplined use. And we can associate with it a T-sentence: ‘*Lying is bad*’ is true iff *lying is bad*. If so, it appeared, maintaining any expressivist thesis is dead in the water: value sentences cannot help but be truth-apt since they meet the minimal conditions.¹³

Clearly minimalism is a repudiation of classical expressivism. But not a repudiation of the spirit of expressivism given in **Expressivism** as an explanatory doctrine. Expressivism is not the view that value-sentences lack truth. Its essence is a claim about how utterances of value-sentence work: that in explaining what speakers do we invoke non-representational

¹³ Wright (1992) takes this line.

states. Those drawn to semantic minimalist expressivism present a picture of expressivism in these terms. Truth-apt discourse comes in two forms: representational and expressive. Some areas of discourse meeting the above minimal conditions for truth-aptness involve contents that function as representations. They involve robust assertions that are reports about structured reality, manifesting genuine beliefs, with genuine cognitive content. Statements about the physical world are paradigm instances of this discourse. Other areas meeting the same minimal conditions for truth-aptness involve contents that do not function as representations. Expressive discourse falls into this second region. Both representational and expressive regions are truth-apt, because both meet the minimal conditions for truth-aptness.¹⁴

Semantic minimalist expressivism—let’s call it SME—may look attractive to some, but it is untenable. Minimalism about truth is open to serious objection, as I argue in §20.¹⁵ The objection I focus on here is that SME fails as a theory of expressing.

Minimal Truth-Aptness: The Shadow of Subjectivism

Let us look again at the theory that SME provides of truth-aptness and assertion. The theory proposed is:

Truth-Aptness: (i) *S* is truth-apt because it is declarative, possesses unrestricted embedding—antecedents of conditionals, negations, belief attributions.

(ii) *S* has a disciplined use—stable rules correlate the speaker *U*’s use of *S* with *U*’s possessing some state Σ .

(*Minimal*) *Assertion:* *U* asserts something with *S* in virtue of uttering *S* meeting the conditions for truth-aptness above and represents herself as having the state Σ .

There are three general concerns with this analysis.

It is doubtful that SME distinguishes itself from subjectivism. Compare the sentence *V*, whose utterance is meant to express possession of Σ , and the sentence *I have state Σ* ,

¹⁴ Divers and Miller (1994), Price (1990, 1994), Blackburn (1998), Horgan and Timmons (2000) are examples of theorists carrying out the new wave in different ways. Wright (1992) does not equate his treatment of value-sentences with expressivism, but that is, in part, because he takes expressivism to be first wave expressivism. For Horgan and Timmons the emphasis is on beliefs and judgements either possessing or not possessing representational content, and only secondarily on sentences possessing such content. Blackburn seems to have made the move from first wave to new wave.

¹⁵ See Gupta (1993) for trenchant critiques.

whose utterance constitutes a report about possession of Σ . Both sentences are embeddable and correlated with a state Σ . Why then are V and *I have Σ* not equivalent content-wise? Why isn't utterance of V just a report about possession Σ , as utterance of *I have Σ* is? A reply is that *I have Σ* is correlated with the belief that one has Σ . For this reason its utterance is a report about a state, and not an expression of a state. Unfortunately, we can also say that V is correlated with belief that *one has Σ* . Speakers who use V must believe that they are in Σ . They represent themselves as having Σ and as believing that they have Σ .¹⁶

Another line of thought is that assertions of *I have Σ* involve substantial representation of a psychological state—possession of Σ . Whereas assertions of V do not involve substantial representation of possession of Σ .¹⁷ What is it to *substantially represent*? In uttering V the speaker U cannot help but represent in some sense her possession of Σ . After all, audiences can infer that U has Σ from the fact that U asserts V . V is a sign in some sense of possessing Σ . If so we are left in the dark about why, in fact, utterances of V are not after all reports about subjective states. SME fails to meet our basic desideratum for a theory of expressing: distinguishing expressing from reporting.

Embedding and Truth-aptness

Another problem concerns the thesis that truth-aptness depends on embeddability. To illustrate the pressing nature of the difficulty here, take orders, for example. In uttering *Leave* to someone U conveys somehow her desire that the person leave. In some sense of *expressing*, U expresses her desire. Clear-cut linguistic rules associate U's expressing that U has such a desire with her uttering *Leave*. So why are orders are not truth-apt?¹⁸ It seems that all the semantic minimalist can say is that imperatives produced in orders are not truth-apt because (a) they are not-declarative and (b) they are not unrestrictedly embeddable in logical compounds—orders cannot embed in the antecedents of conditionals or negations. There are two problems with this reply. The first is that not all truth-apt sentences can embed. For example, rhetorical questions are truth-apt but do not embed. Expletives do not embed but

¹⁶ Jackson and Pettit (1998, 2003) raise this kind of objection against expressivism.

¹⁷ Price (1990) makes a proposal along these lines with his distinction between *when-rules* and *rules of representation*.

¹⁸ This issue here is inspired by the discussion in Price (1989).

look truth-apt. One can respond *true* to utterance of *Oh what a beautiful day!*.¹⁹ One response is to discount rhetorical questions and expletives as counterexamples to the embedding principle for the reason that they are non-standard cases. But this move weakens considerably the explanation of why orders are not truth-apt. It is consistent with this position that orders are non-standard truth-apt sentences, like rhetorical questions and expletives. Why are they not non-standard cases of truth-apt sentences?

A more general concern with the embedding account of truth-aptness is whether facts of embedding can constitute truth-aptness. Why do declarative sentences have a certain pattern of embedding and imperatives another? Is this an arbitrary, conventional fact about languages? If so, we might expect languages different from English to allow different kinds of embedding. But, I expect, we do not find this. If so, we need to explain the actual pattern of embedding. But to do that we need an independent account of the nature of embedded contexts, and an independent account of the nature of potentially embeddable sentences, declaratives, imperatives, and interrogatives, resulting in an account of the pattern of embedding we do find. But that means we need an independent account of truth-aptness. So the embedding account cannot be right.

To conclude: it looks like semantic minimalist expressivism (SME) fails to provide a basic treatment of cognitive and semantic categories.²⁰ It is doubtful then that it can be the vehicle of expressivist aspirations.

§9 Semantic and Cognitive Categories: A Problem for Everyone

Those opposed to expressivism will find these results congenial. For them, this is just evidence that expressivism is a misbegotten project trying to square the semantic circle. But if that is what they think, they think too soon. The problems raised above, the nature of

¹⁹ See Stoljar (1993) who raises this difficulty.

²⁰ Another problem for the minimalist treatment of truth-aptness, which I won't discuss here, is conventional implicature. Conventional implicature is a form of non-truth-conditional content that is embeddable in the scope of logical operators. Particles like *even* or *but* introduce conventional implicatures. A sentence *Even Granny got drunk*, carries the truth-conditional content that Granny got drunk, and the conventional-implicature content that Granny is a surprising instance of drunkenness. It can be easily shown that by minimalist criteria conventional implicature is truth-conditional. See Barker (2003, 2007).

assertion and the conditions for a sentence to be truth-apt, are problems for any theory, and, as I hope to show now, orthodox approaches to meaning fail to explain these phenomena. They are in no better position than the minimalist expressivist is. Let us begin with truth-aptness.

Propositions

In orthodox truth-conditional theories people often assert that sentences are truth-apt by virtue of their correlation with propositions. Propositions might be sets of worlds or Russellian states of affairs. The obvious idea is that truth-apt sentences are sentences correlated by semantic rules with propositions. The rules for a language provide what amounts to a recursive definition of *proposition-encoding sentence*. The set of sentences so fixed will be the truth-apt sentences. There are several problems with this idea.

Orders are also correlated by semantic rules with propositions. *Leave the room, Tom!* is correlated with the proposition that Tom leaves the room. One might theorise that sentences are correlated with propositions in a structurally distinct way. One idea, suggested by Alston (2000), is that declarative sentences represent the structure of the proposition in a way that orders or questions do not. There is a kind of constituent isomorphism between sentence and proposition in the truth-apt case. But any such theory would seem to be off on the wrong foot. Perfectly meaningful declarative sentences can fail to be truth-apt because they are used to perform orders or performatives. Interrogative sentences can be used to make assertions and be truth-apt. This suggests that truth-aptness is more a matter of pragmatic features than syntactic features and the supposed internal composition of propositions.

Even leaving aside these worries the theory does not look promising, for it needs to provide: (a) the right theory of propositions as structured entities; (b) the right theory of isomorphism. But meeting these conditions looks difficult. One concern is logically complex propositions: We must suppose that there are entities corresponding to logical constants somehow functioning as constituents of the proposition. Thus *Not-S*, one presumes, constituent-isomorphically encodes the proposition $\langle \text{Neg}\langle S \rangle \rangle$. Neg is the logical entity corresponding to negation, but what is this? The answer, *primitive negativity*, whilst an answer, is not a happy answer. Improving on it is hard to do. Neg cannot be a truth-function,

since that invokes truth as it applies to sentences, and that just brings in the same problem of truth-apt sentences.

Another concern is abbreviation. One may use a predicate *is F* that has a very complex definition in terms of logically complex conditions. But then *O is F* will not bare constituent isomorphism to the proposition. So the theory needs some qualifications in this respect, but here things get tricky. Take any simple sentence *T is F*, which can have an interrogative form *Is T F?* If both encode the same proposition, $\langle T, F \rangle$, then why don't both sentences have a constituent isomorphic relation to it? One might object: the proposition encoded by the question *Is T F?* is really a disjunction with the structure $\langle T, F \rangle$ OR $\langle \text{Neg}\langle T, F \rangle \rangle$. But why isn't *Is T F?* just an abbreviation of a more complex sentence encoding the proposition in a constituent isomorphic way?

Truth

So why not use truth to analyse truth-apt sentences so that sentences that are truth-apt are those that are uttered with the intention that they be true if and only if certain conditions are met, the thought being that orders or questions are not produced with this type of intention? This analysis has several problems. One is circularity: in intending that a sentence be true if and only if *P*, for some *P*, we are intending that a sentence have a certain property under certain conditions. What property is it? One might reply that it is just the property that propositions are capable of having. But how can we just intend to transfer this property to sentences, entities of a completely different kind? One might suggest that the property of truth that we intend sentences to be capable of having is just a primitive property. But if we are going to intend that sentences have it under certain circumstances, how do we know that it is the kind of property that sentences can possess? For example, I cannot intend that sentences be happy. We must be assured then that it is built into the nature of the property that it is the kind of property that can apply to certain classes of sentences. But that just brings us back to the question: which classes of sentences? To reply: just those sentences that we utter with the intention that they are true if and only if *P* is circular.

The other problem with this analysis is that there are sentences to which, it seems, we can assign truth-aptness, but which are too complex for us to understand. Say a super

computer, Super, generates theorems from a bunch of logical and non-logical axioms. We may judge that a certain sentence \mathcal{S} , that Super proves, is true. But its truth-aptness cannot reside in the fact that it is produced with the intention to utter a truth if and only if P , for some P . No one can grasp what P is. The sentence \mathcal{S} is so big that we fall asleep before we finish reading it let alone understanding it. One might think the judgement that \mathcal{S} is truth-apt is simply a commitment to using \mathcal{S} with such intentions. But surely we cannot have any such commitment: it is impossible for me to carry it out.

Belief

Another idea is that the way of explicating truth-aptness is through belief—see Jackson, Smith, and Oppy (1994). The idea is that sentences are truth-apt because they are associated with possible beliefs; beliefs are themselves truth-apt, and so sentences acquire their truth-aptness from association with belief. This theory is problematic. We might first ask what makes beliefs truth-apt. Desires are not truth-apt. What then distinguishes beliefs from desires? Even if we accepted a representational theory of mind, we cannot say that the difference between belief and desire resides in the fact that belief is representational whereas desire is not. Desires can also be representational if beliefs can be. Can we invoke their role in causing behaviour? Both belief and desire contribute to behaviour. Take the desire that one Fs, and the belief that one will F. The desire is something that can combine with other mental states leading to one's F-ing. But likewise the belief is something that can combine with other mental states to lead to one's F-ing. Even if we could find some generally characterisable difference, why does this explain why beliefs are truth-apt and desires not?

A difference that has often been cited as an explanation of the truth-aptness of beliefs is *direction of fit*.²¹ Beliefs have to fit reality—they are wrong if they don't—whereas desires don't have to fit reality, reality is meant to fit them. That is all very well, but without further unpacking of the metaphor of direction of fit, it just looks like a fancy way of describing what has to be explained, namely, that beliefs are open to truth-evaluation, whereas desires aren't. There have been attempts to do so. For example, Smith (1994: 115) suggests that belief

²¹ Austin (1955) and Searle (1983) introduce this distinction. It goes back to Anscombe (1957).

differs from desire in functional role. Roughly, a belief that P is a state that tends to disappear if one comes to *perceive* that not-P, whereas a desire tends to persist, inducing in the subject a tendency to bring it about that P. There are various things wrong with this idea, but the main one is circularity: Perception that not-P covertly makes reference to believing not-P.²²

I hold that there is a functional analysis of the difference between belief and desire. I shall provide it below in §11 and §15. But direction of fit does not explain the difference. Furthermore this analysis, as we shall see, takes us beyond orthodox conceptions of content. Moreover, even with the illumination of direction of fit, we still need to articulate the relation that beliefs have to bear to sentences for such sentences to be truth-apt. It is sentences associated with the expression of belief that are meant to be truth-apt. But what is *expressing*? All sorts of relations can fall under the term *expressing*. Orders express beliefs but they are not truth-apt. A second matter is that the proposed analysis runs up again the problem of comprehension transcendent sentences. The super computer Super's sentence $\$$ is not correlated with any belief. Super computers do not, necessarily, have belief states.

Assertion

It might seem that we could deal with these last difficulties facing the belief-account of truth-aptness by moving from belief to assertion. It is the correlation of sentences with acts of assertion that makes them truth-apt; sentences used in orders are not correlated with assertions, and so they are not truth-apt. But what does this idea of correlation amount to? Embedded sentences can be truth-apt but are not asserted. One might suggest that they are viewed as *potentially asserted*. But now we encounter again the challenge of comprehension transcendent sentences. How do we view the sentence $\$$ proved by Super as potentially asserted? Super computers do not make assertions.

The other problem is with assertion itself. What is it? Grice proposed: U asserts that *P* if and only if she utters *S* intending that others believe that *P*, or that others believe that she believes that *P*.²³ These intentions are reflexive in form: the speaker U aims to get an

²² See Sobel and Copp (2001) for development of this criticism and application of it to related views, such as Humberstone (1992).

²³ That's rough: Grice (1957, 1969) required that the intending referred to here was embedded in a system of reflexive communicative intentions. We need not concern ourselves with these here.

audience H to gain a certain state r —say a belief, desire, intention—by attempting to get the audience to recognise this very intention.²⁴ The task then for an analysis of assertions—according to Grice—is to determine which state r U intends H to gain when U asserts something. Grice (1971) toys with two proposals:

a belief that P.

a belief that U believes that P.

But both are problematic. It is quite possible for U to assert that P but to be perfectly indifferent as to audience response. Say Jane is paid to make announcements but has no concern about the epistemic states of her audience, and so she lacks intentions of this kind—see Alston (2000).

Bach and Harnish (1979) and Recanati (1986) suggest that in assertion the state that U intends her audience to possess is:

that H have a reason to believe that U believes that P.

But even this is too much, since an assertor U may be indifferent as to whether audiences come to have reasons to believe that U believes that P. Grice's assumption that assertion is essentially an act of communication—and so involves intending audiences to have certain kinds of doxastic or epistemic state—is open to dispute.²⁵

Perhaps assertions are acts of aiming to utter true sentences. But this proposal faces the problem of insincerity. One needs to talk rather of *representing oneself* as intending to utter a truth. But illuminating *representing-oneself-as* is difficult—see Pagin (2005). The proposal also invokes the idea of a true sentence, and thus truth-aptness, which is the very thing we are attempting to clarify.

We might think there are other approaches to assertion that will work. We shall consider Brandom's theory below when we examine *inferentialism*.²⁶

²⁴ See Grice (1969), Strawson (1964), and Schiffer (1972) for examination of the structure of such intentions.

²⁵ See Alston (2000) for a critique of Gricean style approaches to assertion.

²⁶ Another idea is analyse truth-aptness indirectly, by analyzing non-truth-aptness. Davidson (1979) tries to explain why orders are not truth-apt by the hypothesis that they are really sentence pairs of the form: *You will leave. That's an order*. Davidson claims that pairs or triples of sentences that are not conjunctions are not truth-apt. But this is doubtful. Why can't texts—non-conjunctive sequences of sentences—be true? One can say *all that is true* to a text, but not to an order. Davidson might object that texts are implicitly conjunctions, but then he owes us a theory of why orders are not implicitly conjunctions.

§10 Inferentialism and Interpretationism

Inferentialism is the view that the way to understand meaning, content, and thought, is through inference. Perhaps inference, understood appropriately, is the way to a theory of assertion, truth-aptness, perhaps even belief. I shall argue not. Brandom (1994, 2005) offers an inferentialist analysis of assertion according to which assertion is a move in the game of giving and asking for reasons. In a 1983 paper he proposes that:

U asserts that *S* iff (i) U obligates herself to justify *S*, if asked to, and (ii) permits speakers to use *S* as a premise in arguments.

Taken at face value this analysis is problematic. The concern is circularity. Both (i) and (ii) specify conditions that really make reference to the assertion of *S*, and it is not entirely obvious how this reference can be removed. What one justifies is not a sentence but assertion of a sentence. It is not a sentence that appears as a premise in an argument but a sentence that is asserted. So really, to be perfectly explicit, the analysis should be:

U asserts that *S* iff (i) U obligates herself to justify her assertion of *S*, if asked to, and (ii) permits speakers to use *S* as an asserted premise in arguments.

But then the analysis is plainly circular.

Perhaps we are not appreciating the holistic orientation of inferentialism. The explanatory ground for both content and assertion is inferential patterns. Inferences can be material, as in I_1 below, or formal as in I_2 :

Pegasus is a horse \vdash *Pegasus has a head and four legs*;
Pegasus is a horse \vdash *either Pegasus flies or he does not*.

An inference rule is a sequence of sentence productions allowed by a language community. Since sentences gain their meaning from the kinds of inferences they enter into, inference cannot be explicated by appeal to prior sentence-meanings. The production of sentences in these inference patterns is not analysed in terms of assertion, since we are using inference to explain assertion. Inference is a normative relation between sign productions. Thus if $S \vdash R$

holds as an inference for some community, then in tokening an instance of S , U is taken as committed to tokening R , in the sense that if U tokens something incompatible, U is sanctioned. Inference patterns fix meanings. A word or sentence's meaning is its inferential impact on conversational score. To assert is to utter a sentence and obligate oneself to produce sentences, the production of which will in turn carry obligations to provide further such sentences S . The form of the obligations incurred is fixed by inference patterns. Asserting *snow is white* brings the obligation to provide sentences fitting into the inference forms: *X told me and X knows*, or *I saw snow*, etc.

Does inferentialism work? I argue that it is deeply problematic for four reasons below.

1. What is inference? Suppose U produces *O is a cat* and then *O has ears*. Let us suppose that the sequence $X \text{ is a cat} \vdash X \text{ has ears}$ is an inferential sequence for U 's language community. Has U made an inference? Not necessarily. Perhaps U saw the cat O , and did not infer from O 's being a cat that he had ears but from looking at the animal itself. Something else is needed for inference. Perhaps what is required is that U be disposed to insert *therefore* or *so* between the two sentence tokens. But that cannot be right. Particles like *therefore* make explicit an inference, they do not explain what inference is. Where there is inference there is action. That suggests that inference is action caused in a certain way. Inference requires that U 's tokening of *O is a cat* caused the tokening of *O has ears*. But what form must the causation take? In the case above, maybe U 's utterance of *O is a cat* is partly responsible for causing *O has ears*, because it supplied a pronoun antecedent, and also made U think of ears. But U still made the utterance primarily because she perceived an eared cat. So the token was a causal factor, but that still does not make it an inference. What kind of causation makes it inference? I think the only answer is this: What underpinned U 's assertion was a mental element and her processing of this mental element led to tokening of a second mental element, which then underpinned production of the second sentence. What are these mental elements? Whatever they are, they must have something to do with sentence content. But that means content cannot be fixed by inference, *pace* Brandom.

2. Inference commitment is partly analysed through incompatibility. The structured sequence $S \vdash R$ is an inference for some language community because tokenings of the form

S commit speakers to tokening *R*, not in the sense that they must invariably token *R*, but that if they token something incompatible with *R* they are sanctioned. But incompatibility is a content presupposing notion. However, Brandom's explanatory foundation is the production of sentences in such sequences in a normative framework that does not yet presume content, rather content arises from networks of sentences inter-activated by inference.

3. Brandom says that in making an assertion we obligate ourselves to provide reasons, but how do we obligate ourselves? One idea is that we just have a disposition to provide reasons. But that will not do. First, assertions do not require such dispositions. I can assert without any disposition to provide anyone with reasons. I just make an announcement. One might object that *U* must be disposed to provide reason to herself. That may be true. How is that the normative notion of obligating oneself? Another is that the social context introduces obligation. But one does not have to be in any social context to assert. Alone late at night I sub-vocalise: *She does not love me*. Perhaps I am obligating myself to myself. But what did I do to do that? The insight may have come to me in a spontaneous flash, with no prior relevant interior dialogue to provide an interior context.

4. Does the inferentialism account distinguish assertions from orders? One might say *yes*, because orders do not enter into the game of giving and receiving reasons. But is that so clear given what Brandom has told us about what giving reasons is? If one utters: *This is my house, You people are dangerous intruders. Leave!* It seems I do not do anything that attracts censure. If I were to say *please stay* rather than *leave*, that would attract censure. Similarly, utterance of *You, leave! You are able to leave* is fine, whereas *You, leave! You are physically incapable of leaving* involves some kind of incompatibility. These facts suggest that imperatives can enter into inferential relations. So why isn't my utterance of *leave!* an assertion? One might reply that imperatives cannot feature as premises in arguments. One cannot say *Leave! Therefore, you are capable of leaving*. True! But Brandom cannot cite this as an explanation. *Therefore* is just a word that makes explicit inferential sequences. So to explain its function we require a prior explanation of inferential sequence. But that is what we lack. If so, there is no explanation of why orders are not assertions.

Interpretationism

We have failed to find a theory of truth-aptness. We have failed to find a theory of assertion. We have failed to explain the difference between belief and desire.

One response to this predicament is that the concepts we possess of these semantic and cognitive categories are analytically irreducible, but that their natures are somehow revealed in our practice of using them when we interpret the speech of others. The content of these concepts is revealed, say, in the inferential relations between attribution of truth-aptness, assertion, belief, and other interpretative assignments. For example, an inference we accept is: *If S is asserted, it is truth-apt*. However, we need more than just a few inter-linking inferences to provide a framework for a stable application of these concepts. We need to know what interpretative pressures favour attribution of truth-aptness, assertion, belief, and related notions to behaviours.

Moreover, lacking a theory of assertion and truth-aptness we lack a theory of other illocutionary acts. The picture is this: somehow all these different kinds of phenomena—cognitive, pragmatic, and semantic—are clarified by the form of interpretative theory. Maybe this can come about, but it is far from obvious how it can do so. But as we saw in the introduction, interpretative theory cannot tell us what interpretation is, so how can we expect it to inform us about the nature of these cognitive and semantic categories?

§11 Expressivism Returns as a Theory of Language agency

Orthodox semantic theorists are quick to attack expressivists on the grounds that they fail to provide an adequate theory of embedded value-sentences, of truth-aptness, of assertion, and so on. In other words, that they fail to provide a theory of semantic and cognitive categories. But, as we have shown, orthodoxy also fails to provide an adequate theory of such categories. Semantics is hoisted on its own petard. What is the way out of the quagmire? Ironically, I shall argue that the solution lies with expressivism. We have been attempting to analyse truth-aptness in terms of surface syntactic features, such as embeddability, and appeal to propositions correlated somehow with sentences through semantic rules, truth, assertion,

belief, and inference. This orientation has not proven fruitful. There seems to be no way of using some of the elements in the set {truth-aptness, assertion, belief, inference} to explain other members. Moreover, we cannot treat these categories as primitive. Nevertheless, these categories are related to each other somehow. We are led to only one conclusion. These categories are explained by something outside the set: a domain of phenomena that falls outside these semantic and cognitive categories.


What has this got to do with expressivism? Expressivists about value seek to explain value-assertions by appeal to affective states, which is to say, non-truth-apt mental states. This raises the hackles of orthodox theorists, who say assertions are the outer apparel of beliefs, which are the animating cores of assertion. But now we have reason to think that this orthodox response to expressivism is wrong. Belief cannot be used to explain assertion, since belief just raises the problem of truth-aptness at a different level. This is the problem of direction of fit. We can now see that perhaps the expressivists were not entirely deluded in looking outside the domain of doxastic and truth-apt states to explain what goes on in the assertion of value-sentences.

If we entertain for a moment the strategy of the expressivist we see that it fits in with what I have sketched as the form of an explication of our categories. The members of the set {truth-aptness, assertion, belief, judgement, inference} are explicated by something outside the set that is non-doxastic and non-truth-apt. That is reminiscent of what the expressivist was proposing for the case of value-sentences, though the expressivist obscured her own project in attempting to integrate it within Semantics. But the core idea is right: value-assertion is not to be explained in terms of belief. The message to be extracted now is that no form of assertion is to be explicated in terms of belief and other truth-apt states. Here is the basic idea of the program that is based in this seed idea.

The mental element that underpins all assertion is never doxastic or truth-apt. In explaining the activity that is assertion we do not view assertion as an activity that is preceded by a realm of content, embedded in some attitudinal structure, with assertion as the outer apparel of that embedded content. Rather we can only understand assertion in terms of the structures of a language agency system that comprises a system of signalling and its

underpinning network of pre-truth-apt and pre-doxastic states, that make up what I have called the NC—or natural cognitive system—see §2. The term *cognitive* here has no connotation of belief or propositional content. Assertion, judgement, belief, and inference are all on the same cognitive explanatory level. These phenomena are exercises of language agency. Belief is not the inner core of assertion, and assertion its outer apparel. As we shall see, belief, judgement, and inference, like assertion, have functionally specifiable structures. The basic operation that they all share is that they are linked to production of sentences *defending* what I call *Π-properties*, these being one kind of pre-doxastic states that make up the NC system. I use the term *Π-property* because it is a neutral one, without folk connotation, and we are now delving beneath surface of folk concepts. To defend a Π-property is to signal a disposition to indicate *grounds* for that Π-property. These are somewhat vague ideas at this point, but we shall give them substance in the next two chapters. Our concern is with form at this point.

We can represent the form of the present proposal in relation to assertion, judgement, and belief as follows:



On the left I present the orthodox view. Assertion is a surface phenomenon underpinned by some special relation to belief (and maybe judgement). Assertion is the manifestation and outer apparel of belief. In this conception, desire cannot be the mental element of assertion. Hence orthodoxy denies the possibility of moral assertion being *expressive* of desire as the expressivist contends. Truth-aptness is found at the level of belief, and assertions inherit their truth-aptness, somehow, from belief. Contrast now the right-hand side, which is the new hypothesis being formulated. Belief, assertion, and judgement are all at the same level. They are linked to production of symbols that are defences of pre-doxastic, pre-truth-apt states. These states, Π-properties, can be affective states, or they may be non-affective states. Indeed, Π-properties can come in many different functional kinds, as we shall see.

On the new hypothesis, truth-aptness is only found at the upper level and not at the lower. It is only at the level of assertions that we find truth-apt states, that is, symbol-producing acts of defending Π -states. The upper level does not inherit its truth-aptness from special kinds of states below. Or in other words, if we must talk in these terms, the direction of fit of the upper level does not come from the direction of fit of the lower level.

If a speaker U has a belief, U has a dispositions to defend a Π -property that she actually possesses. Judgements are private assertions. In this conception, there is no problem at all with some assertions and beliefs expressing affective states. They are just as much assertions and beliefs as those assertions and beliefs that express non-affective states. Affective states are at a different explanatory level from beliefs. This is the functional difference between beliefs and desires (and more generally, affective states). Some beliefs, those about value as we shall see, have affective states as functional parts. These beliefs are dispositions to defend affective states. This is not to say that some beliefs are desires—*besires*, as some theorists have postulated Altham (1986). We retain the Humean distinction between desires and beliefs.²⁷

This is not to deny that beliefs are states that potentially cause non-verbal behaviour. To be disposed to defend a state alters how one potentially behaves non-verbally. The functional properties of merely possessing a Π -property Ψ , and possessing Ψ and being disposed to defend it are different.

What follows

That is the basic idea for a functional analysis of assertion, belief, and judgement. We shall have to say more to apply this structural account to inference, and generally to what we call *displays of reason*. The semantic category of truth-apt sentences, is a somewhat different story. It has no functional characterisation. As we saw \mathcal{S} , proved by Super, the super computer, cannot be meaningfully tokened by us. Yet it is truth-apt. Truth-aptness is like truth and other semantic properties such as having meaning, because it has no functional specification. This is not to affirm that it is primitive in any objectionable sense. We can fully

²⁷ See Smith 1986 for an examination of that distinction and arguments for it.

explicate our talk of truth-aptness; that is, we can provide a functional analysis of that talk in terms of the structures of language agency.

What follows next is laying the foundations of a theory of language agency. This is our first step on the path of Global Expressivism and the dismantling of Semantics.

Chapter 2

How to be a Global Expressivist

§12 Basic Structures of Language Agency

Here is the hypothesis: assertion, belief, judgement, inferential practice, and production of truth-apt sentences are cognitive superstructures riding upon a mental sub-structure whose elements are not yet liable to truth-evaluation, and not yet in possession of propositional content, so cannot be called beliefs or inferences. The sub-structure comprises a network of mental elements that are pre-doxastic states that are not liable to truth-evaluation. This I have called the NC (natural cognitive) system. It is upon this system that the achievements of assertion, judgement, belief, assertion, and reason rest. What we call assertion, belief, and episodes of reasoning are productions of, or dispositions to produce, symbols caused in the right way by elements of the NC system. To analyse these phenomena, we describe the pattern of causal production and the nature of the elements of the NC system. Note that we are offering functional analyses of what it is to assert or believe, not of what it is to assert or believe that *P* for some contentful *P*. That would require a functional reduction of content, but there is not any. In what follows I sketch out the basic features of this new orientation, in terms of which we can eventually derive a functional characterisation of assertion, belief, inference, utterances of truth-apt sentences, and so on.

We can divide the language agent into two levels: the pre-symbol level—the sub-structure, which I call the NC system—and the symbol-level—the superstructure. The whole is the language-agent or signalling system. The symbol-level—the level at which we find acts of symbol production or dispositions thereto—is the level of truth-aptness and propositional content. Symbol production involves the *signalling* of mental states, elements of the NC

system. Assertions are signalings of certain kinds of states: *Π-properties*. Beliefs are dispositions to signal Π -properties, where the speaker really possesses the state. Inferences are specific sequences of assertoric-signalling, and so on. The symbols used in signalling may be external public speech or internal imagined speech of inner dialogue.

Signalling has a causal structure, which we shall characterize in detail in §13 below. The basic idea is that in a signalling exchange of the simplest type, there is a mirroring structure of this kind: U produces a symbol s caused by a state Ψ in U, and H receives s by simulating a homologue state Ω in her own system. I say *homologue* since I do not claim that for a successful exchange to occur—a successful act of communication—the states Ψ and Ω are identical in functional character. These states Ψ and Ω are not public states. Speakers are not directly conscious of them. Agents that signal have no grasp of the nature of the states they signal and do not require one. These inner states are certainly not meanings or the fundamental bearers of meaning. Signalling is a slightly misleading term, since the term standardly implies that the states signalled are public. However it will have to do since there are simply no folk terms that perfectly fit the phenomena we are describing. Understanding involves such reception and simulation. I call this reception and simulation S-processing—see §1-2. To be a language agent involves two things: production of symbols based in pre-doxastic states, and disposition to S-processing incoming symbol productions.

Assertion, Π-properties, and C-laws

The NC system is a network of mental elements. Speech-acts are signalings of elements of the NC system. Assertions, orders, and so on, are signalings of elements of the NC system of a specific kind. We shall begin soon describing the structure of these elements. We get to know about mental elements of the NC system through reverse engineering. We postulate them because they help us understand surface production, and speech-acts, which are symbol productions part of whose causal antecedents are these states. So our initial understanding of them is relational in this sense, though later we get to know more about the intrinsic natures. I now broadly characterise two general functions of the NC system.

The first surface phenomenon is that we encounter sentences as possessing different contents or subject matters. What underlies that judgement is our processing sentences in

terms of mental states called *Π-properties*.²⁸ Π-properties include familiar states, perceptual states, affective states, dispositions to perform speech acts, and so on. But they also include states based in such familiar states, where *basing* involves a process of derivation that generates a more abstract functional state. (I describe this process below.) Π-properties are not propositional contents and do not possess propositional contents, as we shall see.

The key functional feature of Π-properties is that for the individual speaker U they correlate with the kinds of assertions U can make. If U makes an assertion using a colour predicate, *red*, then there is a Π-property based in perceptual states. If the assertion is evaluative then the characterising Π-property is based on an affective state such as approval. If the assertion is about identity or existence the correlated Π-properties will be based upon dispositions of a metalinguistic kind: dispositions to use symbols in a certain way. If the assertion is logically complex or about reason, then the correlated Π-properties will be based in dispositions to process Π-properties in a certain way. If a negation, it will involve a certain kind of processing, disjunction another, and so on. There are other kinds of assertions with correlated Π-properties. We shall examine these below.

Assertions can then be divided into categories: empirical, evaluative, formal, or logical, according to the structure of the kind of Π-property correlated with them. Within such categories we can, for a given speaker, provide a fine-grained description of the specific Π-properties correlated with sentences on specific uses. All this holds for individual speakers at certain times. We might ask what are the structural commonalities across speakers? The answer is that there is functional identity up to a certain point. For certain vocabularies, those to do with the empirical world and calculation, there will not be functional identity, which is to say, the Π-property U associates with her sentence *The flag is red* may not be the same as that H associates with that sentence. I say more about this below.

²⁸ The concept of a Π-property is not part of our ordinary folk conception of semantics or pragmatics. Π-properties are theoretical entities for what is, I shall be urging, a reconceptualisation of certain commitments of cognitive science. Note by *pre-doxastic* some might mean *sub-doxastic* in the standard sense: unconscious, inferentially isolated, representational states. But this connotation should be set aside in what follows: Π-properties are not to be thought of as representational or inferentially isolated. However, they are sub-doxastic in the sense that their structures are not accessible to ordinary intuition.

Assertion is not simply signalling of Π -properties. To understand assertion we need to examine the second general feature of the NC system. This feature is reflected in another surface phenomenon: *reason*. Brandom (1983) analysed assertion in terms of obligations to give reasons for what is asserted—see §9. We rejected any attempt to reduce assertion to reason. Nevertheless, Brandom is right in identifying, at the phenomenal level, a link between assertion and reason. The NC system is a network of Π -properties. Part of the linking of Π -properties is through what I call *grounding relations*, which are certain law-like relations holding between Π -properties. Grounding is a natural relation that links Π -properties. In a speaker U, Ψ grounds Ω just in case in U's cognitive system there is a kind of cognitive law, or C-law, linking production of Ψ to potential production of Ω . This idea of a network is a familiar one, as already noted in §2. The NC system is a network of pre-doxastic states linked by C-laws, these being, in specific cases, precursors of activities of inference.

Higher mammals, we might suppose, have NC systems with grounding relations. They can reason to this extent: they solve problems through cognition. What distinguishes humans is not so much capacity to engage in rational cognition, but the fact that they engage in signalling the grounding relations of their systems. Signalling of grounding relations helps, in a complex world, in which agents, in acquiring information from other agents, need to be discerning. To acquire information about the world from a speaker U, H needs to have faith in U, and that means some idea of U's reasons (reason and experience being the only access to how things are.) To provide information U needs to have faith in her capacities, and to be able to provide reasons. Reason is the phenomenal surface of grounding.

This is the origin of assertion. Assertion is that act of signalling a disposition to signal states that ground a Π -property Ψ . The proposal then is:

Assertion: U performs an assertion iff U produces a sentence *S* signalling her state of having a ground-signalling disposition in relation to a Π -property Ψ .

Here is an extremely simple case of assertion so understood. If a child says *dog*, pointing, then he asserts that a dog or dogs are present. Utterance of *dog* signals a certain Π -property: a kind of functional state linked to perception that we shall describe in §28-31. (The child is

not referring to this state. Signalling is not referring. Reference has no role to play in functional analysis.) The perceptual state grounds the Π -property signalled by *dog*. U's signalling of her grounding relations may be verbal or non-verbal. The verbal may involve full-fledged assertion of premises. Non-verbal may involve demonstrations of various kinds.

Let us call *defending* a Π -property Ψ the act in which a speaker signals a disposition to signal grounds for Ψ . Assertion is that act in which speakers publicly defend Π -properties through symbols in public space. Judgement is that act in which non-public symbols are produced—in imagination—and we defend Π -properties. Belief is a disposition to defend a Π -property, actually possessed by the speaker. Orders involve mere signalling of desire-states, and so on. Non-asserted but truth-apt sentences are sentences produced in *proto-assertions*. Proto-assertion, as we shall describe below, are instances of proto-signalling, which is an essential precursor to any act of signalling. Such are the rudiments of a structural analysis of assertion, contrasting it with other speech-acts. Assertions are manifestations of agency, not mere reflexes. What we are describing is the ensemble of causal elements that agents bring into being in displaying this agency.

§13 Signalling Systems and Communicative Exchanges

The proposal, the core thesis in GE is that to assert is to produce a sentence defending a pre-doxastic state Ψ . To defend Ψ is to signal a ground-signalling disposition in relation to possessing Ψ . A ground-signalling disposition is a state in which U is disposed to signal, under prompting (either external or internal), her possession of Π -properties that ground her possessing Ψ .²⁹ This theory invokes the following: signalling, prompts, disposition, and grounding. I deal with grounding and its relation to reason in chapter 3. I give an analysis of disposition talk in §33 in chapter 4. For now let us focus on these other elements. We begin with the central communicative notion of signalling.

²⁹ *Signalling* and *indicating* are synonymous. I use them as stylistic variants.

Signalling Systems, Communication, and Interpretation

Central to my proposal about assertion is the notion of *signalling a state*. Signalling is a phenomenon that underpins activities that we call communication. Communication does not require intentions of the kind postulated by Grice (1957), that is, reflexive intentions to affect audiences' mental states. Rather the basic features of what we call communication can be articulated without reference to intention but in terms of functionally specified dispositions holding for communicating agents. If O_1 communicates to O_2 , then certain regularities have to hold in relation to behaviours common to O_1 and O_2 that we can think of as a system. A communicative system involves production and reception of signs. Signs are broadly construed as bodily gestures, chemical emissions, written symbols, verbal symbols, etc. In short, any event type producible and perceivable by the agents.

Agents in a communicative system have abilities to produce and receive items of a common set of signs. Productions of signs, in higher creatures, are not mere reflexes. We can talk of agents being *directed towards* producing s . This being directed towards does not imply intention. In engaging in directed behaviour there is a purpose, in the sense that the agent is directed to an end state, but does not have to have beliefs about that state or intentions or concepts. Typically, an agent O in producing a sign s is directed towards producing s as a causal result of being in Σ along with other causal factors. The production of s is not a mere reflex produced merely by being in Σ . Production of signs is regulated by what I call *repertoire dispositions*. Repertoire dispositions link being directed towards producing a sign as the causal result of Σ and possibly producing s as a result. For the agent O , there may be more than one sign associated with Σ . I denote a repertoire disposition by $RD[\Sigma, s]$. So $RD[\Sigma, s]$ is a disposition of the form:

If O is directed towards producing a sign caused by Σ , O may produce s .

The *may* here is epistemic and not normative permissiveness. It is one possible outcome given the system is functioning normally. Where O produces s through the repertoire disposition $RD[\Sigma, s]$ this will not be an automatic reflex. There may be several options of signs for O to produce given that it has Σ , and O may simply not do anything.

There are two ways in which in production of s can be caused by $RD[\Sigma, s]$. One way is that $RD[\Sigma, s]$ is manifested. O has the state Σ and that triggers $RD[\Sigma, s]$ leading to its manifestation, so that Σ in O is part of the causal history of production of s . The other way is that O lacks Σ . Nevertheless, O is directed towards producing s as a causal result of $RD[\Sigma, s]$.

Here is an example: Let us say that lions have a signalling system of roars. Kimba has a repertoire disposition, $RD[\text{Anger}, \text{roar}]$. To denote this state as *anger* involves an interpretative stance, but we are not supposing that anger across individuals has a functional essence. But it is convenient to talk within a partially interpretative framework for the time being. Later I shall be more explicit about functional structures of states. Kimba may then roar as the result of anger. So, his production of roar r is caused by $RD[\text{Anger}, \text{roar}]$ and his directedness to its being a cause. But Kimba may or may not have an angry state. Let us call production of a sign s through a repertoire disposition a *proto-signalling*. That is:

Proto-Signalling: O proto-signals Σ by producing s iff O produce s and (i) O has $RD[\Sigma, s]$, (ii) O is directed towards producing s as a causal result of $RD[\Sigma, s]$.

Proto-signalling is the foundation of all communicative moves in a communicative system. To proto-signal is not yet to signal, but more of that in a moment.

Let us call participants in a communicative system *signalling systems*. One aspect of signalling systems is their ability to produce signs through repertoire dispositions. The other aspect is the reception and processing of signs. Say an agent O' proto-signals her state Ψ by s , and O may perceive the sign and then interpret it. O does not have direct access to O' 's inner state Ψ . But O does not need that to interpret O' 's sign s . Interpreting is not grounded in processing that allows direct perception of others' states. Rather what happens is that O has a spontaneous tendency to *simulate* O' 's proto-signalling. That is, O perceives s , and a repertoire disposition $RD[\Sigma, s]$ in her system is activated. It is as if O is about to produce s herself. But her system does not go that far. Rather a part of her mind-brain or NC system *lights up*. If we thought in terms of mirror neurones then, we might say, mirroring through

neuronal activation occurs in O's NC system.³⁰ I call this kind of simulation *S-processing*—see §1-2. We define S-processing in these terms:

O S-processes O's production of *s* in which O' proto-signals Ψ , iff O enters a state in which O simulates O's proto-signalling act, that is, O simulates production of a proto-act using some repertoire disposition $RD[\Sigma, s]$ she possesses.

For short, we shall say that O, in S-processing, *assigns* a state $RD[\Sigma, s]$ to *s*. We make no assumption that the internal states of O and O'—that is, Σ and Ψ and the corresponding repertoire dispositions—are type identical to terms of functional structure. They may indeed differ in various ways. At the moment, we are interested in necessary structural conditions for communication, but not in specifying necessary and sufficient conditions. We cannot do that because communication implies contentfulness, but there is no functional reduction of content. This idea is developed below in some detail. Nor have we yet provided any account of why O simulates O's production of *s* with one repertoire disposition $RD[\Sigma, s]$ rather than another, assuming that *s* for O is linked to more than one state.

The basic necessary conditions for a communicative system are then:

A: For each O in the group of communicating agents there is a set of pairs with each pair comprising a sign *s* state Σ such that:

- (i) O possesses a repertoire disposition $RD[\Sigma, s]$.
- (ii) In receiving *s* produced by O', O S-processes *s* in terms of $RD[\Sigma, s]$.

B: O and O' share at least some sign forms.

C: The states that O and O' correlate with a given sign *s* share functional features in common, but do not have to be functionally identical.

Agents that are signalling systems need not only the ability to produce and perceive signs, they also need sensitivity to contexts for production. They have capacities to detect that agents are receptive to proto-signallings of certain kinds, which is to say that such agents

³⁰ Mirror neurones have been the subject of much recent discussion in cognitive science. See Stamenov and Gallese (2002).

have certain informational needs. Let's call *prompts* states that signal receptivity to certain forms of signalling on the part of another agent. Abilities to *recognize prompting* will be parts of the capacities of sophisticated communicative agents. Sensitivity to prompting does not require beliefs about being prompted. It does require broadly speaking perceptual capacities and systems for detecting agency in others. I shall discuss functional structures underpinning perceptual states and their relations to belief in §28-32.

Signalling

The basic feature of signalling activity is proto-signalling and ability to simulate proto-signallings. But what about signalling itself? Kimba the lion produces roar *r*, proto-signalling an aggressive stance, but does Kimba really signal his anger? Kimba can signal anger whether or not he is angry. Perhaps he only wants to scare the adolescent male, Simba, in front of him, but really is happy with Simba. We might define signalling in interpretative terms: O signals a state with *s* if and only if one would be right to infer that O has the state from O's production of *s*. But this is a normative, interpretative specification, which is simply the wrong orientation. It requires having independent access to the state that O has, which, typically, receiving agents do not have. Furthermore, what we call interpretation does not deal with the inner states of other speakers or signalling agents. We need a functional analysis of signalling, but this, as we shall now see, is a subtle matter.

Play and Conflict

Kimba has a repertoire disposition: RD[Anger, *roar*]. In producing *r*, Kimba may be directed towards producing *r* as a causal result of this disposition. This is a proto-signalling of a state: anger. But signalling only happens when further conditions are met: the absence of play and the absence of conflict.

Signalling agents can engage in play, which is a kind of pretence using the system. Suppose some lion cubs are playing. They roar at each other. Their roars are not sincere indications of aggressive stances, nor are they insincere signallings. The cubs do not get scared of each other, and have no inclination to be. Nor are they puzzled or unsure of each other. In this case there is proto-signalling of anger-states. But this is put into the context of play. Their roars are not so loud with minimal facial gestures. Their paw pats are soft with

unextended claws. We find signs, toned down roars, which indicate play. How? Signalling play is like signalling any other state. The signs of play have repertoire dispositions that underpin their use, and through which there can be proto-signalling of play status. Where O produces a sign, a roar caused by $RD[\text{Anger}, \text{roar}]$, but is not directed towards the state Anger causing production of the roar, O can add a signal: *toned-down-roar* to proto-signal play status. Signs of play are there to signal that certain states are not causally active in the production of other signs. In short there are repertoire dispositions of this general kind:

If O produces s caused by $RD[\Psi, s]$ and is directed towards producing a sign caused by her suppression of Ψ as a cause of s , O may produce r .

A receiver O' S-processes production of r just as she does s . That is, O' simulates the structure of the proto-signalling produced with s , and the proto-signalling of a play state. If there is a proto-signalling of a play-state in relation to s , there is no signalling with s .

We still have not articulated the conditions for signalling. Even if there is no indication of play, O may still not be signalling something with s . That is because signalling can be undermined in other ways. The other way is *conflicting behaviour*. Suppose for Kimba that wagging tails proto-signals happiness: $RD[\text{Happy}, \text{wag}]$. Moreover, happy is a state that, in Kimba, excludes anger. We shall say that Happy is *cognitively incompatible* with Anger. Cognitive incompatibility is based in the fact that a *C-law* is in operation in O's system such that tokenings of Happy causally bring it about that tokenings of Anger cannot occur. We can represent this by: $\text{Happy} \Rightarrow \neg \text{token Anger}$. This law is purely causal: the incompatibility is not defined by appeal to logical incompatibility. Proto-signalling through wagging conflicts with the proto-signalling of anger because happiness excludes anger. No signalling occurs. We might say that Kimba is just being confusing or perhaps indicating play. If Kimba proto-signals anger and is directed towards suppressing any sign of play and any sign k , such that state, $RD[\Omega, k]$, where Ω is cognitively incompatible with anger—that is, $\Omega \Rightarrow \neg \text{token Anger}$ —then he signals his anger-state.

We should now note several things. Everything here is a structure specific to O—in this case Kimba. The receiver O' —in this case Simba—has to construct a simulation of O's

production. This is an exercise in S-processing. So O', Simba, searches Kimba for signs of cancellation, that is, proto-signals, and constructs a representation, in terms of his repertoire disposition, of the causal underpinnings of O's state. Simba has no access to Kimba's inner states and the relations of cognitive incompatibility for Simba's NC system. He does not need it. He simulates a structure in terms of his own NC system and its own dynamics. If it turns out that, for Simba, roaring and wagging proto-signal incompatible states then he will be in a position to rightly S-process Kimba as not signalling anger. Or if there is no wagging, and no other behaviour incompatible both by Kimba and Simba's lights, Simba will be in a position to S-process Kimba as signalling anger.

We now have a theory of signalling, which we can sum up as:

Signalling Σ : O signals Σ by producing s iff O produces s and (i) has $RD[\Sigma, s]$, (ii) is directed towards $RD[\Sigma, s]$ causing production of s , (iii) O is directed towards suppressing proto-signals of play-status or conflicting mental elements.

Proto-signalling goes on in signalling. It is a component of it, since it involves the application of a proto-signalling, and the suppression of signs that can be used to proto-signal play or conflicting states. So we can analyse the structure of signalling in terms of proto-signalling as below:

Signalling Σ : O signals s by producing s iff (i) O proto-signals Σ by s , (ii) O suppresses proto-signalling of play-states and conflicting states.

Signallings can be sincere: O' signals Ψ and really has Ψ . Or insincere: O' signals Ψ without having Ψ . If O S-processes O's production of s as a signalling, then O simulates the structure of proto-signalling, with no play, and no-cancellation, in terms of states of her system. That will mean a state of the form $RD[\Sigma, s]$. But what does O' do if he takes O' to be sincere, that is, as possessing a state analogous to Σ behind production of s ? O has no access to the functional state Ψ behind O's production of s . Here again we have simulation. Just as O simulates the structure of O's proto-signalling and signalling, O simulates the structure of sincere signalling, which means, with our mirror neurone analogy, the right part of O's NC

system lights up. Having S-processed O' 's production as sincere, then O enters the next processing phase: what to do with this simulation state? It depends on what Σ is and the sophistication of the cognitive systems involved. If it is a state based in perception, O may come to acquire Σ . We could say then that O acquired perceptual information from O' . A more sophisticated system may require further processing procedures. If Σ is a motivational state, then that may involve taking it on as a desire and acting upon it, and so on.

There are three levels of interpretation that are grounded in three levels of S-processing. The first, which we might call *locutionary interpretation*, is S-processing in relation to proto-signalling. The second, which we might call *illocutionary interpretation*, is S-processing of behaviour as a signalling. The third, which we might call *perlocutionary interpretation*, is the processing in which O processes the state as sincere and then processes the resulting state in some way.

Communication and Interpretation

A communicative system is a necessary condition for the presence of communication. But it fails in being sufficient. Consider the judgement that O communicates successfully with O' . In terms of the systems we have described, a structure like the following must be present:

O proto-signals Σ by s and O' S-processes s with $RD[\Psi, s]$

But have O and O' successfully communicated? One might think it is a matter of the functional type identity of Σ and Ψ . That is not the case. GE offers no ultimate functional reduction of content. The judgement of communicative success is an interpretative judgement. All interpretative judgements are themselves based on the S-processing signalling activity of agents. In short, the judgement that O has successfully communicated with O' will itself partially involve a disposition to S-process the productions of O and O' . I say more about such judgements below in §15 and in more detail in chapter 10.

Let me emphasise: GE offers structural-functional analysis. But here is a crucial idea. We are characterising signals, s , in terms of inner states, Σ . But the significance of the sign s is not Σ . Contents are not inner states. Nor are they possessed by inner states in virtue of

functional features. We give no theory of content, that is, an account of how contents are fixed by functional states.

§14 Signalling Systems in Humans: Assertion and Proto-Assertion

Such are the general structures of what I have called signalling systems that are the functional architecture of communicative systems. We now move onto language-using creatures like humans. Verbal and pre-verbal humans are equipped, in varying degrees, with a system of pre-verbal communication, featuring hand gestures, facial gestures, bodily orientations, voice tones, etc. Pointing to features of the physical environment is a way of ultimately signalling the kind of state one is in—a certain perceptual state. Perceptual states are relational. They link an inner state to an outward thing perceived. Thus, pointing at a perceived external thing is a way of proto-signalling a perceptual, and thus relational, state. But there might be some slackness in regard to which feature of the world the speaker points to. A significant refinement to a signalling system is gained by the introduction of words. Words allow for more fine-tuned signalling of states. Thus, if a U says *dog* U proto-signals an inner state. We are not proposing that U refers to her state with *dog*. U's term *dog* refers to dogs. Reference attribution is a matter of interpretation, and at the moment we are not in an interpretative orientation to speakers.³¹

Generally, what we call speech-acts are productions of symbols signalling states. Our main concern is the signalling act that is assertion. We proposed above that assertion is the act in which speakers signal dispositions to signal grounding relations—§12. As noted there, Π -properties are linked by grounding relations, that is, by C-laws. We may introduce an abbreviation. Let Π^{GI} be the state of having a ground-signalling or indicating disposition in relation to Π . This is a state of being disposed, given prompts, to signal grounds for a Π -property. So, assertions are acts signalling states of the form Π^{GI} . We shall apply what we have learnt about signalling systems to that special form of signalling that is assertion.

³¹ Thus the classical Lockean view would propose that what is signalled or indicated is a contentful state. It would then offer a theory of content in terms of such states, taking these to be truth-apt and propositional in some cases. This is not the path of GE, but for a recent approach along these lines see Davis (2005).

Assertion and Proto-Assertion

As I proposed above, signalings depend on repertoire dispositions, and what I called *proto-signalling*. In the case of assertion, qua signalling act, we have the following kinds of repertoire disposition. We designate such disposition by $RD[\Pi^{GI}, S]$. That is:

$RD[\Pi^{GI}, S]$: If U is directed towards producing an utterance caused by Π^{GI} ,
then U may produce S.

Again, the *may* here is not the *may* of normative permissibility. It is the *may* of epistemic possibility. Amongst the things the speaker may do is produce S, she may also produce R, if there is more than one way of defending the state. It could be *That's good*, *That's cool*, etc, all of which can be used to signal dispositions to signal grounds for the same basic preference state. *That's cool* can be used to signal entirely different grounding-indicating states, i.e., when she asserts that something has a low temperature.

Dispositions of the kind above are grounded in more basic dispositions, which are to do with sub-sentential constituents, like predicates. In later chapters I go into some detail about the mental elements that underpin the use of predicates and referring terms. For now, let Γ be the mental element underpinning U's use of a predicate, say *green*. Let Φ be the mental element underpinning U's referring with a term *N*. Let + be the concatenation operation which is the alignment of mental elements to form a Π -property—I say more about + in §16 below. Let the Π -property for *N is green* be constructed from Φ and Γ , as in $\Phi + \Gamma$. In which case, U can have the disposition:

If U is directed towards producing a sentence caused by a ground-indicating disposition for $\Phi + \Gamma$, where Φ is the mental element for *N*, and Γ for *green*, then U may produce a sentence of the form *N is green*.

Dispositions like these, delineating U's repertoire of sentence- and word-use, are a basic aspect of U's language agency capacities, the basis of what we call *linguistic competence*.

In line with our proposals for signalling systems, we can now analyse certain kinds of proto-signallings based on such repertoire dispositions. These acts will be rightly called *proto-assertions*. In line with *Proto-signalling* in §13 above we have:

U performs a proto-assertion with S iff U utters S where U's act is caused by $RD[\Pi^{GI}, S]$ and her directedness towards $RD[\Pi^{GI}, S]$ being a cause of her production of S .

In §13, we analysed signalling in terms of structures underpinning proto-signalling and suppression of proto-signals of play-status and conflicting states. Following the pattern set there, we can analyse assertion as:

Asserting with S = U utters S and (i) U has $RD[\Pi^{GI}, S]$, (ii) U is directed towards $RD[\Pi^{GI}, S]$ being a cause of the production of S , and (iii) U is directed towards suppressing proto-signals of play-status and conflicting states.

Or, more succinctly in terms of proto-acts:

Asserting: U makes an assertion with S iff U utters S and (i) U proto-asserts with mental element $RD[\Pi^{GI}, S]$ and (ii) U is directed towards suppressing proto-signals of play-status and conflicting states.

Assertion is signalling a state Π^{GI} , that is, a ground-indicating disposition. I shall often denote this complex act by the phrase defending a Π -property—we used this term in §1-2. Assertion then is production of a sentence defending a Π -property.

These are structural analyses of assertion. We are not offering an analysis of what U does when she asserts that P, for some content P, since that brings in a content attribution. Rather we are describing the functional structure of certain kinds of acts: assertion. Assertion is a speech-act structure, with a certain functional character, upon which content can be draped. We have yet to get to the question of content.

Comments:

(a) Assertions are not reflexes. They have, as already noted, an element of directedness in them. Even where U possesses $RD[\Pi^{GI}, S]$ and a corresponding ground-indicating disposition,

her dispositional state is not necessarily manifested. It depends on U's being directed towards producing utterances. That is the non-reflex element.

(b) **Asserting** is not defined in such a way that requires that U have the state Π or even the state Π^{GI} . For example, U can make a public assertion and not have any disposition to defend an element Π , since she is just lying mindlessly. But to sustain deception, a speaker would frequently have to develop a ground-indicating disposition. Sincerity requires that all the signalling going on is sincere, which is to say, U has the states lying behind her signals. So, sincere assertion requires a sincere ground-indicating disposition.

(c) Can there not be speakers who robustly affirm claims but appear to have no reasons at all for the states that may be associated with the sentences they utter? There can. But in these cases we can distinguish between bad reasons and no reasons at all. Dogmatic assertions occur where U defends Π with Π -itself. Dogmatic assertions are not all necessarily objectionable. U may assert on the basis of intuition. As a trained-chicken sexer, I can tell by a look that a chick is male. But I cannot articulate my grounds beyond saying: *it looks male*. We allow this kind of dogmatic assertion because the speaker may in fact be a reliable detector of certain properties. Assertion of the validity of basic logical principles has a similar status.

(d) One might object. Is it not possible to sincerely assert and lack a sincere ground-indicating disposition? Say that U is a doorman who has been told by his employees not to let a certain person H into a club. U says to H *You cannot enter*. H responds *why?*, but U won't respond. U has made an assertion but lacks a ground-signalling disposition. But U does not lack totally a ground-signalling disposition. For example, U has the private state. In response to his own inner prompt he would provide grounds.

(e) Asserters may be totally indifferent in making public assertions in relation to the mental states of others. U can assert before audiences she believes to be indifferent or disbelieving, or U herself can be totally indifferent whether she is believed or comprehended. Assertion just requires that U produce an utterance signalling a state; that is, an utterance in

which U activates repertoire dispositions, and suppresses proto-signalling of play-status or conflicting ground-indicating dispositions.³²

Kinds of Processing

If U says something to H, H must process U's utterance, and there are potentially three levels to H's S-processing of U's utterance: *locutionary*, *illocutionary*, and *perlocutionary*. Let us call this locutionary processing *L-processing* for short. In L-processing, H assigns a repertoire disposition $RD[\Pi^{GI}, S]$ to *S*. That means H's NC system simulates activation of $RD[\Pi^{GI}, S]$, where activation is that state in which it is liable to be a causal precursor to utterance. We can say for short that H assigns a Π -property to *S*, or that H L-processes *S* in terms of $\Pi-[S]$, where $\Pi-[S]$ is the Π -property H assigns to *S*. There is more to S-processing than L-processing. There are also the illocutionary and perlocutionary dimensions of S-processing. H's illocutionary processing of U's utterance can assign assertive status to it. That means, as we saw in §13, that H simulates the structure of U's proto-assertion and directedness to cancelling states—play or conflicting states. Perlocutionary processing occurs where, having taken U to be sincere, H responds to U's assertion. H may already have a disposition to defend $\Pi-[S]$, the Π -property H assigns to *S*, or she may not. Then H has to decide whether to take on the disposition. Her decision is an epistemic matter, and we will not investigate that issue here.

§15 Non-Assertions, Belief, Judgement, and Embedding

Any adequate theory of assertion needs to distinguish assertion from other kinds of illocutionary acts. We now compare assertion with orders. We also need to see assertion in relation to belief, since, the claim was made that assertion, belief, and judgement are at the same cognitive level. Finally, we should address another form of production of declarative sentences: their utterance in sentential compounds.

³² Rhetorical forms are interesting interpretative possibilities. In irony, a speaker can proto-assert with *George is a real genius* and indeed make an assertion. But in S-processing her utterance, U is, for interpretative reasons, unable to treat U as sincere. So what is U doing? U is interpreted as play-acting. She is pretending to make a sincere assertion.

Orders

Orders are productions of sentences but are not signalings of ground-indicating dispositions. Rather they are simply signalings of desires. U may have those desires as the result of grounding relations. But, in an order, the disposition to indicate grounds is not causally active. The purpose is not to signal a ground-indicating disposition. A successful communication in which H understands U's order and acts on it will be one in which H S-processes U's utterance in terms of a repertoire disposition: $RD[\Phi, F!]$, where Φ is some desire state in H's system. H takes on the desire, and acts on it. We cannot define necessary and sufficient conditions for understanding an order on the functional level, since that would require a reduction of content to the functional, which is eschewed in GE. Rather what we can capture are certain necessary structural conditions.

Belief and Judgement

The present theory of assertion seems to be an improvement on those canvassed in §9-10. The theory offered does not analyse assertion in terms of some underlying state of belief. But what then is belief for GE? Clearly, in sincerely asserting a sentence S , U does have a belief corresponding to the assertion. Belief can exist without assertion, but *sincere* assertion cannot exist without belief. So, one might think that something is missing in our analysis: reference to belief. But this is not so. All that these considerations show is that assertion and belief are connected. Here's the connection as outlined in §11. Sincere assertions are manifestations of belief. Belief is a disposition to sincerely assert, either privately or publicly. The disposition may remain unmanifested. We can put it another way: beliefs are dispositions to defend Π -properties actually possessed by the agent. Sincere assertions are public manifestations of such dispositions, whereas, judgements are private manifestations.

The idea that beliefs are dispositions to assert has always been undermined by failure to produce an analysis of assertion that is independent of belief. But now we have an analysis of assertion independent of belief. There is also the worry that non-linguistic creatures cannot have beliefs. This is true to this extent: belief requires ability to think, and thinking requires consciousness of symbol use, and non-verbal animals lack that. Whilst higher animals do not

think in the sense of inner monologue, they have all the hardware of an NC system with pre-doxastic elements in relations of grounding. They can reason in the sense that their NC systems produce outputs. Here they do exactly what we do.

Beliefs differ from desires in that beliefs are truth-apt and not desires. Or in other words, beliefs have one direction of fit, whereas desires have another. That is ultimately based in this fact: beliefs are dispositions to defend Π -properties, whereas desires are not dispositions to defend Π -properties.

Beliefs are often described as states that take objects, the objects being known as *propositions*. The paradigmatic way of referring to propositions is through *that*-clauses. To understand propositions we need to first understand how we refer to them. That is not to suggest that propositions are mere linguistic constructions. It is rather the counsel that any question about what Fs are must be sensitive to the issue of how we talk about Fs. As we shall see in chapters 6 and 7, *that*-clauses are *sentential nominals*. In using a *that*-clause one does not refer to an object. Rather one refers to a *non-object*—see §3—something that lacks a nature. There is then no theory about what propositions are, there is only a theory about the language agency of *that*-clauses.

Embedding and Play

The final structurally significant cognitive category that we need to consider is that of declarative sentences produced in embedded contexts. As we saw in our general account of signalling, the basic act performed in a signalling system is proto-signalling. Signallings with sentences are illocutionary acts. A sub-class of these signallings are assertions. An assertion always has at its core a proto-assertion. Proto-assertions are productions of a sentence S caused by a repertoire disposition $RD[\Pi^{GI}, S]$ and directedness to the causal efficacy of $RD[\Pi^{GI}, S]$ —see §14 above. Assertions are proto-assertions in which the speaker is directed towards suppressing any cancelling proto-signals in relation to the proto-signalling of Π^{GI} , that is, proto-signals of play-status or conflicting Π -properties. Declarative sentences embedded in logical compounds are, in most cases, proto-asserted sentences where U proto-signals that a context of play is present.

Sentence-frames such as *Not...*, *either ...or...*, *if...*, ... signal that a context of play is present. The Π -property of a logical compound $O(...S...)$ is a complex Π -property, one that has Π -properties as its constituents. The Π -properties that are its constituents are those of the sentences S of the compound $O(...S...)$. Let Σ be the Π -property of S . So a negation *Not-S* has as its Π -property an operation on Σ , a disjunction *either S or R* a Π -property on Σ and the Π -property of R , and so on. Let us denote the complex Π -property of any compound $O(...S...)$ by $\{\dots\Sigma\dots\}$. In using a logical operator O in a logical compound $O(...S...)$, the speaker proto-signals a ground-indicating disposition of the form $\{\dots\Sigma\dots\}^{GI}$. In other words, the repertoire disposition underpinning use of a compound has the form: $RD[\{\dots\Sigma\dots\}^{GI}, O(...S...)]$. It is a disposition with this general character:

Where U is directed towards signalling $\{\dots\Sigma\dots\}^{GI}$, she can utter $O(...S...)$.

The whole production of a compound $O(...S...)$ is a complex co-ordinated effort. It involves the production of embedded sentences, S , and the production of the whole. Each embedded sentence S is an act of proto-signalling. That means it is an utterance of S caused by a repertoire disposition $RD[Z^{GI}, S]$. Production of the whole sentence $O(...S...)$ is caused by its repertoire disposition above.

U asserts $O(...S...)$ in virtue of the fact that U produces $O(...S...)$ caused by its repertoire disposition and is directed towards suppressing any signals cancelling proto-signallings. These include proto-signallings of play-status. In asserting $O(...S...)$, U performs proto-assertions with the embedded sentences S . But can she also assert with these tokens? That depends on what O is. If O is *not*, *if*, *either..or*, then U signals that a context of play is present. That means U also proto-signals with O that the element Σ^{GI} , that underpins S , is not causally active in production of S . Take disjunction, for example. Asserting *either S or R* is compatible with asserting S . Thus assertion of disjunction does not mean U lacks entirely a ground-indicating disposition Σ^{GI} . U may well have it. Rather, in producing S in *either S or R*, the state Σ^{GI} is simply not efficacious in the causal production of this particular utterance of S . U proto-signals that with *either..or*. Similar comments apply to *if*-sentences. The embedded sentences in such compounds are annulled as assertions, they remain proto-assertions. On the other hand in a conjunction, *S and R*, U can, and in fact will be taken as asserting S , because

there is no proto-signalling of play. Conjunction is not a play-context, whereas *if*, *either..or* and *not*, are.

Given that S is always proto-asserted, be it embedded or un-embedded, a common underlying act is present. That act is the utterance of S caused by a repertoire disposition and U 's directedness towards such causation. If so, an audience is able to bestow the same locutionary interpretation on S in all contexts. That is, she can L-process S in the same way when asserted or embedded. That is the causal basis of common-content judgements.

Logical compounds $O(...S...)$ are used to defend compound Π -properties of the form $\{\dots\Sigma\dots\}$, whose constituent Π -properties are associated with the proto-assertions performed in utterance of $O(...S...)$. We have yet to say what these compound Π -properties are. Indeed, we have yet to say what Π -properties are generally, in any detail. We provide a theory of the structure of Π -properties for logical compounds in §23 and §24, and get onto Π -properties in general in chapters 4 and 5.

§16 Truth-Aptness and Basing

What have we done so far? GE is a theory of the functional structures of the language agent. We have provided an analysis of the basic agency behind assertion, orders, belief, judgement, and the distinction between belief and desire. Herein we have provided a theory of some basic cognitive and pragmatic categories in functional terms. Our discussion of semantic and cognitive categories in §8-10 was preoccupied with the semantic category of truth-apt sentences. Truth-apt sentences are sentences for which there is no category error in assigning truth or falsity. That does not mean that they are bivalent. A truth-apt sentence may be defective, say, due to reference failure, but it is still a truth-apt sentence, whereas a sentence used in an order is not. We were, given the usual tools for analysis in orthodox Semantics, unable to fashion an adequate theory of truth-apt sentence, nor were we inclined to treat it as primitive. We are now in a position to rectify matters. In what follows I embark on an explication of truth-aptness in GE's terms. What we shall provide is not a reductive analysis of being truth-apt. Rather, it is a functional analysis of talk about truth-aptness. We begin

with the basic instances of that talk, and then introduce some machinery enabling us to extend the analysis beyond these basic cases.

S-processing is the functional underpinning of parsing or understanding a sentence. Judgement that a comprehensible sentence *S* is truth-apt is based in dispositions to S-process *S* in a certain way. Basically, it is that we L-process *S* as a proto-assertion and do not assign any further non-assertoric interpretation to *S*. We do not S-process it as being used as an order or question. It cannot be generally, however, that in sincerely, clear-headedly asserting *N is truth-apt* U defends such a disposition to S-process the sentence referred to by *N*. Consider the super computer case from §8. Super proves $\$$ and we duly believe that $\$$ is true. But $\$$ is so complicated it takes several hours to read out, and we cannot even parse it, let alone understand it. But we judge $\$$ to be truth-apt. This judgment cannot require that we are disposed to S-process $\$$. We cannot. Nor are we betting on anyone being able to do that. Does this mean that we need after all to postulate some primitive semantic category: truth-apt sentences? Fortunately we do not.

What we are about to do now, to solve this puzzle, is introduce some machinery that will appear repeatedly in what follows. I proposed in the introduction and in the last chapter—§12—that Π -properties are *based* in various kinds of mental state and activity. The mental states on which they can be based are affective states, perceptual states, metalinguistic dispositions, and processing dispositions of various kinds including S- and L-processing. We shall not go into what these base-states generally are here—I do so in chapters 4 and 5. Our concern now is with *basing*, and in particular Π -properties based in dispositions to S- and L-process. Basing is a specific kind of relation between a more familiar disposition, such as a disposition to L-process a sentence in some way, and a purely abstract functional state. To explain it we need to delve into the structure of Π -properties for subject-predicate sentences and the form of predicates. We don't introduce basing merely to understand how we can talk of truth-aptness. Basing, rather, is a pervasive feature of Π -properties, and solves a raft of problems. We might say a key thread in basing is that it allows a mind, in some sense, to transcend its own cognitive limitations. It is important then to understand basing. To do this, we need to delve further into the architecture of the NC system.

Π-Property Structure

A logically simple sentence *N is F* contains a referring term *N* and a general term *F*. Just as *U*'s assertions of *N is F* are correlated with a Π -property Ψ , utterances of *N* and *F*, in such assertions, are themselves correlated with prior states. These are elements of the NC system, which are constituents of the Π -property Ψ . Let Γ -[*F*] be the state correlated with the term *F*, and Φ -[*N*] be the cognitive precursor for the referring term *N*. Let + denote the concatenation of the two elements. In which case we have:

$$\Pi\text{-}[N \text{ is } F] = \Phi\text{-}[N] + \Gamma\text{-}[F]$$

This means that the Π -property correlated by *U* with *N is F* is made up of the mental element correlated with *N*, and that correlated with *F* and conjoined by the operation +. (Note that all our talk of correlation of Π -properties is individual specific.) We shall describe the nature of these mental elements, Φ -properties and Γ -properties, as we proceed. Don't complain that this is complex. The language agent, the signalling system with its companion NC-system, *is* complex. But I am not supposing that speakers who embody these systems need to have any idea of these structures. They don't. They are the causal structures they embody.

Concatenation + is that relation that allows a Φ -state and a Γ -state to function together in a Π -property.³³ In order for elements, Φ -states and Γ -states, to play a role, they have to be conjoined. The exact nature of this relation is a matter for a full-fledged cognitive science to determine. But we can say something in broad terms here. In Fodor's (1975) mentalese hypothesis, a modular view of mind is embraced. In that theory one talks of a belief-box, which is to say, the location or sub-system in the functional architecture of the mind-brain in which mentalese sentences are located so as to perform a certain role, the belief role. One might also have a fear box, and other boxes for other propositional attitudes. In the spirit of such an approach, we propose that Π -properties have an acceptance box, that is, a functional place enabling them to play certain roles. If a Π -property is part of a sincere

³³ In what follows I treat state and property as synonymous when talking of the NC system. Thus *Φ-property*, *Π-property*, and *Γ-property* are interchangeable with *Φ-state*, *Π-state*, and *Γ-state*. Furthermore, [...]-brackets or (..)-brackets are used interchangeably.

ground-indicating disposition, that Π -property is part of a belief. If it is involved with a kind of disposition, linked to action, it can be desire or fear. In each case, to enter into these roles, the two elements Φ -[N] must be conjoined Γ -[F]. This is where we find concatenation +.

Basing and Truth-aptness

Let us then return to the problem of truth-aptness for comprehension transcendent sentences. The key lies, I suggested, in *basing*. A statement of truth-aptness, paradigmatically, involves utterance of a sentence of the form $\$$ is *truth-apt*. Taking surface grammar seriously, and using our new notation, we can represent the Π -property defended in this assertion as:

$$\Phi$$
-[$\$$] + Γ -[*truth-apt*].

Φ -[$\$$] is the mental element underpinning use of the referring term. Let's not worry about that now. Our concern is with the mental element of the predicate *truth-apt*. I have so far denoted this element by Γ -[*truth-apt*]. But now I am going to introduce another way of denoting it. I shall represent this state as $D_{\text{Truth-Apt}}$. The change of symbolism is just to give idea that we are denoting a kind of functional state—that is, a kind of dispositional state—and that its nature is not intrinsically tied to the word *truth-apt*. If so, we can represent the Π -property for $\$$ is *truth-apt* using this variant notion as:

$$\Phi$$
-[$\$$] + $D_{\text{Truth-Apt}}$.

Our task now is to describe the nature of $D_{\text{Truth-apt}}$. Its nature is not to represent—representation has no place in GE's description of language agency. The term *represent* can be deployed only when we take up an interpretative stance—see §1-2—but in describing the language agent we do not have this stance. Talk of representation then can play no role. Nor is the character of $D_{\text{Truth-apt}}$ defined by inference, since inference is at the level of assertion and symbol production, and we are at the pre-symbol level. Yet, there is something faintly reminiscent of inference in the functional role of $D_{\text{Truth-Apt}}$. The property $D_{\text{Truth-apt}}$ is functionally defined in terms of its relations within the NC system. These relations are not inference or even the precursors of inference, but they are related to it, as we shall see.

To understand the functional role of $D_{\text{Truth-apt}}$, let us begin with U asserting the truth-aptness of a sentence S she can understand. In this case U is disposed to L-process S in terms of some proto-assertion structure that has no further non-assertoric interpretation. Let us abbreviate *disposed to L-process S as a proto-assertion without further non-assertoric interpretation* by the ghastly $LPA[S]$.

Now for a very important idea: the state $LPA[S]$ is the *primary ground* for a truth-aptness statement. Whenever U verifies by her own lights that a sentence S is truth-apt, it is because she is in a state $LPA[S]$. What has this got to do with our functional state $D_{\text{Truth-apt}}$? It means that $LPA[S]$ is functionally sufficient for a state that is the concatenation of the state $D_{\text{Truth-apt}}$ with the mental element $\Phi-[N]$, where N is used by U to denote S . That is, the following cognitive law, or C-law holds:

$$LPA[S] \Rightarrow \Phi-[N] + D_{\text{Truth-apt}}$$

This records the fact that when the language agent is in the state $LPA[S]$ that agent can generate the state $\Phi-[N] + D_{\text{Truth-apt}}$, given that it uses N to refer to S . The C-law is just a regularity in U's NC-system: one state can be derived as a causal result of another. We might call this the primary functional role of $D_{\text{Truth-apt}}$. It is that kind of state which can combine with $+$ and a mental element $\Phi-[N]$ as the causal result of a state of the form $LPA[S]$. We might say it is a causal trace of the state $LPA[S]$, but structured in a certain way. The other elements in the structure, $\Phi-[N]$ and $+$ will also have their causal roles

We have described the basic functional role of $D_{\text{Truth-apt}}$, but we may now wonder what its intrinsic character is. The answer is that it does not have one, or that whatever its intrinsic character is, it is irrelevant to its function as an element of the NC system. Its nature as an element of the NC system is purely relational. Indeed, this is a general characteristic of elements of the NC-system. This echoes ideas about the holism of meaning. It only echoes that idea, since we are not talking about meaning. GE has no holistic thesis about meaning—it has no theory of meaning. The holism in GE is that the NC system is a network of elements with functional natures; that is, natures defined by dispositional relations to other elements in the system and outside, in the larger world.

To repeat then, $D_{\text{Truth-apt}}$ is that state whose primary ground is given by the C-law above.³⁴ The state $\text{LPA}[S]$, however, is not the only mechanism through which a Π -property of the form $\Phi-[N] + D_{\text{Truth-apt}}$ can come to be instantiated in U. Let us return to the problem of truth-aptness statements for comprehension transcendent sentences. The sentence S is comprehension transcendent, and so U cannot instantiate a state $\text{LPA}[S]$, and consequently, cannot instantiate the primary ground for $\Phi-[S] + D_{\text{Truth-apt}}$. Nevertheless, there is no cognitive barrier to her instantiating the state $\Phi-[S] + D_{\text{Truth-apt}}$. It is not too complicated for her NC system. The question is, what grounds this state if not $\text{LPA}[S]$? The cognitive explanation of her instantiation of this state is not that she instantiates a primary ground, but that she instantiates other kinds of states. U reasons to S is *truth-apt* from her two beliefs:

- (i) Where Super provides positive output for sentences x , then x is truth-apt.
- (ii) Super has produced positive output for S .

The Π -properties of (i) and (ii) generate the Π -property for N is *truth-apt*, that is, $\Phi-[S] + D_{\text{Truth-apt}}$. The Π -property of sentence (i) is, effectively, a disposition to instantiate Π -properties of the form $\Phi-[x] + D_{\text{Truth-apt}}$, given instantiation of Π -properties of the form Π -[*Super provides positive output for sentences x*].³⁵ The Π -property for (ii) provides a stimulus condition for this disposition, which is manifested in the Π -property $\Phi-[S] + D_{\text{Truth-apt}}$, the Π -property of S is *truth-apt*. That's how U does it. No primary grounds are required.

We can now characterise what we might call the canonical Π -property for a truth-aptness statement:

Truth-aptness	Π -Property Specification
N is <i>truth-apt</i>	$\Phi-(N) + D_{\text{Truth-apt}}$ where $D_{\text{Truth-apt}}$ is based in $\text{LPA}[N]$.

This format shall be repeated in later chapters. The aim is to provide Π -property specification, for an agent, of productions of assertions and proto-assertions by that agent. These specifications describe functional features of these acts. Basing is fundamental to the

³⁴ We might say that these correspond to language entry rules in the sense of Sellars (1956) and developed by Brandom (1994). There is an analogy. The difference is that these are not rules and not inferences. I say more about the relation between basing and language entry rules as we go along.

³⁵ I say more about universal statements in §54 below.

workings of the NC system, and to understanding the functional role of predicates and relational terms. It is only now that we can appreciate the abstract nature of Π -properties.

Note that we are providing a speculative cognitive theory of the functional states underpinning U's ability to use the term *truth-apt*, and U's application of it to both comprehensible and comprehension transcendent sentences. We are supplying a theory of language agency for the term *truth-apt*. We are not providing a semantics in the sense proposed by Semantics—see §1. We are not providing informative truth-conditions for sentences of the form *N is truth-apt*, nor are we providing a theory of truth-aptness, or an analysis of the concept. Nor is this a theory of assertability conditions. Assertability is a normative concept, and GE does not talk of norms of assertion. We are providing a theory of the relevant cognitive structures.

That completes our explication of truth-aptness for sentences. What we have said applies similarly to related categories, such as, *grammatical*, *meaningful*, *says something*, and so on. We shall look at some of these in more detail in chapter 10. As we shall see, this orientation to semantic categories, along with the *non-object* treatment of propositions given in chapter 7, enables us to avoid the pitfalls of Platonism about meaning. But these results are a long way off yet.

§17 *Content-Attribution, Truth-Attribution, and S-Processing*

What we have offered so far is a structural theory of assertion, orderings, implicature, the production of truth-apt sentences in embedded contexts, and finally a structural account of truth-apt talk. GE is not a theory of content, but rather a structural theory—a theory of the functional structure of activities and states of the language agent. We have provided a theory of what goes on in assertion. We have not provided a theory of what constitutes the fact that U asserts that *P* for some proposition that *P*. GE does not offer a theory of content. But what of content? How does this get into the picture? The answer is that it doesn't. Instead of providing a theory of content, GE gives a structural analysis of sentences in which content is assigned to words. These are sentences like *H's sentence means that R*, *The term T means N*,

H refers to *O*, etc. Assertions of content-attribution sentences are acts in which speakers defend Π -properties based in dispositions to L-process utterances. L-processing is a species of S-processing. The latter is a kind of simulation and recapitulation of the production of utterances. L-processing is just the locutionary aspect of that—see §13-14. S-processing is the essence of what we might call *the interpretative stance*. Content-attribution is essentially linked to the interpretative stance. No specification of dispositional, causal, and nomological features underpinning use of symbols can ever function as a content-attribution to those symbols. No content-attribution can be a specification of functional character because all content-attributions are poised within the interpretative stance in which speakers defend Π -properties based in dispositions to L-process, or more generally S-process.³⁶

At this point, let us restrict ourselves to attributions of meaning to comprehended sentences. (In chapter 10, §73, I look at the attribution of meaning to both comprehended and uncomprehended sentences. That account requires the apparatus of *basing* just illustrated in §16.) A meaning-attribution for a comprehended sentence has the form:

In asserting *H means R by 'S'*, U defends a disposition to L-process H's utterance using the state $RD[\Psi^{GI}, R]$.

This is the functional structure of a content-attribution statement—in this case propositional content. It is a description of what U does. U will typically have no grasp of the intrinsic character of the state Ψ^{GI} or the structure of her assertion. No such knowledge is required for U to make such assertions. (Of course, as theoreticians we have a speculative theory about the structure of such states.) A statement that both U and H mean the same by their sentences *S* and *R*, for example, an assertion of the form *S and R both mean that P*, has the functional structure below, which is a variation on the above:

³⁶ L-processing of known language is automatic, spontaneous, nearly always non-inferential. In reading or listening, speakers perceive words possessing meanings, in the sense that they have automatic, non-inferential acquisition of informational states caused by worldly occurrences; production of words. This does not imply that we think of meanings as objects some how attached to words, whose relation to which becomes problematic. I say more about this in chapter 10.

In asserting '*S*' and '*R*' both mean *P*, a speaker B defends a disposition to L-process both *S* and *R* using a state $RD[\Psi^{GI}, P]$.

None of these proposals analyse statements of meaning-attribution in terms of judgements about identity of the states that may lie behind various speakers' production of sentences, which is to say, the states that feature in their various repertoire dispositions. But it might be thought, we cannot avoid reference to such claims of identity, for what makes U's assertion, *H means by 'S' that R*, and the like true? Surely, it might seem, the answer to that question can only be given, indeed must be given, in terms of the type-identity of the Π -properties that U and H associate with their respective sentences, *S* and *R*.

This charge is false. What makes *H means by 'S' that R* true is the fact that H means by '*S*' that *R*. This fact is not reduced to any other fact. Such a contention can only become clear when we provide a theory of the activity of truth-talk, fact-talk, talk of truth-making, and talk of objectivity, which I do below. As I show in chapter 8, and 10, GE embraces irreducible facts of meaning. But that, as we shall see, is not a metaphysical thesis.³⁷

Truth

Judging that a sentence is truth-apt is a necessary condition for judging that it is true. But what goes on when we judge it true? GE offers no theory of truth, but a theory of truth-talk. I shall give that theory in some detail in §20. But now I give a hint of the shape of things to come. The practice of truth-judgement in the case of sentences we comprehend is given below. Here $\Pi-[S]$ is the Π -property in terms of which H L-processes the sentence *S* produced by U. (U could be H herself):

ID-True: H judges true a sentence *S* produced by U iff H sincerely defends $\Pi-[S]$, where $\Pi-[S]$ is the state in terms of which H L-processes U's sentence *S*.

³⁷ Statements of *O believes that P* carry implicit content attributions, and content is functionally irreducible. This might seem to suggest a kind of token identity thesis about belief. But that does not follow. We do not have to say that elements of the NC system have content assigned to them. Content may only be assigned to symbols manipulated consciously by agents. Beliefs are dispositions to manipulate symbols in certain cognitively rich ways. Perhaps, we can say, in the end, that beliefs are properties of people. Clearly there are large issues here, but as they are not crucial to our immediate concerns, I shall leave them aside.

I use the letters **ID** because this description of what speakers do has an *interpretative dimension*. That is because truth-evaluation of sentences involves an L-processing relation to a sentence. L-processing (and S-processing) is the backbone of interpretation. **ID-True** raises the issue of whether H understands or interprets correctly U's utterance. As noted above there is no functional reduction of correct interpretation. However, for many cases correct interpretation requires a functional-identity up to a certain point. The mental elements underpinning U's and H's use of *S* have a common skeletal structure. Note also that judgement of truth requires a disposition to *sincerely* defend a Π -property, which is to say, one really has the Π -property Π and the ground-indicating disposition Π^{GI} .

ID-True reveals that the basic use of *true* in which we assert *N is true* where *N* refers to a comprehended sentence *S*, is just an assertion in which we defend Π -[*S*]. **ID-True** is not an analysis of *true*. It is a description of a central instance of language agency with the truth-predicate. We need more to explain truth-attributions to comprehension transcendent sentences like $\$$, as produced by Super. We deal with these cases in §20 by application of *basing*, which was illustrated in the case of truth-apt in §16.

ID-True reveals that there can be an activity of truth-ascription that does not require Π -properties to be truth-apt.³⁸ It is not required that others peer into the states of agents to determine the Π -properties they defend. Speakers, rather, encounter things at the phenomenal level, that is, through the perspective of their S-processing dispositions. That is why we say it has an interpretative dimension, in that evaluation is grounded in S-processing (and more particularly in L-processing). See again §1-2 for general remarks on interpretation.

§18 Wissenschaft und Verstehen, and The Home Language

We have provided a theory of truth-aptness, assertion, and the beginnings of a theory of content-attribution from a functionalist perspective, and a theory of the truth-predicate. These

³⁸ GE's approach to truth has some precursors in Strawson (1950a), and Price's functional theory (1988), though the framework described is radically different.

are the foundational elements of GE, that is, a functionalist theory of language agency that provides, as a sub-component, a theory of meaning-talk and the interpretative stance.

The whole orientation of GE is towards description of the individual. We do not describe language communities or forms of life. We describe the signalling system of particular speakers, that is, the underlying NC system and repertoires of production and reception of symbols. We, as cognitive scientists, describe signalling systems (language agents), and do so by deploying our own signalling systems to describe the functional structure of agents' production, processing, and responses to sign productions. That includes S-processing, which is the spontaneous reception of signals through simulation; S-processing being the functional basis of understanding. We also describe the structure of that kind of symbol production that we can identify as interpretation or content-attribution—whose core is defending dispositions to S-process (and L-process) in a certain way. These content attributions may take the form of reference attributions, propositional-content attributions, and truth-condition assignments.

In taking up this functional stance to agents we give up our own interpretative orientation to our subject of study—see §1-2. In its purist form, our theorising does not utilise acts of S-processing and defending such states in relation to our subjects. Our subjects may be interpreting us, but we are not interpreting them. This is the essence of the perspective of *Wissenschaft*: we renounce the natural spontaneous stance towards language that is S-processing, which is integral to being a language-agent. It is a view radically outside the space of interpretative perspectives. From this view-from-no-interpretative-stance we cannot talk of meaning, truth, reference, predicate-application, or truth-conditions, since we have renounced exercising the language-agency functions that enable one to use that vocabulary, that is, S-processing. Quick analogy: if you temporarily renounce your effective states, you cannot make value-statements. So if you renounce S-processing, you cannot make semantic statements. You cannot even say that certain agents speak a language in the sense of producing signs that have meaning.

We are schizophrenic. Suddenly our S-processing systems will kick in, and we shall *see* meaning. At that moment we are the embodied language-agent spontaneously S-

processing signalings of language agents. We have at the moment given up the view-from-no-interpretative-stance and have re-assumed some particular interpretative perspective. We are back within *Verstehen*. It is only here that we can use intentional vocabulary; that is, assert sentences like: *O assert that P, 'S' is true iff p, or N refers to Fs*, and so on. But now we should be aware that the use of this language is unproblematic. It does not need further explication over and above that which we have provided from the perspective of *Wissenschaft*. In short, no theory of meaning is required, and no theory of reference or truth. (At least this is the claim that will be argued for.)

To describe the home language we must retain the functional view-from-no-interpretative-stance of *Wissenschaft*. Most of this book is concerned with describing matters from this perspective. This is why you will not hear much talk of truth-conditions, meaning and reference and representation. That's why you won't have a theory of meaning. That's why a theory of the home language is not a theory of meaning. Furthermore, the description of language agency, in relation to sentences, predicates, referring terms, and logical constants at no stage work in terms of relations of representation or correspondence, or relations reducible to these. For this reason, in returning to the perspective of *Verstehen*, and taking up again semantic vocabulary, there is no non-banal answer to the questions: what is reference, truth, fact, etc.

Comments and Consequences

1. It is perhaps useful to reconsider why analysis of the home language restricts itself to the perspective of *Wissenschaft*, that is, functional analysis. Our principle argument of §1 made a case for functional analysis. If meaning is irreducibly interpretative, then at some stage, illumination of language, thought and reality needs to provide a theory of interpretation. But that cannot be done within the context of an interpretative stance to language. The only alternative available is that of a functionalist stance. Laying the foundation for this is what the first few chapters of this book has done.

2. Content plays no role in the functional analysis of the language agent. From within the stance of interpretation, symbols, words, and sentences, and dispositions to produce them can be assigned content. But to a limited degree we can direct our S-processing dispositions to

the non-verbal. We can make some room for the idea of non-conceptual content. (I say more about this in chapter 10.)

3. If there is no theory of meaning there is no theory of normativity and intentionality.

There is no theory of mind qua theory of mental content. It may be that GE accepts functionalism about mind. It just rejects the idea that this is a theory of mental content.

4. Just as GE provides no theory of meaning, it provides no theory of concepts. Concepts are not analysed as representational capacities or inferential capacities. Neither of these will do, as we shall see. First, representationalism is problematic since it is difficult to describe in functional terms the conditions for a mental state to represent something in the world, such as a class of entities. Inferentialism is flawed in that speakers with quite different inferential dispositions with the term *red* may both possess the concept of red. As we shall see, interpretation is holistic and sensitive to the cognitive limitations of speakers. Despite the fact that two subjects may differ in the inferential and application patterns of their use of *red*, U may, legitimately, S-process H's term *red* using the state that underpins her use. She assigns to them the same concept across diverse patterns of inference. But no theory of concepts emerges here, just a theory of concept talk and concept attribution.

5. There is no problem of psychologism in GE. Meaning, norm, and rightness are irreducible to cognitive and merely functional fact. GE does not reduce semantics to cognitive science. It simply gives up semantics, or as we have denoted it, *Semantics*.

6. At the heart of all talk, and signalling systems, is the intersubjective orientation. This is tied to the idea that there is no god's eye view on meaning. Meaning is perspectival. The perspective is the language agent fully engaged in S-processing, which is the functional skeleton of interpretation. The function of the truth-predicate can only be seen as a projection from this perspective. But do not get me wrong: I am not proposing that truth is perspectival or relative. GE does not provide a theory of truth. We shall see in chapter 9 that GE is entirely non-relativist about the truth-predicate.

§19 *Expressivism Resurrected*

Let us now step way back to the introduction of this book. The theory provided, GE, I claimed, can be seen to fall within the spirit of the movement in philosophy of language called *expressivism*—§3. I characterised expressivism as an explanatory thesis, summed up in *Expressivism* in §7. This is the view that we explain what goes on in the assertion of sentences in domain *D* not in terms of representational ambitions of speakers but in terms of the expression of states, whose natures are not representational. What GE offers fits this pattern. Representation has no explanatory role in analysing language function, that is, language agency. For GE the domain *D* of expressivism is all assertion. If so, because *meaning* is tied to the concept of *representation*, meaning cannot have any explanatory role in explaining language function. A theory of language function cannot be a theory of meaning.

Expressivism needs to explain the expressing that goes in assertion in *D*, and *Expressing* of §7 summed up the requirements for a theory of expressing. Expressing, we noted, can have many forms—the term in itself is almost useless. But now we have a theory of the *expressing* required by a viable expressivism. In self-standing assertions of a sentence *S*, U expresses a non-truth-apt pre-doxastic state in the following sense:

GE-Express: U utters *S* and expresses a state Π if and only if U utters *S* defending Π , where Π is a non-representational, non-truth-conditional, pre-doxastic state. In other words, U signals the state Π^{GI} , that is, a disposition to indicate grounds for Π .

Expressing in this technical sense goes on with all assertion. The expressing here is not the expressing that goes on in saying *thanks*, or conveying anger in saying *jerk* to someone. Is expressing in this technical sense worthy of the term *express* as the emotivists and others have used it? It is, since it captures what I take to be the essence of expressivism.

We very much want our expressivism not to collapse into a form of subjectivism. According to GE, in asserting any sentence *S*, U expresses a Π -property, in the sense above. That state may be based in affective states or non-affective states. U's assertion does not have

truth-conditions specifiable in terms of her or anyone else's affective states. U's assertion of *S* is not true if and only if U has Π -property. We are assured of this when we look at how the truth-predicate functions in GE. H judges true and assertion of *S* if and only if H defends Π , the state in terms of which H L-process U's utterance. To assert that some utterance is true is just, more or less, to defend the Π -property one associates with that utterance. H does not judge *S* true if and only if H thinks that the speaker U has some Π -property. H may have no idea whatsoever about Π -properties, and certainly they do not enter into her deliberations about truth-evaluation.

A theory of expressing meeting our strictures must also explain embedding. But here again there is no problem. We have provided in §14 an account of embedded declarative sentences in terms of proto-assertion. Speakers do not express states in the sense of **GE-Express** in embedded contexts. However, they perform proto-assertions. Proto-assertions are performed in both asserted and non-asserted contexts.

Truth-Conditions

GE does not claim that assertions lack truth-conditions. How do we get truth-conditions? What fixes them? The answer is that truth-conditions are cheap. To assign truth-conditions to a sentence is just to assert an instance of '*S* is true iff *P*', where '*S*' means that *P* or '*P*' is in some equivalence relation to '*S*'. Truth-conditions have no interesting explanatory role to play in the theory of language agency, and the theory of language agency is all there is to say in clarifying the relation of language to thought and a broader reality.

Everything is Expressive: Another Paradox

Another seeming paradox looms: if assertion is about defending pre-doxastic states, and belief a disposition thereto, and these states lack truth-conditions, then all assertions are expressive. But aren't some statements *reports*? We contrasted in §8 expressing a mental state, say, a desire, from reporting possession of that state. We even affirmed that making sense of that distinction was a key challenge for expressivism. But in affirming that all assertions are expressive, are we not denying there are any reports? Here we have to be careful and prepare for a slight *aufhebung* of the reporting/expressing distinction. There is a distinction between expressing a desire state, in the technical sense defined in **GE-Express**,

and reporting that one possesses a desire. It is just that the latter will be a different kind of expressive assertion. Reportive assertions are a species of expressive assertions. Roughly speaking, what we call reports are non-evaluative assertions. So reports are assertions whose mental elements are non-motivational Π -properties, ones based in perception or dispositions to manipulate Π -properties or linguistic forms—see chapter 4. In reports about desire, as in *I desire X*, the report is based in perception of desire. We can perceive our own desires and theorise/perceive others'. Π -properties can be based on such perceptions. It is Π -properties based in such states that we defend in asserting: *I desire X*.

Self-Application

Much of GE's analysis of speech-act activity is about events of symbol production correlated with dispositions and functional states. GE talks in terms of dispositions, causes, properties, and events. One might be worried. Shouldn't all this talk somehow stand outside the purview of the theory? If it did not, wouldn't the theory be incoherent? The worry is unfounded. There is no requirement that an analysis of language agency, in terms of activities and expressed states, should not apply to its own theoretical statements. To illustrate this, consider the case of dispositions. According to GE all assertions are signallings of dispositional states, including assertions about dispositions. GE's analysis of talk with sentences like *O has disposition D* will invoke dispositions. But recall that the language-agency analysis of such talk is not an analysis of truth-conditions. It proposes nothing like the following where x would make mention of dispositions: '*O has disposition D*' is true if and only if ... x Nor is it required that to signal a dispositional state one needs the concept of a disposition. There is then no truth-conditional, conceptual, or epistemic circularity. These are all the forms of circularity there are, and so GE is not viciously circular in its self-application.

Chapter 3

Truth, Reason, Logical Complexity, and Inference

§20 Truth

We have begun our functional analysis of the language agent with an analysis of the structures underpinning assertion, belief, embedded sentences, implicature, and other speech-act forms. We introduced the idea of the language agent as a signalling system with repertoire dispositions underlying its productions and related receptive systems through which it simulates states in the understanding of others' speech. The signalling system also contains an NC system, which is a network of pre-doxastic states that enter into relations of grounding. Beliefs are constituted by dispositions to signal ground-indicating dispositions. Assertions are manifestations of such dispositions. What we now need to describe before we go into more detail about Π -properties is the truth-predicate and grounding relations. We begin with the truth-predicate.

What is the speech-act structure of truth-claims? Let us first look at sentences of the form *N is true*, where *N* is used to refer to a sentence *S*—attributions to propositions are examined in §55. According to **ID-True**—see §17—that a speaker asserts '*S* is true', just in case she sincerely defends Π -[*S*], where *U* L-processes *S* in terms of Π . However, ascriptions of truth can occur where *U* is ignorant of the content of the sentence *S* which is the ultimate subject of truth-predication. So *U* is not in a position to L-process the sentence to which she assigns truth. For example, *U* might assert *Everything that George says is true*. Told that George uttered the second sentence yesterday, *U* is committed to asserting: *The second sentence is true*. *U* is ignorant of the sentence *R* produced by George, and so cannot defend any Π -property Π -[*R*]. So what is going on? Don't think this: in asserting *N is true* what *U* defends is a disposition of this kind: were *U* to encounter *R*, and thereby produce a Π -

property Π -[R], U would defend that state. Such a view is unattractive because it cannot deal with the case of ascriptions of truth to sentences that speakers are in no position to understand because the sentences are too complicated to grasp. Suppose that our super computer, Super, is doing proofs, and it produces a readout in which a very complex sentence, $\$$, which no human could possibly comprehend due to its logical and grammatical complexity. U knows that $\$$ is perfectly well formed, grammatical, truth-apt, and meaningful. U knows that if Super produces a sentence $\$$ in this way, the sentence is true. Super has produced $\$$, so U judges that $\$$ is true. The question now is what is the structure of U 's assertion of $\$$ is true? It is computationally impossible for U to L-process $\$$. We might settle on using the concept of commitment as a primitive concept to analyse truth. But that option is totally at odds with the orientation of GE, and we do not have to do it anyway. The concept of *basing*, introduced in §16, saves us, since it is already there in our functional analysis of the predicate truth-apt, which U deploys to describe $\$$. So let us see how basing works in the case of the truth-predicate.

Basing and the Truth-Predicate

What we shall say is almost a carbon copy of what we said in relation to the functional essence of the mental element, or Γ -property, that underpins use of the predicate *truth-apt*, the property $D_{\text{Truth-apt}}$ (§16). The structure of truth-ascription sentences we are looking at here is N is true, where N refers to a sentence. Let us denote the mental element behind the predicate *true* by the expression D_{True} . D_{True} is a functional property that can combine with others to form a Π -property. The Π -property correlated with N is true has the form Φ -(N) + D_{True} . The nature of D_{True} qua functional state is essentially relational. This nature has two aspects, as with every other predicate. The first is that it can combine with + and elements of the form Φ -(N) to form Π -properties. The second aspect is that states of the form Φ -(N) + D_{True} can be introduced through C-laws of a certain kind.

The C-laws we are interested in link a primary ground to the instantiation of Φ -(N) + D_{True} in U 's NC system. The primary ground is simply the disposition to defend the Π -property Ψ of a sentence S that U L-processes in terms of a repertoire disposition $RD[\Psi^{GI}, S]$. So that might be a disposition to defend the Π -property for *Snow is white*, which generates

her disposition to assert that the sentence *snow is white* is true. We have then the following C-law, where N is used by U to refer to S :

$$[\text{Disposition to defend the } \Pi\text{-property of } S] \Rightarrow \Phi\text{-}(N) + D_{\text{True}}.$$

The primary ground above is not the only ground upon which $\Phi\text{-}(N) + D_{\text{True}}$ might be tokened. Let us return to our supercomputer case. We believe that every sentence that Super outputs positively for is true. So we arrive at assertion of $\$$ is true by inference grounded in:

(i) *Everything Super positively outputs is true;*

(ii) *Super positively outputs \$.*

The Π -properties of (i) and (ii) generate the Π -property for $\$$ is true, that is, $\Phi\text{-}[\$] + D_{\text{True}}$. As already noted in §16, the Π -property of sentence like (i) is, effectively, a disposition to instantiate Π -properties of the form $\Phi\text{-}[x] + D_{\text{True}}$, given instantiation of Π -properties of the form $\Pi\text{-}[Super \text{ provides positive output for sentences } x]$. The Π -property for (ii) provides a stimulus condition for this disposition, which is manifested in the Π -property $\Phi\text{-}[\$] + D_{\text{True}}$, the Π -property of $\$$ is true.

A conversational interaction between U and H about the sentence $\$$ in which both assert that $\$$ is true is unproblematic. The Π -properties of their respective utterances have the same structural essence. This does not in itself fix the fact that they both grasp the same proposition. To understand matters of content, we need to consider the interpretative standpoint. This is the standpoint in which we, as a hypothetical speaker O , L-process the productions of U and H . It may turn out that we do L-process the utterances of U and H using the same state, a Π -property that has the structure $\Phi\text{-}[\$] + D_{\text{True}}$. It may be that the mental elements that underpin the referring terms that U , H , and O use to refer to $\$$ have a slightly different functional structure. So, the three speakers, U , H , and O do not have identical realisations of the state $\Phi\text{-}[\$] + D_{\text{True}}$. But that does not prevent O L-processing the sentences of U and H using his or her state.

The Theory

We can now formulate a theory of the truth-predicate and the falsity predicate. The Π -property specifications for N is true and N is false are:

Sentence (<i>N</i> refers to <i>S</i>)	Π -Property Specification
<i>N is true</i>	$\Phi-[N] + D_{\text{True}}$
<i>N is false</i>	Reject: $\Phi-[N] + D_{\text{True}}$.

In asserting *N is false*, U defends *rejection* of $\Phi-[N] + D_{\text{True}}$. Rejection is the mental element that underpins negation—I give an analysis of rejection below in §23.³⁹ Judgements that an utterance of *N is true* is true just have the pattern already captured in **ID-True** of §17. H judges U’s production of *N is true* true if and only if H sincerely defends the Π -property in terms of which she L-processes that utterance. This, assuming correct interpretation on H’s part, will be a structure of the form $\Phi-[N] + D_{\text{True}}$. Embedding of sentences of the form *N is true* is unproblematic, in the light of §16. In embedded contexts, U performs proto-assertions with *N is true*, which will be grounded in the repertoire dispositions underpinning her use of sentences of this type.

Substantial and Deflationary Theories of Truth

GE’s theory of the truth-predicate is a functional characterisation of the activity that underpins its use. GE does not offer a theory of what truth resides in. It does not offer conceptual analysis. GE does not reduce truth to any other property. It does not reduce truth to epistemic properties. So the following do not hold:

S is true = *S* coheres with other believed thoughts.

S is true = *S* is assertable in a state of ideal information.⁴⁰

Nor does GE analyse truth as a correspondence to reality. We can accept the following as a harmless platitude about truth: *N is true just in case the thought that S corresponds to reality*.

We shall see in chapter 8 that GE is quite consistent with universal truth-making.

³⁹ The semantic fortunes of *N is true* and *S*, where *N* refers to *S*, are the same. If *S* is true or false, *N is true* is true or false. If *S* fails to encode a proposition, *N is true* fails to encode a proposition. To illustrate the last possibility: suppose we have a sentence (N) *The man is tall*. Suppose that no determinate reference is assigned to *the man*. We should refrain from assigning truth or falsity or any evaluation to *The man is tall*. Thus no Π -property can be assigned to *The man is tall*—see §48. But likewise, no Π -property can be assigned to the sentence *N is true*. This feature is just part of the functional structure of D_{True} .

⁴⁰ Do not think that an epistemic treatment of *true* is being proposed because *true* is used to express broadly an epistemic state: defence of a Π -property. This confuses *expressing* a state with *reporting* one. All predicates are used in defending Π -properties, but they are not used in reporting these states, otherwise all statements would be claims about states of one’s NC system.

One might be tempted to think of GE's treatment of *true* as deflationary, except that *deflationary* is a contrastive term—other terms in a language get inflationary treatment—and the term naturally lives within the domain of *Semantics*. Nevertheless, a quick comparison with deflationist theories, Horwichian minimalism, prosententialism, and inferentialism, is profitable. Each of these approaches fail for various reasons.

Horwich's (1998) minimalism takes it that in asserting *N is true*, the speaker U assigns a property to a proposition. The property *t* is the property supposedly fixed by the T-sentences all of which we are disposed to assert. In the case of our computer generated sentence *S*, U asserts *S is true*, but cannot assert *S*. Furthermore, U cannot assert the corresponding T-sentence: *S is true if and only if ...*, since, she cannot comprehend the T-sentence, given she cannot comprehend *S*. If so, how does her concept of truth extend to *S*? The concept of truth is fixed by the disposition to assert a set of sentences. We cannot say U has the disposition to assert the T-sentence, even though U is willing to assign truth to the sentence—see Gupta (1993). We might propose that minimalism is the view that U has a commitment to asserting the corresponding T-sentence. But what is commitment? There are no satisfactory elucidations of commitment that can deal with the problem of comprehension transcendent truth-ascriptions, and primitivism about commitment is unattractive.⁴¹

Another minimalist approach to truth is prosententialism—Grover (1992), Brandom (1994). Prosententialism is a modern form of redundancy theory, and in its simplest form holds that assertion of *N is true* just is an assertion of *S*. That means *N is true* is a prosentence for *S*—it is just a way of re-issuing *S* as an assertion. Taken this way, prosententialism is highly implausible. U can assert *N is true* but have no grasp of the content of *S*. So her assertion of *N is true* cannot be an assertion of *S*. One might weaken the prosententialist claim. Assertion of *N is true* is a *commitment* to the proposition *P* correlated with the denotation of *N*. But that brings in an unattractive posit of commitment.

⁴¹ One idea of commitment—inspired by Brandom (1994) is this: That in committing oneself to *P* the proposition encoded by the referent of *N*, one is opening oneself to being sanctioned if not-*P*. But that cannot be the content of commitment since one cannot intend that such a conditional be true: it could easily be the case that not-*P*, but one is not sanctioned, since, perhaps, not-*P* is not verifiable for any audience. Perhaps we should take *commitment* to mean that one be sanctioned if one's interlocutors come to believe *P*. This alternative runs into the problem of comprehension transcendent truth-ascription. It is impossible that one's interlocutors come to believe the negation of the sentence to which truth is ascribed: it is beyond their comprehension.

Inferentialists shouldn't really go down this path. Presumably, they would want to suggest the following. *True*'s meaning is fixed by the introduction and elimination inference rules given below—here N refers to a sentence S :

$$TI: S \vdash N \text{ is true} \quad TE: N \text{ is true} \vdash S$$

In asserting $N \text{ is true}$, we produce a sentence whose meaning is fixed by these rules. **TI** gives us a primary ground S . But not all assertions of $N \text{ is true}$ have to be asserted on the primary ground. If so, this inferential treatment of *true* has no problem with the comprehension transcendent truth. Indeed, it is structurally like GE's treatment in terms of basing.

So how does GE improve on inferentialism? The problem with inferentialism is a general one. As argued in §10, inferentialism fails generally because it has no viable theory of semantic and cognitive categories. So although the proposed explanation is structurally pleasing, it fails in practice. We cannot stay at the phenomenal level of sentence production, but must delve below the surface to the *noumenon* of cognitive structure. That's just what GE does.

The T-schema and the Nature of Truth

GE denies that our concept of truth resides in our disposition to assert instances of the T-schema. That simply cannot be what our concept of truth resides in. GE really provides no theory of our concept of truth. What it provides is an analysis of language agency in using *true*. That agency does not require the production of T-sentences. It is independent of it. The speaker can use the truth-predicate even if she lacks any ability to produce conditional sentences. Rather, the intuitive plausibility of T-sentences, an explanation of why a speaker will be disposed to produce them, is given by a prior theory of agency in relation to *true*, and *if and only if*. Very roughly: Where U can L-process S , then she can assign Π -properties to ' S ' *is true* and S that stand in simple C-law relations, meaning, that in tokening one, she can derive the other by processing, and vice versa. In which case, she has all the materials in place to be disposed to assert the biconditional: ' S ' *is true if and only if* S .

That is all there is to say about the basic function of the truth-predicate. Of course there is more to say about truth. No theory of *true* can be complete without a treatment of the

paradoxes. I shall not offer a solution here, but one is developed using the GE framework in Barker (200+). Another topic is truth's relation to being and our talk of truth-making. That is provided in chapter 8.

§21 *Grounding and Expressivism about Reason*

Assertions are signals of ground-indicating dispositions. We have proceeded up to this point without saying much about grounding. We now make up for this omission. Grounding finds its phenomenal counterpart in reason and inference. The Π -properties defended in statements of reason are based in grounding relations. A reason statement might have the form:

R: *That P means that R.*

Not all reason statements have this form. There are entailment statements like *that P entails that R*. There are inferential sequences expressing reason relations: *P therefore R*. There are explicit statements of reason: *Believing P is a reason for believing R*. Assertions of such sentences express states that link Π -property to Π -property. For any sentence *S*, Π -[*S*] is the Π -property in terms of which U L-processes the sentence *S*. In asserting **R** and its kin speakers defend Π -properties that link, by some kind of precursor to reason, the Π -properties of *P* and *R*. Our concern now is with that link.

Consider a person reasoning and making inferences. The inferences involve the causal derivation of one state from another. The actual process of reasoning requires that there are regularities governing cognitive systems. Think of U's mind at the sub-doxastic level—what I have called their NC system—as containing a sub-system that processes Π -properties. Its processing involves extracting causal consequences from given tokened states. That is how it carries out its computations. It is not so much a syntactically driven system as a causally-structurally driven one. I have called a *cognitive law*, or *C-law*, a specific kind of functional regularity in U's cognitive system. Not all functional regularities are C-laws. Where U is willing to assert that *S*, she is also, as a matter of causal regularity, disposed to assert *I believe that S*, assuming her self-monitoring capacities are intact. Or in other words, where she is in a

state Γ , she monitors that she is in that state. But believing S is not a reason for believing that one believes that S , in the sense that it supports an inference from S to *I believe that S*. C-laws are not causal regularities based on internal monitoring. A C-law linking Π -properties Φ to Γ , write this as $\Phi \Rightarrow \Gamma$, holds just in case in a given system:

- (i) Γ is extractable as a causal consequence from Φ .
- (ii) The causation is not the result merely of internal monitoring.

In asserting a sentence like \mathbf{R} , U defends the state $\Phi \Rightarrow \Gamma$. We encapsulate the core thesis through the Π -property specification below:

Reason-Sentences	Π -Property Specification
<i>That P means that R</i> <i>That P entails that R</i> <i>Believing P is a reason to believe R</i>	$\Pi-[P] \Rightarrow \Pi-[R]$

All three kinds of sentences are correlated with Π -properties with the structure on the right.

In accordance with **ID-True**—§17—an audience H's judgements of truth have the form:

ID-Reason: H judges true an assertion of \mathbf{R} —or variation—uttered by U just in case, H sincerely defends possession of $\Pi-[P] \Rightarrow \Pi-[R]$, where $\Pi-[P]$ and $\Pi-[R]$ are the properties in terms of which H L-processes tokens P and R .

I assume here, as with all these applications of **ID-True**, that H interprets correctly U's utterance, and so gets right the basic skeletal structure of the mental elements underpinning U's utterance. That is, both U and H instantiate a structure with the form $\Pi-[P] \Rightarrow \Pi-[R]$, and with structure isomorphism for $\Pi-[P]$ and $\Pi-[R]$. But there will not be functional identity all the way down, if, for example, empirical terms are used in P and R . (I discuss this matter of functional essence and interpretation in §26.)

The Π -property specifications given, and the exhibiting of the interpretative dimension of truth-assessment given in **ID-Reason**, give us the basic features of the language agency of *entails*, *means that*, and related terms. Embedding is not a problem for familiar reasons—see §15. We should note that assertion of reason-statements often involves virtuous dogmatic assertion. In defending a C-law, U may have nothing else but the C-law itself as a

ground. So, in asserting that *P and Q entails P*, U may not be able to do much more than indicate the same ground, using different verbal forms. This does not mean that the assertion is unjustified. We allow that deep, abiding, logical intuition can be a source of justification. I provide an analysis of what that is below.

C-laws, Rational Intuition, and Errors of Reasoning

C-laws are simply causal facts about particular speakers. It is their incorporation into the agent perspective with its functionally defined interpretative orientation that allows speakers to talk of truth and rightness in relation to reason. We talk of rightness in §25 below.⁴² Speakers are not infallible reasoners. Suppose U reasons: *If P, Q. Q. So P*. Here, U defends a C-law of the form:

$$F: [\Pi-[If P, Q], \Pi-[Q]] \Rightarrow \Pi-[P].$$

This is a C-law in operation in her NC system. But the system is failing. That failure is not understood in semantically or logically normative terms, since such concepts have no explanatory role at this functional level. As we shall see the NC systems of speakers are directed towards certain features that C-laws can possess. These are causal-functional features. Just what these features are depends on the domain of reasoning. For logical and analytical reasons, these are certain causal conditions, which I specify below. The normally functioning NC system may fail to identify these conditions for various reasons, including what we call, *inattention, unreflectiveness, and cognitive limitation*. I have no theory about how these kinds of phenomena are realised in systems, but we suppose they are somehow.

Basing and C-laws

Not all statements of reason involve speakers defending C-laws. Generally speaking, the Π -properties of reason statements are *based* in C-laws. We can see why this is so by considering the case of entailment-statements about comprehension-transcendent sentences. Back to our super computer, Super, that has a perfectly reliable theorem-proving program. Its outputs are reliable indicators of the entailments between sentences. Suppose Super indicates

⁴² It is not required that speakers have any grasp of a C-law. Rather this is an account of what *they do* in asserting claims of reason. Nor is it the case that speakers need to grasp the concept of reason.

that an entailment holds between two sentences, $\$$ and \pounds , that are comprehension-transcendent. Given my faith in the computer I am committed to asserting: $\$$ entails \pounds . In asserting $\$$ entails \pounds , I cannot be defending a C-law that features Π -properties corresponding to these sentences because I couldn't possibly comprehend $\$$ and \pounds . So what do I defend in asserting $\$$ entails \pounds ? The theory of basing that we introduced in §17 and deployed in §20 comes to the rescue.

Basing means defining the functional essence of the mental element, Γ -property, that underpins the relational term *entails*. Let us label this D_{Entail} . The structure of the Π -property of a sentence N entails M , where N and M are referring terms picking out sentences is:

$$\Phi\text{-}(N) + D_{\text{Entail}} + \Phi\text{-}[M].$$

Thus the functional role of D_{Entail} is twofold. The first feature is that it can combine with the mental elements of referring terms N and M , used to refer to sentences. The second feature of D_{Entail} 's functional role is that it has a primary ground. The primary ground for $\Phi\text{-}[N] + D_{\text{Entail}} + \Phi\text{-}[M]$ is that state in which U instantiates the C-law of the form: $\Pi\text{-}(P) \Rightarrow \Pi\text{-}[R]$, where N and M refer to the two sentences P and R . We have:

$$\{\Pi\text{-}[P] \Rightarrow \Pi\text{-}[R]\} \Rightarrow \Phi\text{-}[N] + D_{\text{Entail}} + \Phi\text{-}[M]$$

The primary ground is not the only state through which $\Phi\text{-}[N] + D_{\text{Entail}} + \Phi\text{-}[M]$ can come to be instantiated. It is quite possible for U to instantiate the state $\Phi\text{-}[N] + D_{\text{Entail}} + \Phi\text{-}[M]$ through another cognitive pathway, one that involves reasoning about the general reliability of Super. We saw something analogous in the case of application of *truth-apt* and *true* in relation to the outputs of Super in §16 and §20. The details need not be presented here.

We can now characterise the canonical Π -property for an entailment statement:

Entailment	Π -Property Specification
N entails M	$\Phi\text{-}[N] + D_{\text{Entail}} + \Phi\text{-}[M]$, where D_{Entail} is based in a grounding relation \Rightarrow .

Again, we have an abstract functional essence in the mental element D_{Entail} . Think of it as the scaffold that provides an abiding foundation constraining her use of *entail*. The difference

between the Π -property specification in this case and the earlier reason-sentence specifications is that in the earlier sentences speakers use *that*-clauses. *That P entails that R* can only be asserted if U can L-process the sentences *P* and *R*, and thus instantiate the C-law $\Pi-[P] \Rightarrow \Pi-[R]$. Again, truth-assessment of entailment-statements of this form will function, as do all truth-assessments, in terms of the pattern laid down in **ID-True** in §17. Again, embedding is no issue.

§22 *Kinds of Reasons, Kinds of C-laws*

Reasons come in varieties. There are two basic kinds. Let us call them *demonstrative*—which includes logical and analytical reasons—and *inductive*. This phenomenal difference is mirrored by a difference in underlying functional structure, that is, in the functional structure of the C-laws to which NC systems are directed. That means when a speaker is reasoning logically, her system is, at some level, directed towards instantiating certain kinds of C-laws; when reasoning inductively, other kinds. U may not instantiate these states—fallacies can occur—but frequently she will instantiate C-laws that have these properties.

Logical Reasons

Logical reasons are those provided in logical arguments. In providing logical reasons, speakers are directed towards defending C-laws in which the *intrinsic causal powers* of the Π -properties of the premises are those that produce the conclusion's Π -property. The NC subsystem processor is directed towards extracting intrinsic causal consequences from premise Π -properties. Let us call these *intrinsic C-laws*. Let me illustrate the idea. Take the inference-form for one side of conjunction elimination:

$$\text{CE: } P \text{ and } Q \vdash P$$

The Π -property for *P and Q* is a compound state—designate it by $C[\Omega, \Psi]$ —that combines the Π -properties Ω and Ψ of *P* and *Q*. The state $C[\Omega, \Psi]$ is simply the union of the causal powers of Ω and Ψ . If *p* is a causal power of Ω , then it is a causal power of $C[\Omega, \Psi]$. Noting

the powers at its disposal in $C[\Omega, \Psi]$, the NC system can extract a sub-set p that corresponds to the Π -property of P , namely Ω . If so, the following is an intrinsic C-law:

$$C[\Omega, \Psi] \Rightarrow \Omega$$

This C-law corresponds to the inference **CE**. In claiming that **CE** is valid, we defend this C-law. Generally, in seeking logically valid arguments, our NC systems seek intrinsic C-laws.

What provides the stability of our inferential practice? The naturalistic stance of GE is that minds do not have intuitions about Platonic reason relations. Nor, as Inferentialism proposes, is the source of stability simply the surface inference patterns like **CE** as brute facts about our use of sentences. Rather, for GE, the source of stability is the fact that in logical reasoning speakers are directed towards intrinsic C-laws, and their NC systems are reasonably reliable at identifying these in simple cases.

Given that logical constants are underpinned by functional essences, normal functioning of the NC system will imply that a certain set of inferences are the ones that ideal speakers will be deeply disposed to realise. That implies that there is a kind of ultimate constraint on logic. This is not the ground that makes this the right logic. That is a normative matter of rightness, and rightness cannot be reduced to facts about functional states—see §25. It is rather the ultimate constraint on the system of logical intuition. It is the source of conviction that there is a right logic. I will not attempt to describe that logic here.⁴³

Analytic and Conceptual Reasons

I call *analytic reasons* ones expressed in claims like *that X is a bachelor means that X is unmarried*. By *conceptual reasons* I mean statements like *That X is red means that X is not blue*. In both cases we find the NC system directed towards intrinsic C-laws, the difference is the level of *conceptual depth*. Analytic reasons are comparatively superficial, the results of definitions and synonymy, and thus purely to do with association of surface forms with underlying mental elements.

⁴³ I do so in Barker (200+). It is not classical. Many people have affirmed the validity of the spread principle, or *ex falso quodlibet*: $P, \text{not-}P, \vdash R$. There is no way that this can be properly judged an entailment on the current account, since it fails the intrinsicality condition. That explains why there is something unintuitive about this principle—which paraconsistent logicians have emphasised.

Definition is essentially abbreviation. Take the case of defining the meaning of predicates like *bachelor*: For U, *bachelor* is defined by *unmarried man* if the Γ -state which she associates with *bachelor* is derived from that which she assigns to *unmarried man*, which is a compound form from Γ -[*unmarried*] and Γ -[*man*]. That means that judgements using *bachelor* ultimately depend on judgements using *unmarried man*. If U is doubtful about her application of the term *bachelor*, she looks to application of *unmarried man* to resolve those doubts. Or, U first came to use the term *bachelor* through her ability to use *unmarried man*. Definition is causally based on and points to a cognitive asymmetry. We are not defining definition in these causal terms, however. Rather, GE offers a theory of claims of definition: U judges that *F means G*, in the sense of definition, she is defending the state of cognitive asymmetry defined above. Compare synonymy. Definition implies synonymy. But synonymy does not imply definition. The value-terms *cool* and *good* are synonymous, but are not definitionally related. If U judges that *F is synonymous with G*, she is defending L-processing these terms in the same way.⁴⁴

Conceptual reasons are not the results of definition at all. That something is red means, necessitates that, it is not blue. The meaning of *red* is not partly defined as *not blue*. Rather we have to dig deeper. The Γ -states underpinning a speaker's use of *red* and *blue* are based in perceptual capacities: D_{Red} and D_{Blue} . Triggerings of D_{Red} exclude triggerings of D_{Blue} . This is not a cognitively accidental matter. Here we enter into speculation about the structure of the states that underpin predicates, of course. We hypothesise that in human NC systems it is functionally impossible for both D_{Red} and D_{Blue} to be triggered. Why? That depends on details of the architecture of the NC system. So what we are offering is not a reduction or analysis of the necessity conveyed when we say *Red necessitates non-blueness*. This does not mean that the necessitation involved does not exist, or is really only a pale shadow of cognitive regularities. That would involve a category error: in making such

⁴⁴ Quine (1951) tried to reduce synonymy to non-intentional notions, and failing, judged it suspect, discarding thereby the analytic. We can agree with Quine that *means the same as* is irreducible to non-intensional notions, but propose that this is because it is expressive. In asserting '*Bachelor*' means the same as '*unmarried man*', U defends a certain disposition to employ the quoted terms, which is ultimately based on their having comparable Γ -states. The expressed disposition might be to their inter-substitutivity in all non-opaque contexts.

assertions U does not report her possession of a C-law, she expresses it. Nothing is being said about necessity here at all. That's because there is nothing to be said.

Inductive Reasons

Not all reasons are analytic. There are inductive reasons. In this case, the underpinning C-laws have a completely different structure from intrinsic C-laws. When U reasons inductively, her system is not directed towards realising intrinsic C-laws, it is directed to some other kind of causal structure. In an intrinsic C-law, $\Psi \Rightarrow \Omega$, the Π -property Ω is an intrinsic causal consequence of Ψ , which means it is part of the intrinsic causal power of Ψ to generate Ω . The C-laws, $\Psi \Rightarrow \Omega$, underpinning inductive reasons lack this feature. We may judge that something is a cat and so infer that it has a tail, but it is not part of the intrinsic causal power of Π -[*O is a cat*] that we can extract Π -[*O has a tail*]. If we discover that the cat lacks a tail, we do not automatically infer that it is not after all a cat. Adding the information that O is a Manx to our pool defeats the inference from *O is a cat* to *O has a tail*. Inductive inferences are defeasible. The inference pattern from *O is a cat* to *O has a tail*, is more explicitly represented through a total evidence clause, as in:

O is a cat & No defeaters are likely. \vdash O has a tail.

The no-defeaters clause is an open matter. It might be expressed more colloquially by the sentence *There is nothing weird about the cat regarding having a tail*. What goes into the scope of this claim partly depends on the speaker's theory of cats, and, in particular, how cats can lack tails. Thus properties that may spontaneously come to U's mind as defeaters are being a Manx, having visited a vet for tail removal, genetic defects, and so on. It is because these conditions are generally unlikely for cats that the inference has the form above, and there is a presumption in favour of inferring tail-possession from cat-hood. Different speakers may differ in their cat-theory. This does not mean that they mean something different by *cat*. It is not required that one possesses such C-laws to be said to *grasp the concept of cat*.

What is the structure of the C-law in this case? One idea is that a defeasible inference from *O is F* to *O is G*, is a C-law linking Π -[*O is F*] and a state corresponding to the no-defeaters clause to Π -[*O is G*]. That might be best. Although we may not make the no-

defeaters clause explicit, it seems, we do, in such inferences, establish defence of Π -[*O is F*], then establish that no relevant defeater holds of *O* regarding *G*, and so token Π -[*O is G*]. We can represent this as: $\{\Pi$ -[*O is F*], Π -[No-defeaters] $\} \Rightarrow \Pi$ -[*O is G*].

How do we acquire such C-laws? It seems it is purely empirical. We acquire these dispositions through testimony and experience. It amounts to no more than learning that *Cats nearly always have tails*.

Generic Inductive Inference

The inductive inferences that have more complicated kinds of C-laws underpinning them, and a generic inference guiding knowledge acquisition, are ones like the following:

All *F*s observed under conditions *C* are *G*s. *T* is an *F* observed under conditions *C*. No defeaters are likely. \vdash *T* is a *G*.

This form is bound to be contentious. Either way, the speaker guided by such inferences defends the cognitive necessitation by a certain set of (premise) Π -properties of a (conclusion) Π -property.

What constrains intuitions about inductive reason of this kind? To describe what that is we need a theory of scientific rationality. I assume a naturalism thereof is a viable possibility. That means we hypothesise that we are cognitively disposed to instantiate certain C-law forms underlying inductive inference because our cognitive systems have undergone a certain selective pressure. Cognitive systems that have deployed these forms have aided the survival of the individuals possessing them. In short, they are reliable inferences. This is not to define inductive validity is reliability. We need induction to determine reliability, and so that would be circular. Rather, we are offering an explanation of why we are disposed to instantiate certain C-law forms. Our claims of inductive rightness are defences of such C-law states, not claims about their reliability. (I examine rightness in §25.)

§23 Expressivism meets Negativity

We have characterised the functional essence of reason talk for a language agent. We need at some stage to deal with the normative questions of rightness and correctness in reason. But before we get to that we need to say something about the functional essences underpinning logical constants. In what follows I consider negation and then disjunction. I consider generality in chapter 7.

What is the activity or aspect of language agency that is implemented in using *not*?⁴⁵ According to inferentialism, negation is explicated through introduction and elimination rules. The introduction rule is something like the following. Here $P \dots \perp$ is a derivation of an absurd sentence \perp :

$$NI: P \dots \perp \vdash \text{not-}P \quad NE: \text{Not-not-}P \vdash P$$

The problem for inferentialism is the status of \perp . Even if we assume that the sentence P possesses content, the element of absurdity, \perp , has no inferential explication. But that suggests that our concept is primitive. This is a somewhat worrying prospect.

According to GE, the functional essence underpinning negation is an operation on Π -properties. I call it *rejection*. If a speaker asserts *Not-S*, then normally she is deeply constrained from asserting *S*. Doubt about *S* also involves not being disposed to token Π -[*S*]. In the case of doubt, the cause of being constrained from tokening Π -[*S*] is lack of evidence, which is to say, lack of a Π -property or Π -property precursor, grounded in perception or affective states, and so on, generating the state Π -[*S*] through C-laws. Rejection of Π -[*S*] involves being disposed not to token Π -[*S*], where that state is caused by a Π -property. Let $\neg D$ -token Π -[*S*] abbreviate *constrained not to token Π -[*S*]*. To reject Π -[*S*] is to be in a state such that U tokens some Ψ and $\Psi \Rightarrow \neg D$ -token Π -[*S*]. In short, to reject Π is to have a ground Ψ for not tokening Π . The proposal is:

⁴⁵ I assume that *not* is a sentential operator in this book. For an adverbial treatment see Barker (2004).

<i>Not</i>	Π-Property Specification
<i>Not-S</i>	Reject: Π -[<i>S</i>], that is, tokening of some Ψ such that $\Psi \Rightarrow \neg D$ -token Π -[<i>S</i>].

Note that the Π -property of *not-S* is not simply the state $\neg D$ -token Π -[*S*]. It is that state of possessing some Π -property grounding the state $\neg D$ -token Π -[*S*]. This latter complex state is itself grounded by a Π -property that will impose the constraint on not tokening Π -[*S*].

Let us be careful here. One might fear that this account will not properly distinguish doubt from rejection. Suppose U does not know if the inside of a fruit is green or not green. U is in a state of doubt. U could assert: *It may be green or not green, but I cannot say which.* But doesn't U's NC system instantiate a state Ψ such that $\Psi \Rightarrow \neg D$ -token Π -[*It is not green*]? I say no. It is true that U's system, through self-monitoring, registers the fact that there is no grounding of Π -[*It is green*]. But that self-monitoring is not itself a ground for $\neg D$ -token Π -[*It is not green*]. Self-monitoring is not a source of C-laws—see §21 above.

Statements of negation are evaluated in the following terms:

ID-Not: H judges true an assertion of *Not-S* iff H, sincerely defends Reject: Π -[*S*], where U L-processes *S* in the scope of *not* in terms of Π -[*S*].

H in evaluating U's utterance of *Not-S* is not concerned with whether or not U instantiates some state, but with her own, H's, willingness to defend a state in terms of which she L-processes U's utterance. It is this inter-subjectivity, which drags U and H out of their solipsism.⁴⁶

Exclusion Laws

In asserting a negation, U defends a rejective state, but U also implicitly defends her commitments to the C-law underpinning the negation. U could get it wrong (by the lights of others). U might have information inducing in her $\neg D$ -token Π -[*S*] which we would not accept as the right kind of inducement. But through intersubjective engagement she may be

⁴⁶ Price (1990) seeks to explain negation in terms of something called *rejection*. But Price treats rejection as an attitude to sentences or positions, and treats it as primitive.

set right. The character of the underlying C-laws, however, needs commenting upon.⁴⁷ The C-laws underpinning rejection are different from the kinds of C-laws we have examined so far in §22. The C-laws we are looking at now are of the form $\Psi \Rightarrow \neg D\text{-token } \Pi\text{-}[S]$. They link a Π -property, Ψ , to something that is not a Π -property, $\neg D\text{-token } \Pi\text{-}[S]$. This is not a Π -property but a constraint on not tokening a Π -property. We might call C-laws of this form *laws of exclusion*. Apart from this difference, exclusion laws are just like any other C-law. They can be demonstrative or inductive.

Here are some examples of such laws. U is constrained from tokening $\Pi\text{-}[T \text{ is red}]$ through tokening of $\Pi\text{-}[T \text{ is blue}]$, and constrained from tokening $\Pi\text{-}[T \text{ is a cat}]$ through $\Pi\text{-}[T \text{ is a dog}]$ via the exclusion laws:

$$\Pi\text{-}[T \text{ is blue}] \Rightarrow \neg D\text{-}\Pi\text{-}[T \text{ is red}].$$

$$\Pi\text{-}[T \text{ is a dog}] \Rightarrow \neg D\text{-}\Pi\text{-}[T \text{ is a cat}].$$

These exclusion laws correspond to the C-laws underpinning conceptual reasons.

An important set of exclusion laws underpins negative existential statements. Say U searches the room for a cat. U examines all points in the room but at each point her cat detection system D_{Cat} —which is activated with perception of cats—remains untriggered. U's perceptual search has been given every opportunity for U's cat detecting system, D_{Cat} , to be triggered. Yet it remains untriggered. Let this state of thorough search without triggering of D_{Cat} be designated by *Search-No- $[D_{\text{Cat}}]$* . U might defend this state in her assertion of *There appears to be no cat*. U infers, inductively, that there is no cat present. The C-law underpinning this inference, more or less, is the (non-intrinsic) C-law below:

$$\{\text{Search-No-}[D_{\text{Cat}}], \Pi\text{-}[\text{No-defeaters}]\} \Rightarrow \neg D\text{-}\Pi\text{-}[\text{There is a cat}].$$

It could turn out that there is a cat after all, despite appearances.

⁴⁷ Negation is not essentially linked to falsity. There are so called meta-linguistic negations such as *I don't like cricket, I love it*. See Barker (2004) for a discussion of these.

Contradiction and De Re Incompatibility

Where U instantiates the C-law $\Psi \Rightarrow \neg D\text{-token } \Pi$, U cannot, given the normal functioning of her NC system, be in a state where she accepts both Ψ and Π . They are subjectively incompatible for U. This is the ultimate basis for our horror of contradictions, why contradictions seem unintelligible in some sense. The sense of unintelligibility is the cognitive incompatibility of having an NC system directed towards tokening $\Pi\text{-}[S]$ and not tokening $\Pi\text{-}[S]$.⁴⁸

Exclusion laws define incompatibilities that have a purely causal functional basis. That does not rule out the legitimacy of talk of incompatibilities in the world. We can assert that being red is incompatible with being blue, being a dog is incompatible with being a cat, and so on. These worldly incompatibilities between properties are real. They are not fictions. When we claim that being red is incompatible with being blue we are not making a claim about our cognitive systems, we are making a claim about the world. We make this claim about the world by defending an exclusion law. That might seem paradoxical, but there is no paradox. Property talk is perfectly legitimate in GE—see chapter 7. There are relations between properties, which include relations of exclusion. This talk is legitimate. It is just that in explicating the talk—providing an analysis of the language agency underpinning it—we make no appeal to correspondence relations to such relations in reality. That is why, having given the analysis of the language agency involved, there is no metaphysical question left about the nature of properties and relations of exclusion.

Here is a parallel with negation. In asserting *Not-S*, U is not reporting that she has a grounded constraint with respect to tokening $\Pi\text{-}[S]$. U is not reporting on the subjective state

⁴⁸ Priest's (1995) *dialetheism*, the doctrine that there are true contradictions, is not ruled out. It might seem that it is. In asserting *P* and *not-P*, U defends Π and rejection of Π . Is this essentially irrational? Let us suppose the world just is inconsistent in that it has contradictory states of affairs. It could be that in describing this world adequately, U instantiates a complex Π -property Ψ that she defends through assertion of *S*. The state Ψ is such that in U, $\Psi \Rightarrow \Pi$ and $\Psi \Rightarrow \neg\text{token } \Pi$, and U is compelled to defend these C-laws. Clearly, the C-laws together generate a disposition in U's NC system that cannot be carried out. It cannot token both Π and a constraint on not tokening Π . So U cannot token simultaneously both Π -properties of *P* and *not-P*. But that does not mean she could not assert them sequentially. Furthermore, it may be that this is the best she can do. Her best description of reality requires that she does this.

of her rejecting Π -[S]. The parallel with *is true* is exact. In asserting *That S is true*, U is not reporting her acceptance of Π -[S], she is defending it.

Because of the gulf that separates reporting a subjective state and expressively defending one, the sentences *I reject Π -[S]* and *Not- S* are not insubstitutable *salva veritate*. They do not have the same speech-act structure. So, for example, we may assert *Either O is pink or not pink*, but do not assert, and are not committed to asserting, *Either O is pink or I reject Π -[O is pink]*, or indeed, *Either I accept Π -[O is pink] or I reject Π -[O is pink]*.⁴⁹

This also relates to firmly distinguishing pragmatic from formal contradiction. GE does not conflate the two. In asserting *S but I reject Π -[S]*, U produces a pragmatic contradiction because the speaker is defending commitment to Π , but reporting her rejection of Π . This is quite distinct from assertion of *S but not- S* in which the speaker defends Π , and defends rejection of Π .

§24 Disjunction

We examined logical compounds from the point of view of embedding in §15. Logical compounds are compound speech acts in which (a) proto-assertions are embedded and (b) whose Π -properties are compounds formed from the Π -properties associated with the embedded proto-assertions. Our concern now is with the Π -properties defended in such compounds. We have looked at conjunction and negation. We now consider disjunction.

In asserting a disjunction, U utters *Either S or R* and (i) performs proto-assertions with S and R , and (ii) by concatenating *either...or* with S and R , U defends a compound Π -property **D**. What is the state **D**? One might think the inference patterns below will help us to decide what the structure of **D** is:

DI: $P \vdash \text{either } P \text{ or } Q$

DE: $\text{Either } P \text{ or } Q, P \rightarrow R, Q \rightarrow R. \vdash R$

⁴⁹ It is failure to deal with such problems of substitution that plague expressivist theories in metaethics. See van Roojen (1996) who uses a substitution argument against Blackburn. Blackburn has no reply because he has no viable theory of what expressing is.

What we want is a mental element that generates dispositions to follow such inference forms. Our discussion of conjunction above in §22 will help us. The Π -property of P and Q is the state $C[\Psi, \Omega]$, which is the union of the causal powers of P and Q , that is, Ψ and Ω . That suggests that the Π -property of *either* P or Q is the state $D[\Psi, \Omega]$, which is the intersection of the causal powers of Ψ and Ω . The state $D[\Psi, \Omega]$ is simply that functional element that has the common powers of Ψ and Ω . If that is $D[\Psi, \Omega]$'s functional specification, the following C-laws can be deduced to hold for $D[\Psi, \Omega]$:

$$C_I: \quad \Psi \Rightarrow D[\Psi, \Omega]$$

$$C_E: \quad D[\Psi, \Omega], \{\Psi \Rightarrow \Pi\}, \{\Omega \Rightarrow \Pi\} \Rightarrow \Pi$$

C_I holds because if U 's system tokens Ψ , then a state that has a sub-set of Ψ 's powers can be derived as an intrinsic causal consequence. That is what $D[\Psi, \Omega]$ is. C_E holds because if $\{\Psi \Rightarrow \Pi\}, \{\Omega \Rightarrow \Pi\}$ hold in the NC system, then Π is an intrinsically derivable state from both Ψ and Ω , which is to say that a power to produce Π is in the intersection of the powers of Ψ and Ω . So Π is causally derivable from $D[\Psi, \Omega]$.

Inferential Intuitions and Inferentialist Parallels

We see, as in the case of negation, that GE assigns functional elements to logical constants that are not themselves precursors to inference, C-laws, but rather states from which C-laws can be causally derived. Dispositions to perform inferences flows from the causal powers of Π -properties associated with logical constants.

§25 Validity and Rightness

We now examine the functional character of Π -properties of statements of validity. The basic idea is that a statement of the form A is valid, where A refers to an argument, from $R_1, R_2 \dots$ to S is a Π -property that is ultimately based in cognitive processing to do with C-laws. Given the possibility of comprehension transcendent arguments, and assignments of validity thereto, the account we have to give of the structure of statements of validity is given in terms of basing in the following way. Suppose that A refers to the argument $R_1, R_2 \dots$ to S , which U is

capable of comprehending. We then have the following C-law, where D_{Valid} is the mental element underpinning the predicate *valid*:

$$\{\{\Pi-[R_1], \Pi-[R_2] \dots\} \Rightarrow \Pi-[S]\} \Rightarrow \Phi-[A] + D_{\text{Valid}}$$

The left-hand side is the C-law that corresponds to the argument. It takes us from the Π -properties of the premises to that of the conclusion. This is a cognitively sufficient condition for the Π -property structure $\Phi-(A) + D_{\text{Valid}}$. Thus, speakers may base their judgements of validity on intuitive grounds, in which case they can instantiate a C-law corresponding to the argument concerned, or they need other grounds, which may include, say, the outputs of a supercomputer in the case of arguments that defy comprehension by humans.

So again, we can sum up the functional theory in these terms:

Validity	Π -Property Specification
<i>Argument A is valid</i>	$\Phi-[A] + D_{\text{Valid}}$, where D_{Valid} is based in a grounding relation \Rightarrow .

Again it needs to be emphasised that we are giving a theory of that aspect of the language agency that underpins statements of validity. GE provides no theory of validity, or logical truth. It is a theory of the structures that underpin assertions in logic. GE does not attempt to reduce the laws of logic to laws governing cognitive systems. Cognitive laws are C-laws, but the laws of logic are not these. There are laws of logic in the sense that there are facts about what entailment relations between propositions. (See chapter 8 for the examination of fact-talk.) The truths corresponding to these facts are assertions in which the Π -properties defended are Π -properties based in C-laws. But beyond this explication of the truths of entailment and correct inference, there is no further metaphysical investigation into the nature of facts of logical *ought*. There is no metaphysics of *the logical ought*. That might seem like a big claim, but it is a direct entailment of the whole methodological position of GE and its repudiation of representation in semantics. Compare the expressivist position on moral *ought*. Having completed the language analysis of moral talk, there is no metaphysical question left over about the nature of value. Unless you fault the treatment of value-talk as an analysis of

the nature of that talk, you cannot fault the conclusion that no metaphysics of value remains. Of course, many people fault value-expressivism. But that, I submit, is because proponents of value-expressivism have attempted to maintain it as Semantics. As part of a theory of language agency, it works. Moreover, the form of the analysis is global. We have just exhibited it again for the case of entailment-talk.

Truth and Validity, and Modal Facts

Orthodoxy analyses entailment in terms of truth as in **T-Validity**:

T-Validity: The argument $S_1, S_2, \dots, \vdash R$ is valid iff in for every case C in which the premises S_1, S_2, \dots , are true the conclusion R is true in C .

This definition invokes the technical notion of truth in a case. We can interpret this in terms of possible worlds: to be true in a case is to be true in a possible world. Hence an argument is valid if and only if in all possible worlds in which its premises are true its conclusion is true. Alternatively, we have the standard Tarskian definition, according to which a case is an assignment of values. In which case, an argument is valid if and only if on all value assignments on which the premises are true, the conclusion is true.⁵⁰

GE cannot accept **T-Validity** as an explanatory account of what validity is within a theory of language function. This is not to deny that there is something right about the idea that where $P \vdash Q$ is valid, or P entails Q , it is not possible for P to be true and Q false. The following is correct:

$P \rightarrow Q$ iff either (i) If P , it must be that Q ; (ii) If P it could not be but Q ; or (iii) It is not logically possible that P and not- Q

But it does not follow from our affirming this truth that such modal conditions are what we look to in providing a theory about how terms like *entailment* or *valid* function. The language agency in which we assert the obtaining of facts like (i) to (iii) is open to analysis by GE.

⁵⁰ These alternative interpretations of **T-Validity** do not certify the same arguments. Do we need to decide which is the right one? According to some, Beall and Restall (2000)—who accept this conception of validity—we do not have to decide. There is no objective fact of the matter. We should be pluralists about validity and entailment. I will not explore these issues in detail here, though I say something briefly below.

Thus to take (iii), the Π -property of *It is not possible that P and not-Q* is rejection of Π -[*It is logically possible that Q*]. This Π -property is:

Reject: [$\Psi \Rightarrow \neg$ D-token Π -[*Q*]],

where Ψ corresponds to the set of intrinsic C-laws of logical inferences U defends. If so, *It is not possible that P and not-Q* involves expression of a condition involving defence of C-laws. It is then not surprising that we are inclined to appeal to such locutions in articulating entailment or inference. Such modal facts are not more basic than the facts of entailment. Both correspond to truths that involve defence of intrinsic C-laws.

Naturally enough, the standard orthodox semantics treats modal claims as quantifications over possible worlds or comparable entities. It seeks an account of objectivity of entailment of logical law in *entities* or *structures in reality*. A classical instance of this tendency is Lewis's (1986) grounding of logical properties in a plurality of concrete worlds. But this ontic approach cannot be explanatory in the eyes of GE. Just as we cannot explain the validity of bivalence by appeal to a determinately structured reality—see chapter 9—we cannot explain the validity of logical principles by appeal to *in re* modal structures.

We can now clearly state GE's position. GE rejects **T-validity** as an explanatory thesis, but otherwise asserts it.⁵¹ Facts about necessary truth-preservation are not more basic than facts about entailment. Both arise together. The orthodox truth-theoretic conception implicitly holds that truth has an explanatory role in semantics. GE rejects that thesis.

Rightness

Reason brings in the question of rightness. GE makes no attempt to reduce reason to causal regularities. There is no truth-conditional equivalence between claims of reason and claims about C-laws. What then makes reason statements right? A certain mind-set might say that *de re* structures, worlds and their relations determine rightness. They talk of an inference being valid in virtue of the fact that every possible world that contains the premises contains the conclusion. But this kind of answer cannot be accepted by GE. Here again the analogy

⁵¹ The existence of conventional implicature in a language means truth-theoretic conceptions of validity have limitations from another quarter. But I will not examine that matter here—see Barker (2003).

with the ethical expressivist is meant to help. What makes value statements right? Not the motivational states that we express in asserting them. Their role is merely to explain why we are inclined to assert certain value-statements and to assert their rightness. The question *What is rightness?* betrays misunderstanding. Having given the speech-act structure of claims of the form *T is good*, and *That T is good is right*, there is no further question to ask about good or rightness. Claims of rightness are just defences of the Π -properties of the sentences judged to be right, likewise, for reason-statements. The question *What is rightness in reason?* betrays the same misunderstanding. C-law structures do not explain what rightness is. They explain only our intuitions that certain kinds of inference are right.

One might object that this just shows the inadequacy of GE. Consider two speakers who differ on C-laws, and feel compelled to defend different reason statements. Who is right? Each speaker will defend her state. But it seems, the objection goes, it just comes down to opinion. No ontic worldly arbitrator can be appealed to. So, there is no ultimate fact of rightness. If so, GE cannot be maintained. Representation of *de re* structures must be brought into semantics to explain objective facts of rightness in reason.

The argument is flawed. The first thing is that the objector's inference that the question of rightness is reduced to opinion is entirely fallacious. In GE, one's opinion, the fact that one is inclined to defend Ω , that is, to assert *S*, where Π -[*S*] does not make it right to assert *S*, it does not make it true or right that *S*. No opinion makes reason claims true.

GE's position is not relativism. According to GE, an argument's validity is not constituted by its agreement with the inferential behaviours of members of U's community. Might does not make right. That would be a kind of naturalist relativism. So the might-is-right principle cannot be correct. GE replaces such relativism with a conception of reasoners as situated. We judge things by our lights. It could be that all my co-reasoners start to employ a certain inference pattern, which I cannot make sense of. Are they right because they are in the majority? No. Who is right? That question is not answered by counting heads, but by looking within our own breasts, and engaging intersubjectively with fellow reasoners. I

defend my Π -properties, they defend theirs. If agreement is unresolved by debate then we continue to disagree. There is no higher authority to resolve the deadlock.⁵²

The other side of the coin is that appeal to ontic structures does nothing to explain objective rightness. It simply begs the question. This is initially hard to believe, but I make a case for it later—chapter 9, §71.

⁵² Is there not the possibility of relativism? What of a showdown between two reasoners who cannot come to agree. If the view about intrinsic C-laws underpinning inference is right, then we must say they have a different understanding of the connectives. That means pluralism about logical inference—of the kind defended by Beall and Restall (2000), is ruled out. If on the other hand, not all reason is fixed by C-laws of the intrinsic kind then a kind of pluralism is available. However, it does not generate relativity, say relativity of validity or truth. What we find is faultless disagreement, where speakers can come to agree to disagree. Their views about what is right contradict each other, but there is no objective fact about who is right. As I show in chapter 9, such situations, which I do not think arise very often, do not generate relativity of truth, but only local non-objectivity of truth. That needs explaining, but chapter 9 is the place to do it.

Chapter 4

Expressivism about Everything: The Non-Evaluative

§26 Π -Properties, Concepts, and Interpretative Stances

GE is a functionalist analysis of the language-using agent. It describes the functional structure of the acts, cognitive activities, and states that constitute language agency. The basis of the speaker's capacities resides in an underlying functional system, which I have called the NC system. The basic elements of this system are pre-doxastic functional states. The NC system is a system of elements that include Π -properties, which are related by what I have called C-laws—cognitive regularities. These may take the exclusion form or non-exclusion form. The NC system is the sub-stratum of a signalling system—the foundation for a language agent—that links elements of the NC system to productions of signs. In the table below, I display some of the basic structural features of a signalling and NC system:

Language-Agent State/Act	Characterisation
Π -property	Pre-doxastic state based in perception, affection, etc.
Ground-signalling Disposition: Π^{G1}	Disposition to signal C-law relations grounding Π .
Assertion Repertoire Disposition: $RD[\Pi^{G1}, S]$	Disposition wherein if U is directed towards signalling Π^{G1} she may produce S .
Proto-assertion of S	Production of S partly caused by a repertoire disposition: $RD[\Pi^{G1}, S]$.
Assertion/Judgement of S	Production of S in a proto-assertion where the agent is directed towards suppressing cancellation signals.
L-processing	Disposition of the language system to simulate production of locutionary acts using elements of the NC system.

GE is a theory of mind without being a theory of content: it is a theory of mental architecture. Π -properties and related states are aspects of mentality. Whilst GE does not give a functional analysis of what it is to believe a content P , it gives a functional analysis of what it is to have

a belief of a specific kind, say a belief corresponding to a value-judgement or one involving a specific logical form or category of predicate, say a colour predicate.

Π -properties are theoretical entities. Π -properties are, or are based in, affective states, perceptual states, and dispositions to perform linguistic acts of a certain kind. The theory of basing was given in §17 and applied to truth-, entailment-, and validity-statements in the last chapter. In the table below, I list linguistic forms and the corresponding kinds of Π -properties—it is not intended to be exhaustive:

Language Element	Π-Property identical to or based in:
Formal predicates— <i>true, identical to, exists, object, fact</i>	Metalinguistic dispositions—disposition to produce speech-acts with certain forms.
Logical predicates and Logical particles— <i>entails, valid, not, if, either...or.</i>	Mental manipulations and processing of Π -properties.
Empirical predicates: <i>red, dog.</i>	Perceptual states.
Value predicates and modals— <i>good, ought</i>	Affective states and motivation.
Causal and Dispositional predicates	Dispositions or capacities to manipulate reality.
Meaning and Interpretative claims	L-processing (simulation states).

We can divide the states upon which Π -properties are based into two kinds: those that are complex, being operations on Π -properties, and those that are not. Complex Π -properties are those states in rows one, two and six. The states in the table that are simple Π -properties are those in rows three to five: states based in perceptual states, affective states, and dispositions to manipulate reality. The core states from which Π -properties are derived are the latter: perception, affective states, and agency. We have already described in some detail some of the Π -properties specified in the first two rows of the table. The truth-predicate is a formal predicate in that it is based in a disposition to perform assertions. Logically complex sentences, entailment-sentences, and statements of reason have Π -properties of the second kind. Below we look at Π -properties of the first, third and fifth kinds.

Content Attribution and Functional Structure

In what follows we are going to be looking at Π -properties from the side of Γ -states, the mental precursors of predicates. Here is a significant fact about predicates according to GE, that I have not been fully explicit about. Predicates can be divided into two kinds: those

that have underlying functional essences and those that do not. We have described some predicates with functional essences: *true*, *valid*, *truth-apt*. Call predicates, general terms, or relational terms that are based in perception *empirical terms*. So sensory modality predicates, natural kind predicates, artifact terms, and so on, are empirical terms in this sense. Call terms like *plus*, that is, terms associated with the application of algorithms *calculative terms*. Empirical and calculative terms lack functional essences. Predicates and particles whose underpinning states possess functional essences have this feature: If two speakers U and H are interpreted as using negation, then the mental elements underpinning their use of these particles will have the same functional form—we characterised that form in §23. Likewise, the precursor of the predicate *true* also has a functional characterisation—see §20. This does not mean that the *meanings* of *not* and *true* are these states. Any such affirmation would be to embrace a crude form of psychologism about negation and *true*.⁵³ Attributions of content to particles like *not* and *true* have the same L-processing form described in §17. Contrast now empirical and calculative terms. For example, if we take two speakers U and H who both possess the concept of red, they will use terms employing Γ -states, Γ_U and Γ_H , that will enable them to use predicates that can be interpreted as meaning red. GE can point to common structural features of U and H's states, identifying these states as those characteristic of colour terms. But there is not necessarily a functional identity in Γ_U and Γ_H .

Why is it that empirical and calculative terms are such that the underlying mental elements do not correspond to functional patterns? There are two main reasons.

Holism: Our application of empirical predicates partly reflects theories we have of the world. Say we use *red* to mean *Red-Surface*. Then our use of *red* may depend, in part, on our theory of surfaces, of perceptual situations, of ourselves as perceiving agents. One does not necessarily have to have such theories, but one could have these theories, and applying these

⁵³ One might find oneself thinking, from time to time, what work Π -properties do. Are they not useless cogs in a mechanism? As inner states, they are not publicly manifestable, although their effects are. Should we not dispense with them in providing a theory of meaning? This is a familiar argument, but it is misapplied here. GE is not offering a theory of meaning. Π -states and related states are not meanings. They are not useless cogs but vital to the specifications to language activity. Also, the idea of a theory of meaning that informatively provides an account illuminating meaning in terms of public conditions is misplaced.

theories to specific potential instances of the application of *red* can affect inclinations to apply the term.

Resource-Boundedness: Speakers have cognitive and perceptual limits. They have finite memories and attention spans. These factors can condition their application of predicates. The attribution of concepts to speakers is an interpretative matter. Meaning as proposed in §1 is irreducibly interpretative. Interpretative norms take into account both factors summed up in *Holism* and *Resource-Boundedness*, but such factors have limited impact outside empirical and calculative concepts.

GE offers no theory of concepts or meanings. There is none to be had. (Indeed, as I will argue in chapter 10, meanings are non-objects, and non-objects lack natures.) There is no view-from-no-interpretative stance about what people mean. All that the functionalist analysis of language agency can provide is: (i) a functional taxonomy of kinds of predicates; (ii) the form of concept-attribution sentences; (iii) a probabilistic account of the likelihood of content attributions. Below I describe the Π -properties underpinning other formal predicates, empirical predicates, and causal and dispositional terms. I look at value-sentences in chapter 5. Predicate and verbs related to meaning-attribution are discussed in chapter 10.

§27 Formal Predicates

By formal predicates I mean a family of predicates including *true*, *correct*, *identical to*, *exists*, *object*, and *fact*, amongst others. The Γ -properties of these terms are states based in metalinguistic dispositions, which is to say, dispositions to perform certain kinds of speech acts. I shall discuss, in some detail, *exists*, *object*, and *fact* in chapters 7 and 8. For now, I shall say something briefly about *identical to*. In terms of standard semantics we think of identity as a relation, ontically construed. It is a function or set or ordered n-tuples, or perhaps an abstract entity instantiated by every object. Typically identity is treated as a simple primitive relation. There seems to be no theory about what it is, or what our grasp of it resides in. In GE, no predicate or concept can be primitive, since for every predicate there is

the activity and cognitive precursors that underpin use of the predicate. (Only in Semantics is the idea of primitive concepts intelligible.)

In essence the approach within GE is this: In asserting *T is N*, U performs an assertion with two constituent acts produced with *T* and *N*. These may be referring acts, but not necessarily as we shall see in §75. U defends a disposition to substitute speech-acts of these kinds in all non-oblique contexts. Oblique contexts have a specific speech-act structure, which need not concern us too much here. (See Barker 2004). The relevant kind of speech-act type identity is that U assigns the same mental elements to the speech-acts. If they are referring terms then they are Φ -states or corresponding repertoire dispositions. (I examine these states in chapters 6 and 7.) Thus for a speaker a certain set of tokenings of *Hesperus* and terms like *the morning star*, will be associated with the same Φ -state. Again, Φ -states are internal cognitive states. They are not publicly graspable things, and we should not even think of them as concepts. *Concept* is a term from within the interpretative stance, and at this stage our analysis is confined to the functional stance.

Let us represent a substitution disposition as $IS[T,N]$. In which case, the Π -property structure is:

Identity	Π -Property Specification
<i>T is identical to N</i>	$IS[T,N]$ —a disposition to inter-substitute the speech-acts <i>T</i> and <i>N</i> in all non-opaque.

The structure of evaluation then has the general form given by **ID-True**—§17:

ID-Identity: H judges true an identity assertion of *T is N* uttered by U if and only if H is disposed to defend $IS[T,N]$, the state that H deploys to L-process U's sentence.

ID-Identity furnishes us with the language agency associated with truth-attributions to identity statements. It is not a specification of truth-conditions. Truth-attributions are always cognitively perspectival, from the point of view of the L-processing language agent. GE's account is a description of the structure of that perspective.

The Γ -state associated with *identical to* is a metalinguistic one. It is a disposition to inter-substitute speech-acts. This does not mean that identity statements are metalinguistic. Their truth-conditions cannot be formulated in terms of substitution. This is a general feature of GE, summed up in §11: assertions are not about their Π -properties. Just as moral statements are not about motivational states, identity-statements are not about dispositions to substitute speech acts. GE is not a theory of truth-conditions or meaning. It does not claim that the meaning of an identity-sentence is a disposition to substitute speech acts.

Inferentialism and the Metalinguistic

Brandom (1994) offers a theory according to which to assert *T is N* is to make explicit an inferential commitment to substituting *N* for *T* and vice versa in all non-opaque contexts. Here expressivism and inferentialism are close. The difference is the normative-inferential conception applied in Brandom's case. Brandom also must affirm that identity statements are not metalinguistic. Why doesn't *T is N* have the same truth-conditions as the sentence *M*:

M: '*T*' can be inter-substituted with '*N*' preserving truth.

Here there is a puzzle. According to Brandom, a substitutional inferential practice is made explicit by *T is N*. But we might also say that *M* makes that practice explicit as well. In which case, how does the identity claim differ in meaning from *M*? Clearly, *T is N* uses the terms *T* and *N*, whereas *M* mentions these terms through quotation. True, but what aspect of Brandom's theory explains this difference? Brandom gives no substantive theory of referring terms—see §43—so he provides no substance for the view that in *T is N*, the terms *T* and *N* are used rather than mentioned. We cannot resort to the fact that *M* entails that there are terms in existence, whereas *T is N* does not. What is the source of that entailment? One might say the fact that that we can infer from a sentence mentioning a term that a term exists. But this assumes a theory of mentioning rather than using, which is the very issue under question. Note that Brandom cannot reduce mentioning to syntax, say quotation marks, since not all speakers use them or are clear about their use—only analytic philosophers tend to do that—and quotes do not appear at all in spoken speech.

Evaluation

One might wonder what goes into defending a state of this form. What kind of C-laws does a Π -property of the form $IS[T,N]$ enter into? Typically the basis for identity statements is information about qualitative likeness. So a basic kind of C-law that underpins defence of a state of the form $IS[T,N]$ is the following. Let us denote the disposition of being disposed to defend $\Pi-[T \text{ is } F]$ if and only if one is disposed to defend $\Pi-[N \text{ is } F]$ by $DD(\Pi-[T \text{ is } F], \Pi-[N \text{ is } F])$. In which case the following is a C-law of a non-intrinsic inductive kind—see §22:

$$\{DD(\Pi-[T \text{ is } F], \Pi-[N \text{ is } F]), \Pi-[No \text{ Defeaters}]\} \Rightarrow IS[T,N]$$

The C-law is inductive because it might be that indiscernibles are non-identical. The inference is not deductive.

Paradigmatically, judgements of the falsity of $T \text{ is } N$ are grounded in the contraposition of the indiscernibility of identicals. We can have C-laws of the form:

$$\{\Pi-[T \text{ has } F], \text{Reject: } \Pi-[N \text{ is } F]\} \Rightarrow \neg\text{D-token } IS[T,N]$$

In this case the C-law is intrinsic, that is, logical. In an assertion of $T \text{ is not } N$, U defends a rejective state, that is, defends tokening some Ψ such that $\Psi \Rightarrow \neg\text{D-token } \Pi-[T \text{ is } N]$. In this case the state Ψ could be the compound state: $\{\Pi-[T \text{ has } F], \text{Reject: } \Pi-[N \text{ is } F]\}$.

GE's approach does not validate the identity of indiscernibles. Two objects, T and N could share all qualitative predications, but it does not follow that they are identical. Suppose T and N are two spatiotemporally coincident monads, all of whose features are necessary. They are perfectly indiscernible and could not have had properties other than those they actually possess. It does not follow that they are identical. U, apprised of all the facts in relation to T and N, is not committed to substituting T for N in the context $N \text{ is identical to } N$. This is despite the fact that U is committed to inter-substitution for all other open-sentence contexts. (Leave aside for the moment psychological contexts of the form $U \text{ is thinking about } N$.) U may have very good evidence that T and N are identical, but the inference to any non-trivial judgment of identity is empirical. Thus, in terms of inference to the best explanation, there is a very good case that T is N. But this is not entailed.

Attitudes and Embedding

Identity-sentences can embed in logical compounds. In these contexts, U uses the identity-sentences in a proto-assertion—§15. Proto-assertions are production of terms that are proto-signalling acts, and that means they are productions of symbols through repertoire dispositions. Thus the repertoire disposition for a particular speaker for *Hesperus is Phosphorus* is:

If U is directed towards signalling a ground-indicating disposition in relation to IS[*Hesperus, Phosphorus*], U can produce a sentence: *Hesperus is Phosphorus*.

The embedded sentence, say in, If Hesperus is Phosphorous then they share properties. is an utterance partly caused by this repertoire disposition, though it is not one where that disposition is manifested.

Ascriptions of beliefs about identity are also embedded contexts, as in: *Arnold believes that Hesperus is not Phosphorous*. Belief-attribution sentences can introduce opaque contexts. To explain opacity we do not have to suppose that referring terms lack their customary referents in such contexts and denote *senses* or *modes of designation*—as Frege (1892) argues. Rather, all we need to propose is that the verb *believes*, in opaque belief-ascriptions, is sensitive to the syntactic-phonological-graphemic forms of referring acts within its scope. (See Richard 1990). Thus substitution of co-referring expressions may fail to preserve truth because the syntactic-phonological-graphemic forms of the terms differ.

That is the basic treatment of identity talk. Having given the talk, from the perspective of language function and *Wissenschaft*, there is no substantial question left when we return to the perspective of *Verstehen*, deploying our language-agency capacities to produce identity sentences and interpret others as talking about relations of identity. This is because the language agency underpinning *is identical to* does not involve representation. The capacity of a speaker to use *is identical to* is not a capacity to pick anything out, say, a relation of

identity. If so, the question *What is identity?* is unanswerable. It is a question whose presupposition—that identity has a nature—fails. Identity lacks a nature.⁵⁴

§28 *Perceptual Predicates: Colour, a Case Study*

The Π -properties of sentences with empirical predicates are based in perceptual states. In what follows I shall provide a general sketch for how we analyse the language agency of colour predicates. I concentrate on the colour predicate *red* as my case study. We cannot give the complete cognitive story here. That is a matter for a fully-fledged cognitive science. Although we give the overall form of analysis, it is partially speculative. This analysis proceeds in three stages: (i) I give an account of how a visually unimpaired speaker generates a state a property D_{Red} , the mental element for *red*, and generates Π -properties for visually based predications with sentences *N is red*; (ii) I explain how visually impaired speakers assign mental elements to *red*, and derive their Π -properties; (iii) I explain the phenomenon of context-dependence that judgements using *red* seem to be prone to. (iv) I look at the structure of D_{Red} across speakers and explain its relation to concept and meaning identity. I deal with (i) and (ii) in this section, (iii) in §29, and (iv) in §30.

Consider competent users of the colour predicate *red* who are sighted. An assertion of *N is red* has a Π -property with the structure $\Phi-[N] + D_{\text{Red}}$. This comprises $\Phi-[N]$ the mental element underpinning U's production of the referring term *N*. We shall describe the structure of such elements in detail in chapters 6 and 7. The state D_{Red} is the Γ -state of the predicate, and its general nature should be familiar to us. It is functionally defined through its combinatorial powers including *basing*. (We have applied this framework to *truth-apt* (§16), *true* (§20), and *valid* (§24).) First, D_{Red} is partly defined by its capacity to combine through + with a Φ -state, to form a Π -property of the form $\Phi-[N] + D_{\text{Red}}$. The interesting aspects of D_{Red} 's functional nature, those aspects that tell us what kind of predicate *red* is, all come with *basing*. The story about basing for D_{Red} comes in two stages.

⁵⁴ Undoubtedly, there are further conceptual questions about identity. We might be concerned about the relation of identity to sortals—Wiggins (2001). But these are not, according to GE, metaphysical questions.

The first stage is a perceptual state R caused by certain environmental conditions involving an object O . This state is a veridical or non-veridical perceptual state. The information in R generates the joint triggering of the states in $\Phi-[N]$ and D_{Red} . This means that the property $\Phi-[N]$ is partly functionally defined by its ability to be caused by perceptual states. Likewise for D_{Red} . Both these states can be activated; which is to say they are dispositions which can, through perceptual states, be triggered. (I say more about this aspect of $\Phi-[N]$ in chapter 6.) Let us represent this triggering in terms of on-states: $\text{On-}\Phi-[N]$ and $\text{On-}D_{\text{Red}}$. Let the joint triggering by information in R of both states be $[\text{On-}\Phi-[N], \text{On-}D_{\text{Red}}]$. We suppose that the triggering of D_{Red} depends not only on the perceptual episode R , but also on interaction with Π -properties and C -laws; in short, with what we would interpret as U 's other beliefs. (If *red* in the context can be interpreted as red-surface, then these other beliefs relate to U 's theory of surfaces.) In which case, we have the following cognitive law:

$$\text{L1: } \{R \text{ given processing}\} \Rightarrow [\text{On-}\Phi-[N], \text{On-}D_{\text{Red}}]$$

Let us say that $[\text{On-}\Phi-[N], \text{On-}D_{\text{Red}}]$ is a perceiving-something-as state.

That is stage one. Stage two is just another basing law to derive the Π -property of N is *red*—as used by the visually unimpaired speaker:

$$\text{L2: } [\text{On-}\Phi-[N], \text{On-}D_{\text{Red}}] \Rightarrow \Phi-[N] + D_{\text{Red}}.$$

L2 is functionally just like the other basing C -laws we described. On the left-side is the perceiving-as state, which already involves predicate precursors, the right-side is the Π -property, which is not a perceiving-as state at all. One can token the right-hand side when one does not perceive anything. The left-hand perceiving-as state is just that the primary sufficient condition for its generation of the Π -property $\Phi-[N] + D_{\text{Red}}$. The state $\Phi-[N] + D_{\text{Red}}$ has as a constituent, we suppose, the dispositional state D_{Red} , but its tokening and concatenation with $\Phi-[N]$ does not require that it is triggered. Its triggering only goes on in perceptual states—or perhaps at a low level in imagination. Its concatenation with $\Phi-[N]$ is just its tokening. Given that D_{Red} is an abiding feature of an NC system, one may wonder

what this concatenation, which may not abide, amounts to. But this is a question about cognitive architectural detail that we need not explore here.

The primary ground for $\Phi-[N] + D_{\text{Red}}$ is a seeing-as state—as indicated by L2—which in turn is generated by a perceptual state—as displayed in L1. Seeing-as states, however, are just the primary grounds for states of the form $\Phi-[N] + D_{\text{Red}}$. As we have seen with other basing states—§20—a given Π -property, $\Phi-[N] + D_x$, where D_x is the Γ -state of a predicate used by U , can be tokened not only through a primary ground, but also through logically complex grounds with Π -properties that deploy $\Phi-[N] + D_x$ as a constituent. Thus, one might learn *N is red* through a friend describing N as red. In this case one goes through an inferential process just like that which we go through when we learn that, say, *\$ is true*—see §20. That includes our description of the basic functional features of D_{Red} , the Γ -state of *red*, used for the colour. (Note, this is a bit of a simplification, as we shall see in §29, but it will do for now.)

Neither the state $[\text{On-}\Phi-[N], \text{On-}D_{\text{Red}}]$ nor $\Phi-[N] + D_{\text{Red}}$ is a belief state in itself. The belief state is $\Phi-[N] + D_{\text{Red}}$ in combination with a defensive disposition. Defence involves dispositions to use symbols and signalling activity. Likewise, one might be tempted to call D_{Red} a concept. But *concept* is a term associated with content, and content has no natural reduction. We are in the mode of *Wissenschaft* at this point, and cannot talk of concepts. I shall discuss this matter of concepts and concept identity further in §30-31.

Non-sighted Speakers

That is the basic account of the visually unimpaired speaker's language agency using *red*. What of visually impaired speakers? A blind speaker can make assertions using *red*, but cannot be defending the state $\Phi-[N] + D_{\text{Red}}$, since they cannot token that state. Such speakers must have other kinds of Γ -properties that underpin their use of *red*. By virtue of their causal interactions with speakers who possess D_{Red} —which involves testimony about matters red—they are informed that red things are coloured things, pillar boxes are red, etc. They acquire use of a term *red*. How do they do this? My suggestion is that they need to develop a state, call it a proxy state P_{Red} . In a speaker U , P_{Red} has two functional features: (a) it is correlated in the right way with the predicate *is red*, that is, U has the disposition that if she wants to

defend $\Phi-[N] + P_{\text{Red}}$, she utters *N is red*; (b) It is caused through U's testimonial interaction with possessors of D_{Red} .

The states P_{Red} and D_{Red} differ markedly in terms of their internal structure. They are functionally distinct. But then, as we shall see, there is not necessarily functional identity between instances of D_{Red} across sighted speakers. Both D_{Red} and P_{Red} are mental mechanisms that enable speakers to engage in language activities: they are not concepts, meanings, or the conditions that define meanings. What we need to explain is the conditions under which speakers assert that, for example, O_S a sighted speaker and O_B a non-sighted speaker mean the same by *red*. That amounts to analysing L-processing sentences, ones of the form *O_S and O_B both mean RED by 'red'*. In asserting such a sentence, a visual speaker U defends a disposition to L-process U and H's utterances of *red* using her own state D_{Red} . Indeed, U will, in the right context, spontaneously L-process utterances of *red* by O_S and O_B in terms of her state D_{Red} . Unless O_S or O_B act in ways that disagree with U's own dispositions to apply the term—and these ways cannot be explained through belief-divergence and cognitive or perceptual limitations—U will retain that L-processing disposition. (More will be said about interpretation in §30-31 and chapter 10.)

§29 Context-Relativity and Open Texture

I suggested above that I was simplifying a bit. The general term *red* is really a family of colour terms. One term corresponds to the pure appearance concept, which allows for no distinction between appearance and reality. But *red* is used in other ways than this, ways that allow for a distinction between appearance and reality. There is a use of *red* according to which white objects under red light are not red, they are white. We have moved away from *appearance-red* to *surface-red*. The mental element of *surface-red* is a compound state that has D_{Red} as a constituent. The Π -property of *N is red* with *red* so understood is: $\Phi-[N] + D_{\text{Surface-Red}}$. Here there is cognitive architecture. This state comprises D_{red} along with a dispositional state registering a complex appreciation of lighting conditions and causal dependencies. The state $D_{\text{Surface-Red}}$ is a kind of conditional state. Given activation of certain

cognitive factors, the triggering of D_{red} counts as the on-state of $D_{\text{Surface-Red}}$. Using *red* so understood, it is possible to assert: *X looks red but is not really*. We can apply *appearance-red*, but not *surface-red*.

There are developments of *red* other than *surface-red*. Indeed, there is no limit to the developments that can arise. Consider:

A: This apple is red.

We might think that having fixed the referent of *this apple*, and fixed the meaning of *red* as surface-red, that for any possible world, leaving aside vagueness, the truth-value of *A* will be determinate. But there is an argument that this is not the case. Facts about the interests in conversational contexts in which *A* is evaluated affect our truth-value assignments to *A*.

Suppose there are two contexts C1 and C2. In one we are interested in apples that are red all the way through, and in the other, C2, we are just concerned with apples whose skins are red. In C1 we judge *A* false, and in C2 we judge *A* true. Or C1 might be the context where we are interested in apples that are red by natural means only, and not red, say, by being painted. If *this apple* picks out an apple covered in red paint, *A* is false in C1 but true in C2, and so on.

Closely related to the phenomenon of context dependence is the phenomenon of *open texture*. Go back to *red* understood as *surface-red*. Consider now a range of circumstances under which we might consider whether a certain object is red. We could ask the following questions:

1. Red if red overall except for speckles of yellow?
2. Red if momentarily flashing yellow every thirty seconds?
3. Red if it looks black from a certain angle?
4. Red if in loud noise it turns black?

Given the *surface-red* reading of *red*, it is not obvious what our answers should be to these questions. One would have to say that answering the questions 1 to 4 depends on what our conversational purpose is. If one is keen on not having yellow speckles, then the answer to 1 is *no*. But if one is not bothered by them the answer can be *yes*. If one is not bothered by the thought that the object loses its red-colour when bathed in loud noise, then *yes*, but otherwise

no. Without further specifications of purpose, it seems that there is no determinate answer. Let us say we make the further specification explicit. So, let us suppose that our surface-red object should have no speckles of yellow, not be painted, not look black from a certain perspective. Then we can always generate a new set of questions about conditions of evaluation for the predicate so understood, which again, will lack determinate answers. The indeterminacy cannot be stamped out by further descriptive specifications added to the predicate. Indeed, the further specifications induce as many new avenues of indeterminacy as they close off. This indeterminacy of predicates is called *open-texture*.

We have raised two related issues. The same sentence *A*, it seems, can, in different contexts, convey different contents so that *A* can be judged true relative to one but not another. These different readings can be made explicit by adding descriptors to the predicate, as in *surface-red-non-painted-with-no-speckles*. But any explicit form still suffers the same context relativity, and furthermore, it suffers from open-texture. We can always conceive of a set of conditions under which there is no basis for judging true or false *A*, or *A* with further descriptive specification. The open-texture of a predicate drives the introduction of more refinements, since as new cases of indeterminacy are discovered, new refinements can be added to deal with such cases, generating new forms of indeterminacy.⁵⁵

The phenomenon of open-texture is a deep problem for orthodox Semantics and for the view that the propositional content of a sentence is given by truth-conditions. Truth-conditions are the set of possible conditions under which the sentence *S* is true. These can be conceived of as a set of possible worlds. Typically, the understanding is that the set of worlds fixing the conditions under which *S* is true will be a partition on the universe of worlds. Its complement will be the set of conditions under which *S* is false. But if the proposition is such a set, open-texture is a real problem, since open-texture raises the prospect of worlds in which no determinate evaluation is fixed. But this is excluded by the partition model.

It seems natural then to give up the idea that the set of worlds defining *S*'s truth-conditions is a partition in this sense. Thus the sub-set of worlds in which *S* is true, and the

⁵⁵ See Travis (1989) for a deep extended discussion of open-texture.

set of worlds in which it is false, are not complement sets. There is also the set of worlds in which *S* is neither true nor false. But now we face the problem of articulating a theory of falsity. We need to distinguish worlds in which we have a failure of bivalence from worlds in which there is falsity. What is the difference? We can formally distinguish such worlds by stipulation in our models of the semantics. But if we do that, the falsity of *O is red* at a world *w* is not fixed by the reality at *w* and the meaning of *not* applied to the meaning of *O is red*. The falsity of *O is red* is just a primitive fact but then so is truth of *O is not red*. *Not-S* is an idiomatic lump.⁵⁶

Standard semantics has, then, some difficulty with open-textured predicates.⁵⁷ GE in contrast has no difficulties. The predicate *red* is underpinned by a family of distinct Γ -states, unified by a common thread, the basic Γ -state D_{Red} . This means that the Γ -states of predicates, *surface-red*, *red-all-the-way-through*, etc are complex states that have D_{Red} as a constituent. The structure of these states is basically conditional: if certain perceptually-based states are triggered, and D_{Red} is also triggered, then that constitutes the triggering of the complex state. If the complex state is $D_{\text{Surface-Red}}$, the antecedent conditions are appreciations of surface conditions. If it is $D_{\text{Surface-Interior-Red}}$, the state underpinning *red-all-the-way-through*, then the antecedent conditions are appreciations of surface and interior conditions. The production of sentence *A—This apple is red—*can be associated with any one of a whole family of Π -properties:

$$\begin{aligned} \Pi^0: & \quad \Phi\text{-}(\textit{this apple}) + D_{\text{Red}} \\ \Pi^1: & \quad \Phi\text{-}(\textit{this apple}) + D_{\text{Surface-Red}} \\ \Pi^2: & \quad \Phi\text{-}(\textit{this apple}) + D_{\text{Surface-Interior-Red}} \end{aligned}$$

⁵⁶ Another response is to claim that the appearance of indeterminacy is a kind of illusion. There really is a fact of the matter about what the answers to all the questions are, even if we have no idea. But what determines this fact of the matter? It must be that the rules of use determine these answers. But why are we not privy to the deliverances of these rules. It is not that the cases under consideration are complex, and so speakers' cognitive limitations explain the failure to recognise the right answers. Another conundrum is why would the rules carry information about whether it is right to call something red if, when one squeals the worlds, *Zobra* it turns blue.

⁵⁷ There is a lot of discussion of context dependency in the literature—see Cappellen and Lepore (2005). Most see it in terms of the issue of whether semantics should be context sensitive—and thus introduce what appear to be pragmatic elements in the semantics—or whether semantics should be context insensitive.

Each speaker will have a set of Π -properties structured in this manner, where the sets are sets of homologues. Contexts are simply those conversational factors that will determine implicitly that U will produce \mathcal{A} with one Π -property from one set rather than another. U judges \mathcal{A} true in one context, false in another, because she L-processes \mathcal{A} with distinct Π -properties from the list. On the other hand, for any particular Π -property Π^x assigned to \mathcal{A} , we do not find determinate dispositions to defend it or reject it under all describable circumstances. There will be circumstances under which U is disposed to defend neither Π^x nor its rejection. The Π -property Π^x is a state whose primary ground is a perceptual disposition. That perceptual disposition may get triggered, in which case, U is disposed to defend Π^x . Rejection requires more than simply being untriggered. It requires that the speaker also defends a Π -property Ψ that through a C-law constrains U from tokening Π^x . (See §23 and the nature of rejection.) But there may be, and inevitably will be, given open-texture, contexts in which neither defence nor rejection of Π^x are possible for the agent.

GE then has no problem with open-texture, basically because it is not attempting to define truth-conditions as part of the way predicates function. Moreover, GE preserves the important compositional idea that there is a general function assigned to negation.

Why is there open-texture? Because any given Π^x in the list is a disposition triggered under some finite sub-set of circumstances. There is simply no point in a creature developing a disposition that will be sensitive to every possible discriminable circumstance. Since few of these might ever be encountered it is a waste of cognitive space to develop dispositions that cater for them. That is why there is open-texture. GE has no problem catering for it, whereas it is a significant issue for standard semantics.

§30 Colour Terms: Representation, Inference, and Function

We have sketched out, from the perspective of the functional view, the language agency underpinning predicates like *red*. But now we face a question about how these structures are related to meanings. Take a group of sighted speakers, $\{U_1, U_2, \dots, U_n\}$ using the term *red* on its colour reading and translations of *red* such as *rouge*, *rot*, and so on. Each of these speakers

instantiates a Γ -state, Γ_x , that underpins the predicate she uses. Let us write these out neutrally as, $\{\Gamma_1, \Gamma_2, \dots, \Gamma_n\}$. All these speakers use terms with the same meaning. A reductionist orientation—not one that is accepted in this book—would be that all these Γ -states, $\Gamma_1, \Gamma_2, \dots, \Gamma_n$, tokened in different speakers using *red* or translations thereof, have the same functional essence. Two candidates for functional essences are: (a) a representational property that links these states to a class of things in the world; (b) a causal pattern governing each Γ_x that underpins inferences using the predicate. Such forms of reductionism would be attempts to find a functional basis in $\Gamma_1, \Gamma_2, \dots, \Gamma_n$ for the fact that U_1, U_2, \dots, U_n all mean the same by their terms. I argue that the reductionism in (a) and (b) does not work, and furthermore that it does not have to work. GE's language-agency analysis does not require that meaning the same by *red* requires identity in the functional character of the Γ -states that underpin uses of *red*. As I argue in §31, neither representation or inferential pattern, or some precursor to it, are needed to explain predicate function. Indeed, they cannot play any such role because talk of representation or inference presupposes the interpretative perspective and thus S-processing.

Representation and Causal Co-variation

The first reductionist option is that the functional essence of D_{Red} is its representational character. It is a state that tracks a certain kind of similarity or uniformity in the world, or fixes a set of red things. Fodor (1987) proposes a theory of information flow, based in the idea that internal mental symbols are activated by the presence of objects with certain properties, such that there is a kind of causal co-variation of the presence of certain worldly characteristics and these mental symbols, and this is the basis of these symbols representing certain properties or classes of objects.⁵⁸ Such approaches have been thoroughly criticised.⁵⁹ I will just give some rough indications of why I don't think that representationalism works.

Let us treat *red* as having its red-surface reading. Then a representationalist might accept:

⁵⁸ Such theories are proposed by a range of philosophers but most notably Fodor (1987) and Dretske (1988).

⁵⁹ See Gauker (2002).

Red-Track: A subject x possess the perceptual concept of *red* if and only if x possesses a disposition D_{Red} , such that where x is in a perceptual situation and directed towards triggering D_{Red} , the latter is triggered if and only if a red surface is present.

This analysis is really open to doubt for familiar reasons:

(a) Speaker can respond *red* under circumstances in which a perceived object is really white but the lighting makes it look red. In short, their cognitive and perceptual resources can let them down. We might hope to rectify this by introducing *normal conditions*, but it is hard to know what normal conditions are;

(b) Speakers can have beliefs that lead them to respond in diverse ways. Thus O_1 and O_2 under identical external conditions can both possess the concept of surface-red but respond differently.⁶⁰ One might be tempted by the idea that we need a theory of looking-red. So what fixes the concept is that under certain conditions, U responds with the mental concept-symbol *looks-red*. That may be the Γ -state corresponding to the pure appearance term: D_{Red} .

Triggering of *looks-red* corresponds to a pure appearance judgement. Speakers go from looking-red states to red-states based on other Π -property states (which correspond to beliefs). The question now is, what is the correlation that holds between triggering of *looks-red* and the world? We have two choices:

(a) U possesses RED iff U tokens *looks-red* before x iff x is red;

(b) U possesses RED iff U tokens *looks-red* before x iff x looks red.

(a) cannot be right, since U will frequently token *looks-red* when it is not red. (b) cannot be right since looking red is not an objective feature of the world. In relation to this last point one might object that looking red can be reduced to some physical specification of lighting conditions. But it is doubtful that all speakers possessing the concept red will respond with the same mental symbol *looks-red* if and only if some precise set of physically specified conditions obtains.

⁶⁰ One might say this: were all users of the concept to have the same set of beliefs they would react in the same way to identical perceptual circumstances. But this assumes an idea of possessing the same beliefs but that judgement presupposes an interpretative stand point, requiring a holistic orientation to the subject.

Inferentialism/Conceptual Role

The second idea was that $\Gamma_1, \Gamma_2, \dots, \Gamma_n$ are all functionally identical in virtue of entering into the same causal patterns underpinning dispositions to inference using the relevant predicates. In other words, $\Gamma_1, \Gamma_2, \dots, \Gamma_n$ enter into the same C-laws (or, strictly speaking, C-law homologues). The problem with this approach is that different speakers rightly interpreted as meaning red by *red*, can differ in relation to the inferences they make with *red*. For example, U might lack the concept of colour or surface, or have a range of theories about conditions under which things look red that makes U differ inferentially from another speaker. Perhaps then we have to restrict the set of inferences that are relevant to concept individuation to some sub-set of the possible inferences speakers are disposed to make. But that is without doubt problematic, for what is the essential concept-fixing set?

We must also allow that the states Γ_x are linked to perceptual circumstances. But now we face the problem of identifying the conditions under which the properties Γ_x are triggered. But for reasons already given above, it is doubtful that there is a set of such conditions.⁶¹

Extrinsic Functional Character

Here is another related way in which we may think of the type-identity of $\Gamma_1, \Gamma_2, \dots, \Gamma_n$. We have been looking at what we might call intrinsic functional character, but what of bringing in extrinsic functional character? We have given a theory of the general functional Γ -state of a particular kind of predicate, states D_{Red} that underpin uses of *red*. This character stamps *red* as predicate of a particular kind: an empirical, perception-based predicate. It identifies the mode of perception as visual and colour-sensitive. This is a kind of structure that will be shared by other colour term Γ -states. So what we have said is not specific to *red*, but to colour terms generally. We can identify the spectrum to which the term is related.

Perhaps we can, for the basic D_{Red} that corresponds to the pure appearance concept, go some way towards a functional essence, though I still think we won't get there. Γ -states for empirical terms mesh with a whole set of C-laws, of both a logical and an empirical kind, that have a bearing on the grounding conditions for the Π -properties underpinning utterance

⁶¹ Brandom (1994) argues for an inferentialist treatment of concepts: they are patterns of inference. Both kinds of problems given in the main text are issues for him.

of *N is red*. Beyond general constraints, there is no predicting the pattern that these C-law relations may take. As we noted in §26, there is an element of holism here.

§31 Two Perspectives and the Non-Objectual View of Predication

GE rejects the reductive approach outlined above. In the case of our speakers, $U_1, U_2 \dots U_n$, all of whom can be said to grasp the concept of red and to use predicates that mean red, there does not have to be a functional essence realised by all their respective Γ -states, $\Gamma_1, \Gamma_2, \dots, \Gamma_n$. GE is not a theory of meaning or concepts, and does not attempt to reduce attribution of identity of meaning to identity of attribution of functional identity. Some terms have functional essences—formal predicates like *true*—but not all do. Empirical terms, as suggested in §26—lack underpinning functional essences.

The judgement that H possesses the phenomenal concept of red, or means red by her term *rot*, is an assertion in which U defends a disposition to L-process H's term using a certain Γ -state. L-processing is a kind of S-processing—as described in §17. S-processing is that activity in which a speaker uses the mental elements of their own NC systems to simulate the signalling acts of others. L-processing is that form of simulation focused on locutionary acts. To L-process *red* as used by U is to simulate U's production of that term through a homologue. To make such attributions speakers must have an intact signalling system, they must have abilities to L-process, and have Γ -states and repertoire dispositions that provide the material for the L-processing. However, we do not define successful L-processing in functional terms.

Being a homologue

From the position of Wissenschaft, looking down upon signalling systems, we describe the speakers $\{U_1, U_2, \dots U_n\}$ as sharing homologue states, $\Gamma_1, \Gamma_2, \dots, \Gamma_n$. The basis of that judgement will be:

(a) There are functional similarities between the states. As we noted, there will be features that indicate they are visual-perceptually based states, that they are surface-based, non-count terms, that they are sensitive to certain points on the spectrum, and so on;

(b) A study of the way in which each state Γ_x is activated by lighting conditions can indicate the spectrum, for the basic red-surface red.

To say that states are homologues is to say that we can predict that, very likely, utterances produced by a homologue in one speaker will be L-processed by a homologue in another. And very likely, speakers will L-process our utterance of *red* in terms of these Γ -states, and we can assign some probability on our L-processing dispositions in relation to their overt verbal activities.

Imagine you are a language agent L-processing another's use of *red* using a state D_{Red} . To the extent that there is functional order of a certain kind, we predict convergence of interpretative opinion. We are not defining interpretative rightness in terms of functional states. As cognitive scientists we make judgements of probability: given functional descriptions of U and H, then very likely U will L-process H's productions using such and such terms in a certain way. Science is concerned with predication, not interpretation.

Representation

At no stage in describing the state D_{Red} have we appealed to representation. The perspective of functional analysis sees signalling systems causally embedded in a broader reality, but does not talk of representation. Talk of representation, talk of truth, talk of applying, presupposes language activity that deploys dispositions to S-process signalling activity—the interpretative stance. But these systems are not utilised within the pure functional stance. We do not see the interaction of the agent and the world in semantic terms. The language agency behind *red* does not involve being dispositionally related to a class of red things, a mereological fusion of redness, a universal of red, or related kinds of entities.⁶²

How does GE escape from the need to use representation in characterising predicates? This is how: firstly, GE gives no theory of meaning identity across speakers. So it does not need say, to introduce the world to define *means-the-same*, as we might do in a representationalist framework. Secondly, GE provides no theory of truth for simple subject-

⁶² One might object that a predicate could have an extension, in that there is a set of things of which it is true. Indeed that might be the case. What GE denies is that the function of the predicate is explained in terms of picking out such an extension.

predicate sentences. The fixing of an extension is not required, because in GE, the function of the truth-predicate is not given in terms of how it contributes to conditions of truth. GE entails no theory of truth in the sense of entailing T-sentences or sentences stating truth-conditions. Rather in GE the function of the truth-predicate is given by a description of the underlying cognitive structures and activities of the language-using agent deploying the predicate—see §17 and §20. Thus U judges true an assertion of *T is red* uttered by H iff U is disposed to defend a state with the structure $\Phi-(T) + D_{Red}$, where (a) U L-processes H's utterance through this state. This description of activity does not entail that the term *red* functions in terms of an extension—U's L-processing is not the assignment of an extension. Her assignment of truth is not based in a judgement that a referring term falls within the extension of a predicate.

This does not mean it is ever wrong to talk of *red's* applying to or being true of red things. Talk of reference disappears in the mode of Wissenschaft—the functional stance—not because such things do not exist, but because we are renouncing, temporally, the mode of speech in which we can talk that way—defence of S- and L-processing dispositions—yet explaining that mode of speech.

§32 Expressivism about Natural Kinds

The speculative cognitive model that we have introduced to analyse the states associated with use of general terms like *red*, and the general reflections on the interpretative and functional perspectives that we have just outlined can be applied to what are typically called *natural kind terms* like *dog* or *water*. One might doubt that expressivism about such terms is a possibility. That is because one might think that expressivism is not a doctrine about meaning or language activity that reaches out to reality itself. It deals, one might think, with subjective, reality-independent states, like affective states, whereas natural kind terms have meanings determined by objective mind-independent features of reality. If O's utterance of *This is water* is made true by the demonstrated sample being H₂O, it must be that the meaning of *water* reaches out to and is fixed by the structure of an external reality. If so, it seems,

representation must play a role in explaining the meaning of such terms, and expressivism will fail. This impression about expressivism arises from a certain distorted concept: thinking that expressivism is a theory of meaning, competing with representationalism about meaning. But expressivism does not compete at the same level with representationalism. GE is consistent with talk, from the interpretative perspective, of terms like *water* corresponding to hidden features of reality. It is just that correspondence plays no role within the account of language agency. If that is so, the quick argument that expressivism cannot deal with natural-kind terms fails.

In fact, what we have said about *red* can be applied with little alteration, except in detail, to the case of natural kind terms. Thus take the term *Dog*, for canine creature. We find the same structure of linguistic division of labour between speakers who possess perception-based states, D_{Dog} —states that underpin their capacities to have perceptual encounters with dogs—and those speakers acquiring their use of *dog* through testimonial interaction with speakers possessing D_{Dog} . The Π -property defended by U, a perceptual speaker, in asserting *N is a dog* is $\Phi-[N] + D_{Dog}$. D_{Dog} is functionally defined in terms of concatenation and basing in the way already illustrated.

The perceptual capacities underpinning D_{Dog} are multi-dimensional, using a range of sensory modalities, and differing across speakers in terms of the modalities used. Blind speakers may use smell, touch, and sound, and so on. But GE rejects the idea that there will be either an inferential or causal-co-variation analysis of the functional essence of D_{Dog} . For those judged to be possessors of the phenomenological concept of dog, there will be no functional essence in their underlying states D_{Dog} that fixes the identity of the concepts they each possess. One might wonder what guides speakers in their interpretative judgements that others mean dog by ‘dog’. That is, what guides their L-processing dispositions—§17? In most cases these are purely contextual or situation based. Encountering U who uses *dog*, and taking U to speak English, and using terms like *animal*, H L-processes U in terms of the Γ -state she associates with one disambiguation of her term *dog*. There is a deeper question of what justifies meaning-attribution. We examine this in chapter 10.

In this account of the language agency—given our retention of the functional perspective—acts of predication using *dog* are not analysed in *objectual terms*—see §3 and §31. We do not analyse acts of predicating *dog* of things by appeal to structures that might enter into an analysis of truth-conditions for simple sentences of the form *T is a dog*. Thus we do not assign correspondence properties to U's acts of deploying *dog*, such as reference to the set of dogs or a property of doghood. Rather our analysis of predication is simply an analysis of the structure of Π -properties that are defended in sentences of the form *T is a dog*. Our analysis of application of the truth-predicate to sentences of the form *T is a dog* has the model given in **ID-True**—§17—and re-emphasised in §20. None of this involves assigning referential properties to the term *dog* or D_{Dog} . This does not rule out, from the interpretative stance, asserting sentences like: '*T is a dog*' is true iff *T instantiates Doghood*, etc. In the functional perspective we analyse the language agency behind such talk of truth and instantiation of properties. But we do so in non-representational terms. Representational language is banned from the functional perspective.

GE does not deny that there are referring terms of the form *dogs*, *the dogs*, etc, that refer to dogs. But it denies that correspondence relations have any role to play in explaining language activity with *dog*. Thus we are not offering a semantics according to which *dogs* would pick out the entities in the extension of *dog*. That would be to embrace a correspondence semantics. GE is not offering a semantics at all. As we shall see in chapter 6 and 7, where referring terms are examined, we offer an expressive treatment of reference compatible with this non-correspondence treatment of language agency of general terms.

This account needs supplementation with an account of property talk, which in the case of terms like *dog* and *water*, is related to talk of *kinds*. That is provided in chapters 6 and 7. We may be further worried about *truth-making*. U's assertion *This liquid is water* is made true by something in the world, namely, that the liquid referred to is H_2O . I provide GE's analysis of facts and truth-making in chapter 8. In short, the correspondence relations of truth-making do not have to be part of the account of the function of natural kind predicates. To allow for the objective truth of claims like *This is water* we do not have to think that

water functions representationally. (The issue of objectivity and representation is examined in some detail in chapter 9, §71.)

Theoretical Identity, and Real and Nominal Essence

The worry about truth-making is related to a worry about respecting the fact that natural kind terms are meant to correspond to hidden essences in reality. The phenomenon of theoretical identity brings out this idea of hidden essence. Consider the theoretical identity *I*:

I: Water is H_2O

The terms *water* and H_2O flanking identity have their referring term function, not their general term-predicate roles. Here they are used to denote stuff. We might think that justice can be done to theoretical identity only if our analysis of the language agency of *water* and H_2O appeals to representational relations. This is not so.

Identity claims for GE—§27—are assertions in which we defend inter-substitutional dispositions. In asserting *I*, U defends a state $IS[water, H_2O]$. This is a disposition to inter-substitute speech-acts U associates with *water* and H_2O in all non-opaque contexts. There is more to theoretical identity than just identity. *I* sums up the theoretical reduction of water to H_2O . We do not say H_2O is *water*. Order is significant. In the context of theoretical claims, it signifies that we allow our use of *water* to be guided by our use of H_2O . If there is a conflict in the deliverances of the perceptual tracking dispositions D_{water} and D_{H_2O} , we opt to be guided by the latter. What we call the nominal essence of water is explained by its real essence. We can extract a set of predicates from D_{water} that capture the nominal essence of water, which is to say, predicates about features like transparency and potability. The real essence of water is extracted from the predicates that are associated with D_{H_2O} . We then have explanatory statements of the form: *That this sample is constituted by H_2O and has such and such a structure explains why it is transparent*. Such explanatory statements may involve reference to events, states of affairs, and facts. But GE is perfectly able to explain the language agency behind deployment of such talk. (See chapters 6-8.) We need to provide a theory of the language agency of *explains*. But that does not require introducing representationalism about kind terms.

Putnam (1975) argued that the structures in the physical environment in which speakers find themselves partly determine the meaning of their terms by fixing the extension of those terms. GE will reject the appeal to extension here, but it will embrace the externalist spirit of the remarks without the metaphysical letter. Once we have accepted the theoretical identity, *I*, we may view others' practice, including others under distinct possible circumstances, in terms of this theoretical reduction. For example, we might take the ancient Greeks to be referring to H₂O, if we discern an added theoretical dimension to their *water*-term use. They use their water-term with the theoretical commitment that water is the underlying stuff explaining water-appearances. That commitment means they have theoretical interests manifested by seeking properties, reference to which, explains manifest features of water, of the kind given above.

Sortal/Mass Term Distinction

Natural kinds divide into *count* and *mass* terms: *hydrogen atom* is count, *gold* is mass. How do we understand that distinction? The difference between count and mass terms lies in the differences in the perceptual sensitivity of predicate precursor states underpinning use of these terms in perceptual speakers. Take the count noun *dog*. Here we need to know how a given dispositional state, D_{dog}, reacts to environmental stimuli. It will differ from the way in which that for *dog-stuff*, D_{dog-stuff}, and that for *dogs*, D_{dogs} react. D_{dog} reacts in an individuating way, with grain, whereas D_{dog-stuff} does not. Thus, in the presence of a dog part, there is no triggering of D_{dog}. In the presence of more than one dog the system D_{dog} no longer triggers, but rather D_{dogs} and D_{dog-stuff} trigger. Division downward of a dog into parts, and upward with the addition of more instances will preserve the activation of D_{dog-stuff}. We can apply Evan's (1975) technique. Facts about predication, say colour terms, will distinguish the terms *dog*, *dogs*, and *dog-stuff*. With a dog half-dipped in chocolate, U accepts *The dog-stuff is chocolate*, and so on. What we find here is a limning of the response profile of certain perceptual dispositional systems. This limning allows us to deduce that *dog* is a sortal, *dogs* a plural sortal, and *dog-stuff* is a mass term.

It is part of our concept of individuals, such as dogs, that they are persisting entities with identity over time. Mastery of *dog* requires an ability to re-identity things as the same

dog. To have such an ability amounts to possessing a set of criteria that generate dispositions to inter-substitute referring expressions of the form *the dog at t₁* and *the dog at t₂*, where t₁ and t₂ are different times. (See the treatment of identity in §27.) The criteria that dispose U to develop such dispositions refer to features of qualitative continuity and causal connectedness. If the dog at t₁ looks like the dog at t₂, shares distinguishing characteristics, and shares certain spatiotemporal conditions, a disposition to inter-substitute *the dog at t₁* and *the dog at t₂* might form.

§33 *Expressivism about Disposition and Cause*

In §26, we provided a table classifying predicates and relational terms according to the kinds of Π -properties and related Γ -states that speakers associate with them. We are moving through this classificatory scheme. We have considered formal relational terms, like *identical*, empirical terms like *red*, and natural kind terms like *dog*. We now look at another set of empirical predicates and terms. The mental elements of these predicates and terms are not based in perception. Rather, they are based in another feature of our cognitive life. This is our agency-grounded manipulation of reality. The terms I have in mind are dispositional terms such as *brave* or *soluble*, and causal terms, like *cause*, as well as thick causal terms such as *break* or *harm*. What I say here is brief and speculative.⁶³ I begin with *cause*, but as we shall see, dispositional and causal cognition are intimately related. So we shall find ourselves explicating both forms of talk in an interconnected way. In this discussion we make free use of causal and dispositional idioms. GE, as a theory of language agency, is essentially a causal-dispositional theory. But, as already argued, there is no circularity arising when we turn this strategy on causal and dispositional language, since—as we saw in §19—GE offers

⁶³ There are those in the literature who have attempted to analyse both dispositions and causation in terms of counterfactuals. See most recently Lewis (1997) and Yablo (2002). If they were right, we could propose an analysis in which asserting *N is G*, for some dispositional predicate *is G*, is to defend a disposition to engage in certain counterfactual reasoning, expressed by *had N been F it would have been G*. Matters would hold similarly for *cause*. But counterfactual analysis is notoriously problematic, and I will not spend time reviewing the potential of this approach here. See Martin (1984). Counterfactual analysis fails miserably for chancy dispositions. See Barker (200+a).

neither truth-conditional or conceptual reduction. There is then neither conceptual nor ontological circularity.

Cause

The expressivist analysis of the verb *cause* invokes the subject's own agency. This does not mean that we truth-conditionally analyse causation in terms of agency. Rather, it means that the states we express in asserting *A causes B* are based in states in which agency plays a part. Agency theories of cause and causal modelling are well known, and what I have to say shall draw on some of this work.⁶⁴ Details will differ but not greatly. What differs fundamentally is the expressivist framework in which the ideas are implemented.

Through interaction with its environment, a person can acquire a set of dispositions with regard to how it manipulates events. The basis form of these dispositions is:

U is disposed to do F in order to try to get G.

That means that U has some practice in which, given a desire to realize the instantiation of *G*, U acts so as to bring about an instance of *F*, and with some high degree of reliability achieves that goal. For U, doing *F* is a *recipe* for inducing *G*. Where U has a recipe disposition in her repertoire of interactions with the world around her, there is generated in U's NC system an internal cognitive map of the pattern of her motor activities combined with capacities that underpin her reference to *F*- and *G*-events. Let us denote this state by *Recipe[F, G]*. This state, in a normally functioning agent in touch with the world, will be induced only with successful interactions with her world in which U has achieved recipes for manipulating the world. But it may also be induced by error. The state *Recipe[F, G]* is then akin to a perceptual state in that like a perceptual state, it can be veridical or non-veridical.

The state *Recipe[F, G]* is the primary ground for the Π -properties underpinning generic causal statements; that is, statements about one kind of event causing another kind. Let D_{Cause} be the Γ -state of the verb *cause*. The relata of the verb *cause* are event terms. Let $\Phi-[F] + D_{\text{Cause}} + \Phi-[G]$ be the Π -property formed by the concatenation of $\Phi-[F]$ and $\Phi-[G]$

⁶⁴ See Gasking (1955), Von Wright (1971), Price (1991, 1992). In relation to modelling there is the work of Pearl (2000), and Woodward (2003). The latter eschew analysis, whilst the former attempt analysis.

with D_{Cause} . Our proposal then is that there is a *basing* C-law of the form below present in the cognitive systems of speakers who have the concept of cause:

$$\text{Recipe}[F,G] \quad \Rightarrow \quad \Phi-[F] + D_{\text{Cause}} + \Phi-[G]$$

This is a C-law linking U's possession of a recipe for getting G by bringing about F to the state $\Phi-[F] + D_{\text{Cause}} + \Phi-[G]$; that is, the cognitive state that is the concatenation of D_{Cause} with the Φ -properties $\Phi-[F]$ and $\Phi-[G]$ of the event terms F and G .

If all U's causal judgements were restricted to claims about events that she is able to manipulate, then her causal judgements would be identical to assertions expressing states of the form $\text{Recipe}[F,G]$. However, many of U's causal statements will be about event kinds, instances of which she cannot manipulate. These statements cannot be defences of the Π -property $\Phi-[F] + D_{\text{Cause}} + \Phi-[G]$ based on $\text{Recipe}[F,G]$. If so, there must be alternative grounds for the production of $\Phi-[F] + D_{\text{Cause}} + \Phi-[G]$. These alternative grounds are the deliverances of analysis of and reasoning about causal processes. Where U instantiates $\text{Recipe}[F,G]$, she may examine the process linking F and G and extrapolate thence to other causal judgements not about F and G . (We shall look at this cognition of causal processing below in some detail.) Thus, information that there are processes that are relevantly structurally similar to processes that U can manipulate is another path through which $\Phi-[F] + D_{\text{Cause}} + \Phi-[G]$ can be induced in U. So there are two kinds of sufficient conditions for $\Phi-[F] + D_{\text{Cause}} + \Phi-[G]$. One is the possession of $\text{Recipe}[F, G]$. The other is the possession of information that a manipulable process is relevantly like that linking F and G .

We have been looking at causal statements that are general such as *Pushing pebbles into sand causes dents*. But causal judgements can be singular: *Tau's pushing the pebble caused the sand to dent*. In the singular case, the ground for the causal judgement is identification of a process that, structurally speaking, is an exemplar of a process type. In both singular and general cases, the relata are events, it is just that in the singular case the relata are concrete events that are instances of event types.⁶⁵

⁶⁵ One-off brute causation is not ruled out. One might have theoretical reasons for applying the term *cause* to singular episodes without a general background pattern.

Dispositional Predicates

As we shall see, an ability to analyse causal processes requires an ability to use dispositional predicates. So, before we get onto the conceptualizing of causal process, we need to look at dispositional predicates. Like causal talk, dispositional predicates have mental elements that have an agency basis, that is, a basis in recipes. Take *water-soluble*. We can imagine that U has a recipe in her repertoire of getting a salty-water solution by placing salt in water. That recipe is both the primary ground for Φ -[*Salt in water*] + D_{Cause} + Φ -[*salty water*], but it is a primary ground for a Π -property of the form Φ -[*Salt*] + $D_{\text{Water-soluble}}$, where $D_{\text{Water-soluble}}$ is the Γ -state of the predicate *water-soluble*.

For other dispositional predicates the relevant manipulation practices are not quite recipes. Take a term like *impenetrable*. The grounding states for $D_{\text{Penetrable}}$ are interactions involving attempts to move one object through another object and *failing*. Thus U could build up a kind of practice of failure in relation to achieving a certain result. Let us call such practices in relation to attempts to get G from F *failures*. Failures generate in U's NC system an internal cognitive map of the pattern of her motor activities in combination with the cognitive systems she needs to refer to F and G. Denote this by *Failure*[*F, G*]. Such states are the primary grounds for the Π -properties underpinning sentences of the form *N is impenetrable*. So where U deploys *N* to refer to object O, and U uses various objects X as potential devices for penetrating O, we have the basing C-law:

$$\text{Failure}[\text{Pushing X, Passing X through O}] \Rightarrow \Phi\text{-}[N] + D_{\text{Penetrable}}$$

Many predicates like *hard*, *inert*, and so on, have mental elements based on failures.

Not all judgements about dispositions are based in primary grounds. Not all judgements about objects having dispositions are about objects that the speaker can manipulate. U may assert that O is penetrable, soft or flexible, simply by looking at O and identifying visually the stuff it is made of. Sand over there is judged penetrable, even though we may not be able to manipulate it. The giant rock beyond is hard and solid, even though we cannot touch it. In short, a Π -property like Φ -[*N*] + $D_{\text{Penetrable}}$ can also be tokened on the basis of Π -properties that underpin judgements about what the object *N* is made of.

Judgements that some entity O is made of the same stuff as an object we can manipulate do not necessarily depend on judging that O has certain dispositional properties. Rather, empirical predicates like *rock* and *sand* are not in themselves dispositional, or at the very least, have significant components that are non-dispositional. They can be applied to bridge the gap between the manipulable and non-manipulable. In other words, crucial to our use of dispositional predicates is that some predicates are not (completely) dispositional. That means they have empirical components that are mental elements based purely in perception.⁶⁶

Causal-Process Cognition

We can now return to the question raised above of a speaker U's ability to cognize causal processes. U needs such cognitive capacities to make causal judgements about events she cannot manipulate. Let us take the case of a speaker U able to express verbally her cognition of causal processes. U can describe processes. U, I submit, does not need to use causal vocabulary that is itself deployed through application of causal cognition. What kinds of terms does the speaker need to be able to conceptualise these processes? Consider an example. By using small stones and pressing sand to form dents, U develops a recipe for denting sand through pressing stones, and as a result instantiates Recipe[Push pebble, Dent sand]. Contemplating her manipulations and elements thereof, the little pebbles, the sand, the pressing down, and so on, U is able to extract information that enables her to judge that a giant bolder falling onto sand caused the sand to dent. What is the nature of the description of the process that enables U to do this? The terms U deploys come in kinds.

U will require: (a) spatial and temporal predicates, *like next to, above, in the next moment, touching*; (b) dispositional predicates; (c) purely empirical predicates or mixed predicates. I suggest that the application of all these predicates is independent of any causal thought that is not based directly on recipes. If this is right we have a system of cognitive tools that provides us with a rich system for identifying processes. It is in terms of this system that we can discern process type-identity, and analogy allows us to extend our causal claims beyond events that are within the sphere of our agency.

⁶⁶ It may be that most empirical predicates like *sand* or *dog* are hybrids. Their Γ -states are complex with perception-based elements and agency-based elements. Some elements will involve a mix of agency and perception. The modality of touch is an example. I press an object and feel pressure.

Agency and Metaphysics

GE does not require that the speaker have a concept of agent or action to have the concept of *cause*. To possess a recipe of a state of the form Recipe[F,G], U does not have to have any concept of trying, agent-causation, etc. Even in being able to discern relevant similarities between processes, the speaker needs no grasp of her own agency as such. Unlike other agency approaches, GE's expressivist treatment does not have to suppose any peculiar account of agency. It is not reducing cause to agency. In fact, agency might be explained in causal terms, though I shall remain neutral here. It may be that O's doing F is not analysable in causal terms. The language agency that underpins talk of agency may not relate directly to the language agency underpinning causal talk.

This analysis evidently needs more development as a theory of the language agency underpinning talk of cause and disposition, but this is not the place to do that. What I am offering is not a complete theory but a sketch of the form of a theory.⁶⁷ The analysis assumes an account of event-talk. I provide an analysis of the language agency of event-talk in §49.

Having provided a theory of language agency for causal-talk and disposition-talk, there is no further question about the nature of causation or dispositions. We saw, for example, that in the case of identity—§27—having provided the language agency for *identical*, there was no question left over about the relation of identity. That is because the language-agency analysis of *identical* does not assign to the mental element underpinning deployment of *identical* any worldly correlate that from the perspective of interpretation could be called *the referent of 'identity'*. It is exactly the same with *cause*. The return of the interpretative stance means we can ask *What is the referent of 'cause'?* or *What do we really talk about when we talk of 'cause'?* But there is no answer to these questions, not because they are unfathomable, but because they presuppose something false. Roughly, this is that the language-agency analysis of *cause* assigns to the mental element underpinning *cause* a worldly correlate that from the perspective of interpretation could be called *the referent of 'cause'*. This does not mean that there are no deep questions to ask about causation, such as,

⁶⁷ There are also thick causal terms, like *break* or *move*. They can be agent verbs, as in *Tim breaks the window*, or event verbs, as in *The falling of the stone broke the window*.

why it has a certain prevailing direction. Answers to these questions do not require having a theory of the nature of causation.

§34 Blueprint for Expressive Analysis

Without doubt, this treatment of predicates is sketchy and incomplete. There are many other domains in which one would need to undertake the hard task of analysing predicates. We shall look at value-predicates in the next chapter. But we can also ask about purely theoretical terms, artifact terms, spatial predicates, temporal predicates, psychological predicates, mathematical predicates, predicates associated with modality, and so on. Global expressivism is an extremely ambitious project. I cannot offer a fully worked out system in this book. My aim is rather to describe a framework in which the idea can really live as a possibility. That framework provides us with a kind of recipe for generating expressive analyses, that is, language-agency analyses for cases beyond those that we have examined. As such we have a kind of blueprint for application in some domain of talk, which is the following:

- (i) Find what are broadly speaking the epistemic sources for application of the predicates or particles in the domain concerned;
- (ii) Treat these as primary grounds for Γ -states in terms of the theory of basing described above in §16, §20, and so on;
- (iii) Allow for linguistic division of labour;
- (iv) Offer a non-objectual functional analysis of predication.⁶⁸

So, for example, spatial, temporal, and physical magnitude predicates involve measurement based in various kinds of perception, and so why cannot these mental states underpinning measurement be the basis for Π -properties for sentences about space, time or mass?

⁶⁸ As a methodology one can see, as we have already remarked, parallels with Dummett's verificationism and with inferentialism. The difference is that these are theories of meaning, but GE is not a theory. Dummett's reduction of content to conditions of verification generates anti-realism. Inferentialism as we have seen fails in its foundations—see §10.

As we have noted, a central feature of GE is that the analysis of the language agency of predicates does not require that extensions are assigned to predicates. There is no such thing as interpretation in the sense of model theory. So truth is not defined by reference to a model. That is fortunate, since there are well-known paradoxes that attend any model-theoretic conception of interpretation—see Putnam (1983) and Gauker (2002). A standard view is that content should be cashed out in terms of representation. Let the bearers of content be internal mental representations we can think of as mentalese sentences. These sentences represent possible states of the world. Take the true beliefs of the agent, which will be some sub-set of these sentences that are in the speaker's belief box. Then these sentences will have a model. That is an interpretation, in the model-theoretic sense, assigning sets to the predicates and objects to referring terms that makes these mentalese sentences come out true. Such a model meeting certain constraints fixes the mental, representational content of the agent's mental states.

The problem is that there is no plausible set of constraints that will fix any plausible set of models—see Gauker (2002). I will not go through the arguments here. That is because, fortunately, GE is not presented with this problem. GE does not provide any theory of content in terms of interpretations in the model theoretic sense. Interpretation just does not work that way. One cannot emphasise too much how alien the model-theoretic idea of interpretation is to the spirit of global expressivism. The model-theoretic problem of interpretation is a strange pseudo-problem generated by a misapplication of the concepts of reference and representation—their illicit importation into the domain of explanation of language function. The semantic vocabulary is based in the cognitive foundation of simulation—see §1-2. In the functional stance such cognitive functions are excluded.

What we have said about model-theoretical paradox carries over to Quinean inscrutability of reference (1960). Is the speaker's term *gavagai* to be translated as *rabbit* or *undetached rabbit part*? Translation does not depend on assigning sets of entities to

predicates. If that is right, the whole issue of inscrutability of reference is an artefact of a representationalist framework.⁶⁹

⁶⁹ This is not to rule out under-determination of theory by evidence. That is a general matter that applies to any scientific theory construction and testing, and not a specific matter about meaning.

Chapter 5

Expressivism about Everything: Value

§35 The Return of Value Expressivism

We now return to the heartland of expressivism: value. We return transformed, for now we are no longer looking for a theory of the meaning of value-sentences. We are looking for a theory of language agency for value-sentences. Expressivism, as we noted §7, is often presented as a denial that there are facts of value, that value-sentences are truth-apt, that there are any beliefs to do with value, that value-statements are reports about reality, that there are properties of value. But this is wrong. Expressivism is not essentially a denial of these obvious facts about value-sentences. Expressivism can maintain that value-sentences are truth-apt, can be produced in real assertions, manifesting real beliefs, and that there are facts of value making true value-sentences. That is because expressivism, properly seen, is a purely explanatory thesis. That thesis is summed up in *Expressivism*—§7. This is the thesis that in explaining how sentences in an expressive domain *D* function, we do not bring in representational relations between sentences or mental states and reality. The particular mental elements in play in the case of value-sentences are based in motivation and affective states.⁷⁰ This does not imply that sentences lack representational properties. Just that representation-talk—being a form of speech depending on language agency modes not available in the functional perspective—has no role in explaining patterns of language agency. We shall see that this is perfectly consistent with value-statements corresponding to facts and being made true by reality, with there being moral properties, moral beliefs, and moral truth. Unlike the claims of minimalism—§8—none of this involves any idea of a

⁷⁰ Though GE does not require this as we shall see. There is a form of expressivism about value that rejects the view that value-sentences express affective states, but might express perceptual states of some kind.

distinction between insubstantial fact, beliefs, or properties. Indeed, we can assert there are substantial properties of value, such as goodness, it is just that such properties, as with all properties, lack a metaphysical nature.⁷¹

The fundamental error of ethical expressivists has been to treat expressivism as a theory of meaning. That's why they have attempted to force desires into playing the role of propositions, so as to devise theories of logical validity and compatibility in terms of relations between desires. Such an approach must fail. GE's position is not simply that there is no theory of meaning for value-sentences, but no theory of meaning *tout court*.

Below we examine value-predicates such as *is good*, *is right*, *is tasty*, and so on. We then look at thick evaluative predicates like *nerd* or *courageous*. After that we consider modals, such as *ought*, *must* and *may*. After that we examine issues of truth-evaluation, negation, belief-ascription, embedding, and validity, and then examine some metaethical puzzles. The application of GE to value-sentences is essentially routine.

§36 Value-Predicates

The orthodox view in expressivism is that what is expressed is a motivational state. GE does not quite embrace this idea. Just as we saw empirical predicates as being *based* in perceptual states, we can say something similar in the case of value-predicates. The Π -properties of value-sentences are identical to or based in affective states of various kinds. The following are Π -property specifications for a range of value-sentences:

Value-Predicate	Γ -Property Identical to or based in:
<i>T is tasty</i>	Gustatory preference for T.
<i>T is funny</i>	Having a state of mirth induced by the actions, etc T.
<i>T is interesting</i>	Being disposed to find intellectual satisfaction in engaging with T.
<i>M-ing is right</i>	Being disposed to seriously approve of M-ing and being disposed to seriously disapprove of omissions of M-ing.
<i>M-ing is wrong</i>	Being disposed to seriously disapprove of M-ing and being disposed to seriously approve of omissions of M-ing.

⁷¹ GE is not a form of quasi-realism, Blackburn (1993), or a form of irrealism, Skorupski (1999).

Generally, where V is a value-predicate its Γ -property is a state D_V that is functionally defined through *basing*—§16 and §20—in the familiar way. D_V combines with the mental element of a referring term, a Φ -property $\Phi-[N]$ to form $\Phi-[N] + D_V$. Such states have, through the action of C-laws, primary grounds that are affective states. Let X be an affective state type. Let U pick out O s by the general term N . Basing laws have the form:

$$\{\text{Being in affective state } X \text{ in relation to } O\text{s}\} \Rightarrow \Phi-[N] + D_V$$

Thus basing C-laws for the predicates *tasty* and *wrong* are:

$$\{\text{Being disposed to eat with relish things } O\text{s}\} \Rightarrow \Phi-[N] + D_{\text{Tasty}}$$

$$\{\text{Being disposed to disapprove of } O\text{s}\} \Rightarrow \Phi-[N] + D_{\text{Wrong}}$$

Given that the mental element D_V has this functional character, $\Phi-[N] + D_V$ can either be instantiated through its primary ground or through other avenues. These alternative epistemic routes will be paths of inference grounded in, say, testimony. Thus, I may assert *Haggis is tasty* even though I have never tasted haggis. I do so because I have read descriptions of haggis and its mode of preparation. Or I have been told that Haggis is tasty. (Compare earlier discussions of basing in §16, §20, and §21.)

One might object that the abstraction of basing is a step back from the spirit of value expressivism. It is and it is not. The states expressed in value-assertions, according to GE, are affective states or states based in such states, so they may fail to be affective states as such. Yet, the Π -properties of value-sentences are never *cognitive* states in the sense of being belief states. That fits in precisely with the general spirit of value expressivism.

The structure of truth-evaluation of value-sentences has the general interpretative form given by **ID-True**—§17. In the case of statements about haggis's tastiness we find:

ID-Tasty: H judges true an assertion of *Haggis is tasty* uttered by U if and only if H sincerely defends $\Phi-[haggis] + D_{\text{Tasty}}$, the state that H deploys to L-process U's sentence.

We defend taste-based states only by invoking other taste-based states. But that just means that gustatory value cannot be grounded in anything that does not invoke gustatory value.

Likewise for the moral utterance about wrongness:

ID-Wrong: H judges true U's utterances of *Meat eating is wrong* if and only if H sincerely defends Φ -[*meat eating*] + D_{Wrong} , the state H deploys to L-process U's sentence.

Structurally speaking statements of taste and wrongness are identical—the difference is in the kind of affective states upon which they are based.

As with all analysis of language function, we are concerned with the structure of the individual's perspective—there is no God's eye view of meaning or truth-conditions, just a God's eye view of the language agent engaged in the activity that we call truth-evaluation and interpretation. Consider the dialogue between Smith and Jones:

Smith: Eating meat is wrong.

Jones: True. Eating meat is wrong.

The respective states that they defend are based in disapproval states about meat-eating. Our judgement that Smith and Jones talk of the same thing is interpretative. We interpret both Smith and Jones as thinking about a kind of event: meat-eating. As we noted in §26, empirical predicates do not have functional essences. Smith and Jones have states of the form $D_{\text{Eat-Meat}}$ that underpin their uses of the term *meat-eating*. But $D_{\text{Eat-Meat}}$ in Smith may not be functionally identical to $D_{\text{Eat-Meat}}$ in Jones. The interpretative judgement that both are thinking about meat-eating is not fixed by some functional essence that both instantiate. What then about disapproval states? Is there a functional essence to disapproval?

A property disapproval state—one might also have object-disapproval states—is an emotional response to property instances. We shall give an analysis of property and property instance-talk in chapters 6 and 7. The disapproval state is analysed in these terms. An emotional response to property instances of meat-eating is an emotional response caused by perceptual states in which $D_{\text{Eat-Meat}}$ is triggered. That response is a feeling of a certain kind. How we analyse the functional character of feeling is a substantive issue that will not be

examined here, but basically a rough idea is: (a) U is strongly inclined to desist in perception of the property instantiation; (b) U is disposed to stop the worldly correlates being instantiated. There are many patterns of response that might fall into this format, but a basic dispositional pattern is present. If so, we might say that there is a functional essence here. I do not expect that all kinds of disapproval has this form.

Let us represent the cognitive structure of such states as: [Disapproval: $D_{\text{Eat-Meat}}$]. In their dialogue Smith and Jones token states which have this form. Overall their states may not have a functional identity since they contain components of the form $D_{\text{Eat-Meat}}$, and they may differ in functional detail. A third party interpreting Smith and Jones, nevertheless, can judge that they do not speak at cross-purposes. That is not a judgement about their inner states. Rather it is a simulation-based judgement. The third party will L-process both their utterances using the same Π -property state, again of the form: [Disapproval: $D_{\text{Eat-Meat}}$].

§37 Disagreement, Negation, Irrationality, and Rejection

Falsity-attributions introduce issues that are normally thought as deeply problematic for value-expressivism. But nothing is problematic here in GE's terms, since, effectively, our analysis of negation and falsity in non-evaluative cases takes care of the issue. Take a dialogue:

Smith: Eating meat is wrong.

Jones: That is false. Eating meat isn't wrong.

In this case, Jones's response, given the analysis of negation offered in §23, is one in which Jones defends rejection of a Π -property based in an attitude. I defined *rejection* in these terms. U rejects a Π -property, Π , if and only if: U's NC system tokens a property Ψ such that $\Psi \Rightarrow \neg D\text{-token } \Pi$. Roughly, rejection of Π is to be constrained, on the basis of a Π -property Ψ , from tokening Π . Any kind of Π -property can be rejected; that is, any state, based in perception, metalinguistic dispositions, processing of other Π -properties, or as proposed here, affective states. Affective states, like desire and approval can be rejected since there are C-

law relations that hold between pre-doxastic states. Nothing distinguishes affective from non-affective states in this respect.

A property-disapproval state [Disapproval: $D_{\text{Eat-Meat}}$] is a disposition to respond to property instances. Suppose Jones rejects this state. Then she must token some state Ψ such that $\Psi \Rightarrow \neg D\text{-token} [\text{Disapproval: } D_{\text{Eat-Meat}}]$. What is the state Ψ ? It could be that Jones approves of meat-eating for sustenance and pleasure. In this case there is an exclusion law—in the sense of §23—of the form:

$$[\text{Approval: } D_{\text{Eat-Meat}}] \Rightarrow \neg D\text{-token} [\text{Disapproval: } D_{\text{Eat-Meat}}].$$

This exclusion law has the same status as exclusion laws of the form $\Pi\text{-}[T \text{ is blue}] \Rightarrow \neg D\text{-}\Pi\text{-}[T \text{ is red}]$. It is grounded in the cognitive incompatibility of [Approval: $D_{\text{Eat-Meat}}$] and [Disapproval: $D_{\text{Eat-Meat}}$].

Alternatively, it could simply be that she discerns no morally objectionable property in meat-eating, finding it to be neutral with the same status as scratching one's head. In this case, Jones mentally inspects the property of meat-eating and her disapproval system does not trigger. This is analogous to U's searching the room for a cat, finding her dispositional state, D_{Cat} does not trigger—see §23. That non-triggering in the cat case grounds assertion of *There does not appear to be a cat*. In the present case the non-triggering of disapproval—represent it by *Search [No-Disapproval-Meat-Eating]*—grounds assertion of *There appears to be nothing wrong with meat-eating*. We have then the C-law:

$$\{\text{Search [No-Disapproval-Meat-Eating], } \Pi\text{-[No-defeaters]}\} \Rightarrow \neg D\text{-token} [\text{Disapproval: } D_{\text{Eat-Meat}}].$$

The C-law here has an inductive character—§22. It corresponds to an inference from appearance to reality, the kind of inference we make all the time. As with all such inferences we need a no-defeaters clause.

Incompatibility

Let us be clear about what GE is offering. GE supplies a theory of the language agency of disputing parties. Smith's *Meat-eating is wrong* is logically incompatible with

Jones's *Meat-eating is not wrong*. But GE is not providing a theory of what logical incompatibility is constituted by. It is certainly not constituted by the incompatibility of attitudes. Just as, for GE, there is no theory of validity or entailment—§25—there is likewise no theory of incompatibility, either in the value or non-value case. So GE makes no attempt to reduce facts of logic to facts of rationality or conflict of attitude. No such reduction is possible. Yet this is the error of many forms of value-expressivism, in particular Blackburn (1984, 1998), who seeks to reduce logical incompatibility of value-sentence to conflict of attitude. The whole problem with expressivism in orthodox treatments is that it attempts to treat expressivism in the mould of orthodox Semantics—§1. It seeks to provide a theory of validity and entailment. This is a recipe for failure.

GE does provide a functional analysis of the necessary conditions for an agent to have: (i) conflicting attitudes, and (ii) logically compatible beliefs. The structures required for conflicting attitudes, in the case of disapproval of eating meat, have the form *AI*:

AI: U's NC system instantiates Π -properties and C-laws that can generate tokening of [Disapproval: $D_{\text{Eat-Meat}}$] and rejection of [Disapproval: $D_{\text{Eat-Meat}}$].

The functional backbone for being in a state of logical incompatibility of belief is:

LI: U is disposed to defend both [Disapproval: $D_{\text{Eat-Meat}}$] and rejection of [Disapproval: $D_{\text{Eat-Meat}}$].

A speaker can find herself in a state of the kind *AI* but not progress to a state of the form *LI*. U may simply lack dispositions to defend the Π -properties she possesses; which is to say, she may put assertive production in abeyance for Π -properties in the relevant domain. (I discuss this phenomenon in §68.) But typically U will progress from *AI* to *LI*.

States of the form *AI* and *LI* provide the functional structure of states of attitude conflict and logical incompatibility. But GE does not give a theory of attitude conflict and logically incompatible beliefs, and certainly not a theory reducing the latter to the former. Just as it only gives a functional analysis of the skeletal structure of asserting and believing, it only gives an analysis of the skeletal structure of attitude conflict and logically incompatible

belief. That is because the claim that U has conflicting attitudes or logically incompatible beliefs is an interpretative claim. It presupposes exercise of the interpreter's simulative systems of S-processing. Thus **GE** does not give a reductive truth-conditional analysis of:

Smith's statement is logically incompatible with Jones'

Rather, GE provides a functional analysis of the act of asserting such a sentence. This will be an assertion in which a speaker H defends the state:

Where H L-processes Smith's and Jones's utterances in terms of the states Π_1 and Π_2 respectively, there is some Ω , such that H tokens the C-laws (i) and (ii) below:

(i) $\{\Pi_1, \Pi_2\} \Rightarrow \Omega$ (ii) $\{\Pi_1, \Pi_2\} \Rightarrow \neg D\text{-token } \Omega$.

GE explains our disposition to assert that Jones and Smith say things that are logically incompatible with each other. But it does not analyse that incompatibility.

§38 *Thick Value Terms: Value Content and Implicature*

There are other kinds of value-terms than value-predicates like *wrong*. There are thick evaluative terms, like *courageous* or *nerd*. These sentences carry some descriptive information about the kind of features of which there is disapproval or approval. In the case of *nerd* the descriptive features are *intellectual, uninterested in fashion or sport*. So, *T is a nerd* has a dual content—a descriptive component—which we assume to be some set of natural features—and a value-component—which, assuming expressivism, is affective. One idea is that sentences like *T is a nerd* are conjunctive: they comprise an assertion that T has F, the natural feature, and an evaluative assertion in which disapproval of F-things is defended. The problem here is that such sentences don't seem to be conjunctions. One can say *T is a nerd* is false because one thinks that T really is a socially adept sports star, but the following seems odd: *It is false that T is a nerd, being badly dressed and intellectual is OK*. Rather, if one disagrees with the attitudes expressed by utterance of *T is a nerd*, one does not pronounce the assertion false. Indeed, one queries the framework of attitude in which it is made.

An alternative view about the value-component of *T is nerd* that explains this response is that it is *conventional implicature*. Conventional implicatures are introduced by operators like *even* and *but*, and contribute non-truth-conditional content to sentences. A sentence like *Even Granny got drunk* carries a said-content component—that Granny got drunk—and an implicature component—a subjective comparative probability of Granny’s getting drunk introduced by *even*. If one cannot accept the implicature component of the sentence, the sentence is not judged false. If one thinks Granny is a huge drunkard, and so expect Granny to get drunk on every occasion, one does not judge *Even Granny got drunk* false. Rather, we are simply inclined to say something like: *What do you mean even Granny got drunk: she drinks like a fish*—see Barker (2003, 2007). In performing an implicative speech-act, U simply signals her possession of a Π -property. She does not defend that Π -property. This is why conventional implicature is not truth-conditional.

We propose that a sentence like *T is a nerd* has a similar form. It has a truth-conditional content that *T is F*, where F is *intellectual and uninterested in fashion*, but also implicates the speaker’s attitude to these descriptive features: that they are disapproved of. Thus utterance of *T is a nerd* has the structure:

Assertive Π -property:	$\Phi-(T) + \Gamma-[F]$
Implicature Π -property:	[Disapproval of F-things].

Implicature content is presuppositional. When H S-processes an utterance carrying implicature, the perlocutionary effect comprises instantiating the homologue of a state in U, which, because no defensive stance is taken towards it, is marked as uncontroversial, and so something that can be taken for granted. This leads to a kind of *objectivity effect*. In uttering, say, *T is a Kraut*, I implicate that I dislike Germans, but also convey that my audience shares, or finds uncontroversial, this attitude. My audience may not, and I may realise this, but I can still use the term as a rhetorical ploy. I present the attitude as something they accept or can accept without argument. I attempt thereby to pre-empt dispute. If social conditions are right, the fact that F-attitudes are marked as uncontroversial or shared may effect considerable non-rational influence over audiences.

Acceptance and Rejection

Implicature-bearing sentences are not only liable to truth-evaluation but also *felicity* evaluation. In judging an assertion of *Tim is a nerd* felicitous, H S-processes the utterance, simulating a structure like that which underpins U's production of the utterance, and is disposed to defend a homologue of U's assertoric Π -property, and signal a homologue of the implicature Π -property. On the other hand, having S-processed U's utterance, H may reject it. That means, she rejects either a homologue of U's assertive Π -property or a homologue of U's implicative Π -property. Negation, in other words, comes in two forms. Perhaps this is best seen with racial epithets. U may say *Tim is a kraut*. H may reply. *Tim is not a kraut. He is a German*. In this case, U rejects the attitude. Made explicit it would be: *Tim is not to be disapproved of because he is German*. The predicate, *to be disapproved of because he is German*, is a kind of assertoric counterpart to *kraut*. Assertions using that predicate convey in the form of truth-conditional, said-content what assertions using *kraut* convey in terms of implicature. That is one kind of negation. The other is straightforwardly descriptive: *T is not a kraut: he's a Yank*. In this case, H rejects the homologue of U's assertive Π -property.

The Implicature Model: good, right

We might propose that a thin value-predicate *good* has the simple speech-act structure given above for *wrong*. But now we have an alternative treatment according to which sentences like *T is good* have a dual-content structure of the kind just described. *Good* can be used adjectivally, as in *a good knife* or *a good person*. In each case, different properties *F* are being approved of. In the case of *good knife*, features like hardness and sharpness in knives are approved of. That might suggest that in uttering *T is good*, U reports something about T possessing some natural property *F*, but also implicates something about her attitudes to things possessing *F*. This is the proposal made in Barker (2000). Thus *T is good* has the structure:

Assertive Π -property: $\Phi-(T) + \Gamma-[F]$

Implicature Π -property: approval of *F*-things.

The descriptive predicate *F* that a particular speaker may be appealing to in production of a sentence *T is good* is the same property about which U implicates that she has an attitude. There is no constraint upon what *F* is beyond its fitting into U's moral perspective and its being a natural property. It is not built into the account that there is any natural property that all speakers are picking out. Nor is it assumed that ideally rational speakers will pick out the same property. I show below how this can be reconciled with moral agreement and disagreement.

Agreement and Disagreement

Agreement and disagreement, truth-evaluation and falsity evaluation have the same basic structure that we considered above for thick evaluative terms, but there is a further twist. There are two components to the procedure in terms of which U comes to assign a descriptive content to U's utterance of *good*:

F: *Good*, as uttered by U, is associated with *F* such that:

- a) U is committed to (moral) approval of *F*-things;
- b) This *F*-attitude is shared by or uncontroversial for audience/interpreter.

F a)'s source relates to what one might call the *character* of the word *good*. **b)**'s source is pragmatic to the extent that it arises from the presuppositional aspect of implicature. Given the disparate sources of the constraints on interpretively fixing *F*, there are complexities in the general dynamics of content fixation for an utterance of *T is good*. In cases where interpreter and interpretee are in agreement about moral attitude, then the deliverances of **F a)** and **b)** will agree. So take:

Smith: Giving away 20% of your income is good

Jones: True. Giving away 20% of your income is good.

In this case, Jones accepts Smith's utterance. That means U accepts the attitude to *F*—in this case maximises general happiness.

Where there is disagreement in attitude the results lead to an interpretative instability rather than failure of any determinate interpretation at all. Say racist Norm gestures at SS officer Schmidt in a film and claims **G** below. We have two possible reactions, E₁ and E₂:

- G**: Schmidt was good.
 E₁ What Norm said was true.
 E₂ What Norm said was false.

If, in interpreting **G**, we focus on **F a)**—the character of *good* corresponding to a function from speaker F-attitudes to an F-property—then Norm denotes by *good* the property *racist* and Schmidt had that property. So **G** and E₁ are true. But in privileging **F a)** to fix F in this way, we interpret Norm's utterance as signalling the uncontentiousness of the approval of racism, since F-attitudes are meant to be shared. If so, although the truth-conditions of **G** on this interpretation hold, it is infelicitous or incorrect, for us, since there is implicature failure. Now as a general rule, where a sentence *S*'s truth-conditions hold, but *S* is defective because its implicature-condition fails, we are disinclined to say '*S* is true', since asserting that '*S*' is true *conversationally* implicates that we accept that *S*, so interpreted, has passed all the tests on correctness. Although **G**'s truth-conditions hold, we are disinclined to assert E₁ for that would conversationally implicate that **G** is correct, and thus that we share Norm's F-attitude. At best, we could say: *Schmidt was 'good' for Norm*, using scare quotes to cancel the unwanted conversational implicature of shared attitude.

There is, however, another way of interpreting **G**. In interpreting **G**, we may focus on **F b)**, and reason that as Norm is signalling the F-attitude as shared by or uncontroversial for audience/interpreter, we can assign one of our own F-attitudes to Norm's utterance. But then the F-property which is the denotation of Norm's use of 'good' in **G** is an F-property we take to be defining the good. In which case, what Norm says is false, since Schmidt lacks that property. If so E₂ is uncontentiously assertable for us, which, in fact, it is.

Summary

Implicature can play a role in evaluative-sentences, for thick evaluative terms, and possibly thin ones. In the case of *good*, should we go for the simple theory or the dual-content

theory? I am not sure. The simple theory is simpler. In asserting *T is good*, U defends some specific kind of approval state of the thing O picked out by *T*. The kind of approval may be partly identified by some set of properties, whose instantiation U approves of. So, in asserting *T is good* U will communicate that she believes that O has certain properties F. So, like the implicature theory, there will be a kind of dual content. It may be that there is not much between the theories, but we haven't space to examine these issues here. We are displaying possibilities.

§39 Value and Epistemic Modals

We need to move onto another kind of value sentence. There are value-sentences that employ modal auxiliaries such as *ought*, *must*, and *may*, as in *T ought to Φ*, *T may Φ*, and *T must Φ*. Modals have two kinds of reading: value and epistemic. In what follows we look at value-readings—I briefly examine epistemic readings in §70. The basic Π-property specifications for modals on a value-reading are given below. The kind of modal in play depends on what kinds of attitudes of approval and disapproval are being defended:

Value Modal	Π-Property Specification
<i>O must F</i>	Disposition to approve of O's F-ing and strongly disapprove of O's not F-ing.
<i>O ought to F</i>	Disposition to approve of O's F-ing and disapprove of O's not F-ing.
<i>O may F</i>	Rejection of disapproval of O's not F-ing.

Moral *oughts* (and *musts* or *mays*) concern attitudes about conduct directed towards others regarding their well-being. What we might call rational *oughts* are to do with attitudes to agents as reasoners or rational beings. As someone who evaluates rational activity, U will approve of agents with capacities to draw logical inferences within the cognitive resources of the agent. U approves of speakers who can draw simple logical consequences from their beliefs, disapproving of speakers who do not do this. In saying that O ought to believe *Q*, given she believes *P* and *if P, Q*, U defends approval and disapproval states of these kinds in relation to O inferring *Q*. Similarly, U approves of speakers who monitor their own

production of assertions. Thus if O asserts *P*, O ought to recognise that she has asserted *P*.

Here there is no question that O asserts *P* entails that O recognises that she does.

Nevertheless, the *ought*-statement is an expression of approval that speakers have monitoring capacities in simple cases.

Issues of truth- and falsity-evaluations, negation, agreement and disagreement, and so on, need not be examined here. The results of prior discussions can be applied.

Ought, Must, Right and Good

It may seem that the treatment of *right* and the treatment of *ought* are almost identical. But there is an argument that they cannot be the same. Moral reasoners can explain their commitment to assertion of *We ought to M* by appeal to their belief that their doing M is right. On the other hand, we do not explain the fact that our *M*-ing is right, by appeal to the prior fact that we ought to *M*. *Ought*'s flow from prior judgements of right, rather than the other way round. This idea is captured in *because*-statements:

O ought to M because M-ing is the right thing to do.

M-ing is the right thing to do because M ought to do F.

If *ought*- and *right*-statements are semantically identical, it might appear that this is inexplicable.

The right response to the problem seems just to be this: A statement of the form, *It is right to M* expresses approval *in general* of M-ing. *O ought to M* expresses approval of O's *M*-ing. The general approval state can ground the particular, but the particular not the general. There is then no need to change the theory of *ought* and *right*.

§40 *Propositional Attitudes, Embedding, and Validity*

We can have beliefs that are about value. GE allows that beliefs about value are just as substantial as any other belief. There is no essential difference *qua* belief states. A belief state is a disposition to defend a Π -property really possessed. That Π -property can be an affective or motivational state. As we shall see, there is no question of an analysis of *the moral*

proposition. That is because there is no question of analysis of any proposition. I argue in §53 that propositions have no natures, whether they are moral or not.

In uttering a *that*-clause, *that P*, in an evaluative belief attribution, as with any belief attribution, U performs an embedded proto-assertion with the embedded sentence. We have already considered belief-attribution in the case of identity sentences—§27. Everything said there applies more or less to this case.

GE's form of expressivism, evidently, rates as a kind of cognitivism. But again, this is perfectly consistent with the essence of expressivism, which I have presented as an explanatory thesis summed up in §7 with *Expressivism*. It is the view that assertions of sentences in the expressive domain *D* are to be analysed in terms of pre-doxastic mental states, without appeal to representational content. But that is consistent with their being truth-apt and capable of manifesting belief, according to GE.

Value and Motivation

For any affective state there is, potentially, a corresponding belief. In other words, where U has a motivational or affective state, and U is disposed to defend that state, then U has met the structurally necessary conditions for having an evaluative belief. It is not sufficient because generally functional states are not sufficient for the fixing of contentful states. We can see the phenomenal correlate of this in action with the way in which attitudinal verbs are depersonalised. Thus the first-person sentences—*I desire that T Ms*, *I prefer that T Ms*, and *I accept that T Ms*, *I like T*, *I despise T*—can all be transformed into the evaluative impersonal sentences—*It is desirable that T Ms*, *It is preferable that T Ms*, *It is acceptable that T Ms*, *T is likeable*, *T is despicable*. In asserting the impersonal sentences we defend, and express, the states that are reported through assertion of the personal sentences.

It does not follow from the fact that one has an affective state that one has the corresponding value-belief. That is because one may not be disposed to defend the affective state. Say U disapproves deeply of sexual acts of a certain kind. Her reason for her disapproval may lay in approval of sex playing a certain role in people's lives. But she recognises that her disapproval is not shared by many individuals, and she has no overwhelming disapproval of these people. Although she cannot abide the thought of these

acts, she does not condemn them, even privately. There are moments in which, in a state of clear reflection, U withdraws from defending her attitudes, either publicly or privately, although there may be less reflective moments in which she does assert that the acts are wrong. Does she believe they are wrong? U is undecided; we might even say, fluctuating. But it is not her disapproval state that fluctuates. What fluctuates is her preparedness to defend the state, even privately to herself. I say more about this in phenomenon in chapter 9, §68. Having a value-belief is not to be identified with possessing an affective state. Possessing an affective state is only part of what it is to possess such a belief.

Motivational Internalism

Related to these questions is the issue of motivational internalism. Many people think that expressivism is committed to a form of motivational internalism—see Smith (1994). I take motivational internalism to imply that if U really believes that doing F is good, and U is perfectly clear-headed, then U has some degree of motivation to ensure that F gets done. Is expressivism committed to internalism? A cursory glance at GE might incline one to think it condones this idea. But it does not. We said above in §36 with the example of assertions of taste, that one can learn by testimony that haggis is tasty, and so, believe that it is tasty, even though one lacks a gustatory preference state, since one has never tasted haggis. Something similar can happen in the moral domain. The Π -properties of moral-sentences, like those of taste-sentences, are *based* in affective states. They are not identical to affective states. It is possible for a speaker to gain moral knowledge by testimony, but lack the corresponding affective states. Thus it is conceivable that a psychopathic individual asserts sincerely that what he is doing is bad but he has no motivation not to do it.

Embedding in Logical Compounds

Value-sentences embed and enter into logical arguments. Dealing with such facts about value -sentences is typically seen as a serious challenge to the viability of expressivism. GE's version, as one should perhaps suspect at this stage, has no difficulties. There is no special problem of embedding and validity for the case of value-sentences. Take a conditional like *E*:

E : If Fred is a liar, he ought to be punished.

In *E*, U performs two embedded proto-assertions with *Fred is a liar* and *he ought to be punished*. This means that their utterance is caused by repertoire dispositions—see §14— which in the case of the second sentence has the form below:

If U is directed towards a ground-indicating disposition in relation to approval of Fred's being punished, U may utter *he ought to be punished* (with anaphoric reference to Fred)

When embedded, the ground-indicating dispositions for the two sentences are not causally efficacious, and U is not signalling them. Which is to say, U makes no assertions with the sentences. (See the discussion of embedding and play in §15.) We shall not discuss the structure of conditionals in detail here—I say more in §56. We just need to note that what backs up the conditional *E* is a C-law:

$$\Pi\text{-[Fred is a liar], } \Pi\text{-[Every liar ought to be punished]} \Rightarrow \Pi\text{-[Fred ought to be punished]}$$

In this case, the crucial hidden Π -property is that underpinning *every liar ought to be punished*. I examine universal claims of this kind in §56. But for now, crudely, in making the generic assertion about liars, U defends a disposition to token and defend $\Pi\text{-[}T\text{ ought to be punished]}$, for arbitrary *T*, given that she tokens and defends $\Pi\text{-[}T\text{ is a liar]}$. The basis for such a disposition is just U's generic disapproval state of liars.

Validity

Our talk of the validity of arguments like *A* below is perfectly explicable:

A : If Fred is a liar, he ought to be punished. Fred is a liar. Therefore, he ought to be punished.

In asserting that an argument is valid, U defends a state based in C-laws—§25. So, the basis for asserting that *A* is valid can be instantiation of the C-law:

$\{\Pi\text{-}[\text{Fred is a liar}], \{\Pi\text{-}[\text{Fred is a liar}] \Rightarrow \Pi\text{-}[\text{Fred ought to be punished}]\}\} \Rightarrow \Pi\text{-}[\text{Fred ought to be punished}]$

There is then no problem whatsoever in explaining talk of validity. Of course, we are making no attempt to analyse validity. Certainly, unlike other approaches to expressivism, GE does not attempt to reduce validity in arguments featuring value-premises to facts about rational relations between effective states.

Embeddings and the Independence Problem

There are other problems associated with embedding. We think that there are objective facts of value independent of our attitudes. Part of our conception of the objectivity of value is that we are willing to assert sentences like (i) and deny (ii):

- (i) It might have been that M-ing was right but I did not approve of M-ing.
- (ii) If M-ing is right, then I approve of M-ing.

Given normal functioning of a language agent, there is absolutely no disposition to perform assertions of:

- (iii) M-ing is right \rightarrow I approve of M-ing.

To think the entailment holds is to think that part of the truth-conditions of *M-ing is right* is that I approve of M-ing. But that is false. (We say more about related matters in chapter 11, §80.) In which case (iii) provides no basis for affirming (ii) or denying (i), and there are no other such bases. If so, we can see facts of value and facts about people's motivations as independent. A world of value does not require a world of motivated people. However, a world of valuers requires a world of motivated individuals.

§41 Wishful Thinking, Objectivity, and Value's Reality

Dorr (2002) has presented an interesting objection to expressivism. This is that expressivism commits itself to a form of wishful thinking. In wishful thinking, one thinks, irrationally, that the world is a certain way as a result of a desire that it be so. Consider *W*:

W: If George did something wrong, God will punish him. George has done something wrong. Therefore, God will punish him.

Dorr sees the problem in these terms: if one accepts the conclusion of *W* then one believes that the world is a certain way—George will be punished—by virtue of the fact that we do not like what he has done. If we do not think that God acts according to our desires, or that there can be a common cause of our desires and God's, this is a case of wishful thinking.

Dorr's argument trades on confounding an assertion of *George has done something wrong* with *We disapprove of something George did*. If we substitute the latter sentence for the first in the conditional in *W* we go from (i) to (ii):

(i) If George does something wrong, God will punish him.

(ii) If we disapprove of something George did, God will punish him.

One salient reading of (ii) is that God punishes as a result of our likes and dislikes. But there is another reading: our disapproval of an act is a reliable guide to God's disapproval of the act, and thus of how God will act: we overlap in attitudes with God. In which case, the conditional (ii) is not irrational. Still, it carries a different connotation from (i). But that is OK. Utterance of (ii) reports what (i) expresses. If there is overlap in God's disapproval dispositions and our own disapproval dispositions, and disapproval is linked to punishing behaviours, then it is rational to have a C-law that links a disapproval state in us to a Π -property underpinning assertion of God's punishing someone. If our feelings of moral outrage—the phenomenal description of the approval states that underpin moral utterances—were no guide at all about other's states, including God's, we would not have a moral community. We do not make this explicit in (i), but (ii) is a conditional that is about the coordination of affective states required for moral community. This is the difference between (i) and (ii). (ii) is a conditional report about the mechanism, (i) is an expression of the kinds of C-laws that will be instantiated in speakers who participate in it. I conclude that there is no wishful thinking, rather only hope of moral community.

This discussion brings out a general feature of human systems of approval and disapproval. Consider taste-statement making. In judging that this pasta is tasty, I offer some to you, in the belief that you will like it. I reason: *If it's tasty, you will like it*. I might similarly be accused of wishful thinking. But what rationalises my conditional is my participation in a community of taste, in which there is significant agreement in gustatory preference, which can be useful to exploit for certain, entirely commendable, socially oriented purposes.

Disagreement

According to GE, judgements produced by asserting sentences of the form *T is tasty*, *T is cute*, *T is cool*, *T is beautiful*, *T is fantastic*, *T is superb*, or *T is (aesthetically) good* will all have the same basic structures given above. What changes are the kinds of affective states, which are partly differentiated by matters of topic, the kinds of properties concerned, or the strengths and patterns of behaviours. Talk about a meal's being tasty—or a jacket's being cool, or an installation beautiful—can lead to dispute and dialogue with the same structure as those for moral value-discourse.

There is, however, a difference. We think that the objectivity with respect to taste discourse is less robust, or one we take less seriously, than in the case of morals. This is to say that we opt out of debate fairly early, if there is disagreement, in the case of assertions of taste. We might put aesthetic dispute somewhere in between that of taste and morals with regard to the degree of robustness. Why is it that we are inclined to opt out in the case of disagreements about taste, less so in the case of aesthetics, far less in the case of disagreement in morals, and even far less so for serious moral disagreement?

One answer, which I reject, is that for taste, there is no mind-independent reality or objective facts corresponding to the discourse, whereas there is in the case of serious morals and aesthetics. One can talk in these terms, but this is talk within the domain of the interpretative stance. We frequently talk of objective fact. But appeal to objective fact does not do any explanatory work. What we need to do is examine the language agency that underpins such talk. We shall see that it reflects features to do with how we manage our activities of defending Π -properties in certain domains of talk. (I do this in chapter 9.)

Moral Reality

Our expressivism has taken it that the Π -properties underpinning value talk are affective in their character. That does not imply that there are no moral properties, that value-sentences are not made true by facts of value, or that there is no moral reality. We talk in terms of properties of value, facts of value, and moral reality. These things all exist. What GE rejects is that correspondence to properties, facts, or realities have any role in explaining the language agency of value-sentences. As we shall see in chapter 7 there are moral properties. These are just as real and robust as physical properties. In chapter 8, we shall see that there are moral facts as much as there are physical facts. On the other hand, we shall see that there is no *location problem*—to use Jackson’s (2000) phrase—the problem of where to place moral properties in the natural order. There is, essentially, no such problem because there is no metaphysical problem about the nature of properties and facts.

Perhaps not all value claims are assertions whose Π -properties are based in affective states. Perhaps moral epistemology will show that we need to introduce something more like perception of value in our analysis of the language agency of value-talk. That is not a problem for GE. All that would mean is that the Γ -states for certain value-predicates should be analysed in using structures that are closer to, say, the Γ -states of colour terms. We would still maintain expressivism about value, it’s just that it would not be a motivational expressivism about value. The case for expressivism is independent of any particular thesis about the epistemology of assertions in any domain. See §34.

Chapter 6

Word and Object

§42 Reference, and Causal and Interpretative Perspectives

GE, within its perspective of functional analysis, provides a theory of agents causally embedded in a broader reality. We are now going to turn that analysis of functional structures towards the activity underpinning referring terms. To that end, what follows is a functional analysis of the semantic category of referring terms, and various divisions within that category. Here some surprises are in store for us. First, what this analysis is not. In what follows, we provide a theory of the functional structure of attributions of reference. But we provide no theory of reference. Talk of reference requires that we step down from the heights of the functionalist perspective, and take up a stance within some language agent signalling system and engage in the simulations of L-processing, that is, the perspective of *Verstehen* or interpretation. GE does not reject reference, it just rejects its playing a role in a theory of the home language. In order to make assignments of reference we must take up an L-processing stance towards agents, but in analysing the home language we are not taking up this stance, but describing the functional structure of such a stance. So we cannot use the term *refers*. (The home language is not analysed from an interpretative stance—see §1.) Furthermore, in analysing the language agency behind *refers*, no worldly correlate is assigned to *refers* or the mental element behind it. So, the question, posed within the interpretative perspective, *What is reference?*, has no answer. Certainly, there is no causal reduction or even theoretical reduction, implicitly defined by its theoretic role in the context of an interpretative theory.

GE's approach closes off theoretical gratification in any theory of reference. But opens up a surprising distinction in how we understand our talk of things. Referents in orthodox semantics are those things to which referring terms refer and which play a role in

relation to explaining predication. In explaining how a simple sentence *T is F* functions, we display how its truth-conditions are fixed by the referent of *T*, and the property or extension assigned to *F*. This is the essence of an objectual view of predication—§3 and §31. In that view, every referent is an object: it is just something in the world to which predicates can apply. GE does not deny that a sentence *T is F* is true because the referent of *T* satisfies predicates *F*. It just denies that such talk has a role in illuminating the nature of predication. For GE all the work illuminating predication is outside the semantic domain. For that reason, GE can make space for the curious idea of the referents that are *objects* and those that are *non-objects*. That space is opened up through GE's analysis of the cognitive precursors of referring terms, that is, what I have called Φ -properties and repertoire dispositions that link Φ -properties to terms. As we shall see, some Φ -properties are capable of entering into causal relations with worldly things, where those things can be referents of terms. But some Φ -properties are incapable of entering into such relations. The referents of terms with these Φ -properties never have any explanatory role in language analysis. The first kind of Φ -properties underpin object terms, the second kind underpin non-object terms. I do not expect that what I am talking about will be immediately obvious. But as we proceed, matters should clarify themselves.⁷²

§43 *Inferentialism and Reference*

First a brief comment about inferentialism. Fans of inferentialism may discern in what I have proposed about referring terms a parallel with inferentialism as developed by Brandom (1994). Brandom does not give a theory of reference. He gives an inferential analysis of referring terms and of the concept of referring. The category of *referring term* is defined inferentially, in the following way:

N is a referring term iff for every sentence *S* in which *N* is a constituent, if *R* is the result of substituting *M* for *N* and *S* entails *R*, then *R* entails *S*.

⁷² My analysis of referring terms largely restricts itself to singular referring terms. There are also plural denotation and plural referring terms. That there is plural referring has been argued elsewhere. See Boolos (1984), Rayo (2002), Barker (2004), and Sainsbury (2005). I do not provide a treatment of plurals here.

Thus, according to this theory, *Ted* is a referring term, because for any sentence containing *Ted*, as in *Ted is a dog*, if we replace *Ted* by a term, say *Jane's dog*, to derive a new sentence *Jane's pet is a dog*, and *Ted is a dog* entails this sentence, then the reverse also holds. In contrast, the term *a dog* in *Ted is a dog* is not a referring term, since we can derive by substitution for that term a sentence *Ted is a mammal*. Our original sentence *Ted is a dog* implies the latter, but not vice versa.

Having thus analysed referring terms, Brandom analyses the verb *refers*. He does so by first analyzing phrases of the form *the referent of x*. Such phrases are devices of anaphoric reference. The term, *the referent of 'Ted'* is simply a way of re-issuing *Ted*. Presumably, *T refers to O* makes explicit a practice of substituting *the referent of 'T'* for *O*.

So much for the basic elements of the theory. As should come as no surprise given the exiguous materials being used to carve out the category of referring terms, this theory does not work. According to the theory, if *Ted* is a referring term, then for every sentence *S* in which *Ted* is a constituent, if *R* is the result of substituting *M* for *Ted* and *S* entails *R*, then *R* entails *S*. But, we can easily find terms *M*, which substituted for *Ted*, generate sentences *R*, which *S* entails but which do not entail *S*. Example: let *S* be *Ted is intact* and *M* be *Ted's foot*. We accept: *Ted is intact* implies *Ted's foot is intact*, but we do not accept that *Ted's foot is intact* implies *Ted is intact*. One cannot object that *Ted's foot* is not an appropriate substitution instance because it is a quantifier phrase. Even if one were to accept that diagnosis based on an independent characterisation of quantifiers, one can run the example with a name for *Ted's foot*, like *Footy*.

I conclude that Brandom has not provided us with an adequate analysis of referring terms.⁷³ To understand referring terms what we need to do is go beyond the phenomenal realm of surface forms and inferential practices, and enter the noumenal realm of cognitive structures. This is where we find GE.

⁷³ For related objections to Brandom's theory see Taylor (2005), and Fodor and Lepore (2001).

§44 Referring, Proto-Referring, and Denoting

Like inferentialism, GE does not offer a theory of reference just as it offers no theory of meaning. It provides a language-agency analysis of reference attribution, a theory of referring terms, and a theory of referring acts. Referring is like asserting and other illocutionary acts like ordering or implicating: it is the production of a symbolic string signalling a state. In asserting—see §14—the state signalled is a ground-indicating disposition Π^{GI} in relation to a Π -property Π . (We called this *defending* a Π -property.) In the case of referring acts, the state signalled is what I have called a Φ -state. Referring acts are acts of signalling Φ -states. We are not proposing, in making this claim, that the Φ -state is the referent or semantic value of the term. That would be complete confusion. And anyway, in undertaking functional analysis, talk of reference and referents and semantic values have no place.

Just as we had to go to some trouble to explain what Π -properties and ground-indicating dispositions were, we need to go to some trouble to explain Φ -states. I do this below. Like assertions, referring acts can be sincere or insincere. I refer to George Bush, and really intend to denote the President. In this case I signal a Φ -state, which I really have, and my audience L-processes my utterance using a homologue state. On the other hand, I can signal a Φ -state I lack. I am under the delusion that George Bush does not really exist, but believe everyone else thinks he exists, but I do not want anyone to realise that I disbelieve. I therefore pretend, in the way that a liar pretends, that I have a state I lack when I use *George Bush*. I perform insincere referring acts.

Assertions are produced by performing proto-assertions and being directed towards suppressing any cancelling signals—see §14. A proto-assertion is an act caused by a repertoire disposition of the form $\text{RD}[\Pi^{\text{GI}}, S]$. Referring acts are similarly constituted through proto-referring, and thus through repertoire dispositions. Proto-referring acts are productions of symbols caused by repertoire dispositions analogous to those that underpin proto-assertion. The repertoire disposition underpinning a proto-referring act is a disposition $\text{RD}[\Phi, T]$ with the following character:

RD[Φ, T]: If U is directed towards producing a string caused by a state of the kind Φ , then U can utter T .

Here Φ is a Φ -property, as I have called it. The utterance of a term T is a proto-referring act using T just in case it is caused in a specific way by a disposition state RD[Φ, T]. More precisely, a proto-referring act can be characterised thus:

U performs a proto-referring act with T iff U utters T and (i) U instantiates a state RD[Φ, T]; (ii) RD[Φ, T] is a cause of her utterance of T ; and (iii) U is directed towards making RD[Φ, T] a cause of her utterance of T .

This is more or less parallel to proto-assertions—§14. Note clause (iii). Human language acts are not automatic or mere reflexes. What drives production of acts is the speaker's being directed towards production. In the case of referring terms, that means that the agent needs to be directed towards her dispositional state RD[Φ, T] manifesting in an utterance and thus being a cause of her production of T .

A referring act is an act in which a proto-referring act is performed, in which the speaker suppresses any cancelling signals—signalling of play or conflicting states. If the speaker is sincere she really has the Φ -state. On the other hand, a speaker can perform a mere proto-referring act. Say U asserts *Pegasus does not exist*. The production of *Pegasus* is a proto-referring act, but not a referring act, in the sense we have defined, since through the sentential context in which U places *Pegasus*, U generates a cancelling signal, that is a sign indicating that no Φ -property causes production of her term.

In ordinary vernacular, both speakers and words used by speakers can be said to refer. Both words and people can refer to the existent and the non-existent. A person can refer to that which she believes to be non-existent. That is ordinary usage. For our technical use of the term *refers* from within the theory of language agency, we shall stick to the usage that U performs an act of referring if and only if U performs a proto-referring act without any signals of cancellation. I note that we are not giving a theory of O referring to some object. That

would require a functional reduction of content and GE does not provide that. We are giving a theory of the structure of a category of act: referring.

In what follows, I shall provide a theory of the language agency of our ordinary talk of speaker's referring and words referring. I look at talk of speaker's referring in §73. Below I examine briefly the language agency of talk of words referring. In ordinary usage, *refers* has both existential and non-existential uses. We can subscript these, *e* and *n-e* respectively, as we do in the table below, which represents the basic Π -property structure of referring attributions:

Referring/Denoting	Π -Property Identical to or Based in:
' <i>T</i> ' refers _{<i>e</i>} to <i>O</i>	Disposed to L-process H's tokens of <i>T</i> in terms of RD[Φ, O] for some Φ that U instantiates.
' <i>T</i> ' refers _{<i>n-e</i>} to <i>O</i>	Disposed to L-process H's tokens of <i>T</i> in terms of RD[Φ, O], for some Φ that U may or may not instantiate.

The term *refers*, like any verb or predicate, has a mental element that we can represent as D_{Ref} . And like any such mental element it will be based on mental states pertaining to motivation, perception, agency, metalinguistic dispositions, or processing thereof. D_{Ref} is based in L-processing. If U asserts that the term *the Philosopher* refers to Aristotle, teacher of Alexander, U defends a disposition to L-process that token term using RD[$\Phi_A, Aristotle$], where Φ_A is the Φ -state U deploys with *Aristotle*. In asserting that 'Pegasus' refers to the flying horse of Greek legend, U defends a disposition to L-process the token presented of *Pegasus* in terms of a state RD[Φ_P, O], where Φ_P is the property she deploys with *the flying horse of Greek legend*. But note the difference. In the case of *Aristotle*, U really instantiates the property Φ_A , for in this case U is disposed to use *Aristotle* on this reading sincerely. On the other hand, U may be a non-believer in the case of *Pegasus*, in which case the property Φ_P is one she does not instantiate.

In the case of *Aristotle* the property Φ_A is an informational state that is instantiated through certain epistemic pathways that we shall describe below. In the case of *Pegasus*, the property Φ_P is an informational set that corresponds to a possible state of the kind U instantiates with Φ_A , but which U does not instantiate. Rather, in using *Pegasus*, U will be disposed to perform some set of proto-assertions that will define an informational set

corresponding to a possible Φ -state. It would be fixed by predicates like *horse, flies, spoken of by such and such a text*.

Let us now take stock. We say that inferentialism attempted to analyse what being a referring term is in terms of superficial features of inferential activity. But this failed. GE offers a cognitively deep theory. Referring terms are those terms used in proto-referring acts.⁷⁴ At the heart of these structures is the mental property kind, the Φ -property. The rest of this chapter, and much of chapter 7, revolve around issues of the structure and nature of these Φ -properties. Like Π -properties, Φ -states come in distinct kinds. Φ -properties come in two basic kinds which fix two basic kinds of referring term: object and non-object terms. We now examine this distinction.

§45 *Objects and Non-Objects*

In orthodox semantics referring terms refer to objects, if they refer at all. We have been taught by Quine that one of our criteria for things being objects is that they are quantified over, which means they can satisfy predicates. Predicate-satisfiers are objects. Referring terms pick out predicate satisfiers, so, referents are objects. Orthodoxy with its objectual treatment of predication, must treat terms like *redness, being red, that T is red*, etc., as denoting objects—see §3.⁷⁵ And then with quantification over their kinds, ontological commitment is incurred. Even if we are austere truth-theoretic practitioners, metaphysical puzzles can only ensue in the form of questions: What is redness? What is a proposition that T is red? What is the fact that T is red? It seems we must either remove these idioms from canonical speech or face the metaphysical consequences and provide some theory of these entities. This is what makes being a Quinean naturalist hard—see §4.

GE is not forced down this path. GE clears space for the idea that there are non-objects referred to by basic subject terms, terms for properties, relations, facts, and

⁷⁴ Some make it a necessary condition for referring terms that they involve some kind of causal connection—Evans (1983) and Neale (1990).

⁷⁵ In a standard ontic semantics, abstract substantives like *redness* would denote universals, sets, or fusions of tropes, and so on, and gerunds similar entities—see Moltmann (2004).

propositions. By non-object we do not mean Frege's unsaturated entities, or fictional entities, or minimal entities. What we mean can only be clarified—though not defined or analysed—by looking to cognitive character of the Φ -properties associated with referring terms.

Referring terms that we call *object terms* are terms whose associated Φ -properties are derived from the cognitive structures involved in the perceptual tracking of particulars. I call these Φ -properties *ι -properties*. In perceiving a particular, say a cat, U's cognitive system brings to bear a state whose activation is causally dependent on features of the particular. Perceptual states carry information that, at some level of processing, correspond to activated generic Γ -states, that are the mental element underpinning kind-predicates—see §28. What I call *ι -properties* come in two forms. The basic kinds are cognitive elements that are derived directly from perceptual states, preserving some of the information in those states. Object-terms whose ι -properties are so based refer by *acquaintance*. That is, if U's object term *N* refers by acquaintance to O, then U has a state ι that is the causal consequent of a perceptual state caused by O. The second kind of ι -properties are not derived directly from perceptual states, but constructed from elements of perception-based states. For example, *the physical universe* is an object term but its ι -state is not based directly in a perceptual state, but is a state constructed from elements of primary ι -states.

Let us be very clear that we are not analysing what it is to be an object, nor giving a reductive analysis of the term *object*. Rather, we are providing an account of the language activity underpinning use of the term *object*. Summed up, the basic idea is that the term *object* is a formal predicate like *identical to T*—§27. That means that the Π -property of *N is an object* is a metalinguistic one, as in:

Object	Π-Property Specification
<i>N is an object</i>	Disposition to deploy <i>N</i> in a referring act whose Φ -state is an ι -property.

The resulting sentence *N is an object* does not have metalinguistic truth-conditions. It is not a sentence about how we use the term *object*. As we noted, sentences are not about their associated Π -properties—§11.

Non-Object Terms

Not all referring terms are object referring-terms. There are a range of terms, those for properties, relations, propositions, and facts, whose Φ -states are not based in ι -properties, which is to say in perceptual tracking states. I illustrate the idea in the case of property terms, such as *the property of F* or *F-ness*.

Consider a predicate like *is red* and the counterpart property terms: *the property of red* or *redness*. Clearly there must be a relation between predicate and property term. Whatever we assign as the reference of a property term, *F-ness* should relate to the function of the corresponding predicate *F*. In truth-conditional orthodoxy, predicates have reference in the sense that they are correlated with things that enter into the definition of truth for simple sentences featuring the predicate. So, in the case of *red*, referential assignments to the predicate might be: (i) a class of red things; (ii) a plurality of red things; (iii) a universal of red; (iv) a mereological fusion of red-instances—see §31. If so, the referent of *redness*—if we allow such a term at all—needs to be drawn from entities of the kind (i) to (iv) that are correlated with the predicate *is red*. So *redness* must refer to something like a class of red things, or a universal, and so on. In short, given that predicates are correlated with entities, properties must be entities. We also, apparently, quantify over such things. So, assuming Quine's criterion, they must be entities (if that is, we accept this mode of speaking at all).

But now consider GE. In GE, the function of the predicate in its role as the constituent of a simple sentence does not involve its correlation with worldly items, which then enter into the characterisation of truth. So nothing like (i) to (iv) above is correlated with the predicate *is red*. The Γ -state D_{Red} , although responsive to objects in perception, does not itself function by correlation with some worldly thing. But then how do we understand use of *redness* as an (abstract) singular term? GE would appear to have nothing to draw upon to explain its function, since no entities are correlated with predicates. However, the conclusion that GE has no resources to explain the singular terms *redness* or *the property of red* is false, and reflects an implicit acceptance of objectualism about predication. That predication is analysed by representation or truth-conditions.

GE's analysis of abstract singular terms is as follows: GE analyses *redness* as a nominalization of the predicate *red*. Think of a nominalization as a simple grammatical transformation of the term *red* enabling it to take up subject position in a sentence. In that transformation, the mental element that underpins the use of *is red*, is exactly that which underpins use of *redness*. How, then, can *redness* function in predications, as in *Redness is a property* or *Redness is possessed by this object*? No problem! To function as a referring term, all that a term needs is to be underpinned by a mental element that can combine with the mental element of a predicate, to form a Π -property, something which is defensible in assertion. It turns out that the mental element underpinning the predicates *is a property* and *possessed by this object* are the right kinds of cognitive material to combine with the mental element of *redness* to form Π -properties.⁷⁶

The core claim then is this: In judging that U's term *Redness* refers to the property of redness, the Φ -state underpinning the term, D_{Red} , is not one whose function involves correlation with something in the world. Nothing is causally correlated with it that is assigned as the referent of the term. This is not to say that it does not have a referent. It is rather, that in explaining the functional sub-structure of successful reference, no worldly correlate is present. In contrast, the Φ -state of *this cat*, as used by U to refer to a cat is causally involved with a thing in the world. In explaining how that state functions—when we examine matters from the perspective of language agency—we bring in a thing, which is the referent of the term.

There is more to say about this distinction. But that is better said after we have described the character of ι -properties and nominalization in more detail. But the basic idea now should be clear that GE is not buying into any obscure Fregean distinction between saturated and unsaturated entities. It also brings with it a radical thought: referring and quantifying over certain things does not bring with it ontological commitment as such.

⁷⁶ Some might be tempted to say that according to GE, terms like *redness* refer to linguistic expressions. But this is not the position, and thinking it must be, just reflects adherence to objectualism about predication.

§46 Reference to Objects and ι -Properties

I now sketch out how individual ι -properties work. The detail is partly speculation, bringing in postulated structures of the NC system. What is important is the spirit of the remarks, not the letter, but the letter helps us see the spirit.

The origin of ι -properties, the mental elements of object-terms, is perception. To perceive an object is to have an internal perceptual state that causally depends upon features of the thing tracked. A characteristic of perception of an individual object, say a red ball, is the production of a state that at some level of processing can issue in the joint triggering of generic tracking states—see §28. Thus in the veridical perception of a red ball, at some stage, D_{Ball} and D_{Red} will be caused by property instances in the object. The object is a bundle of instances of properties unified by spatiotemporal and causal features. The mental state that tracks the object is a cognitive bundle of states that correspond, in the veridical case, to instantiated features in the object.

We call a ι -property a cognitive bundle that is a causal result of a perceptual state. It is an information selection from the perceptual state, and so its informational richness is deeply impoverished in comparison. The ι -property is a cognitive bundle of Γ -states. I have been referring to Γ -states in two ways: Using the symbolism D_{Red} and so forth when discussing their dispositional characters, and the symbolism of the form $\Gamma\text{-}[red]$ to indicate their association with a predicate. I use the second form in what follows. So ι -property is a bundling of cognitive properties of the form $\Gamma\text{-}[F]$, for predicates F . Let us represent the bundling through the symbol, $^{\circ}[\dots]$. Thus we can represent an ι -property based in a perception of a red ball, as a state below:

$$^{\circ}[\Gamma\text{-}[Ball], \Gamma\text{-}[red], \dots]$$

That is, the Γ -states are $\Gamma\text{-}[Ball]$ and $\Gamma\text{-}[red]$ and the bundling relation is given by $^{\circ}[\dots]$. Such a state may be instantiated in the midst of perception or based in perceptual memory. Such states persist after perception has passed.

The ι -property above is comprised by Γ -[*Ball*] and Γ -[*red*] related cognitively in a specific way, that designated by $^{\circ}$ [...]. What is the cognitive bundling relation designated by $^{\circ}$ [...]? The ι -state is a cognitive bundle of Γ -states that is ultimately caused by a cognitive bundle of activated precursor states of incomparable richness, which is the perceptual state. So we have:

$$\{\text{Perceptual state of a red ball}\} \Rightarrow ^{\circ}[\Gamma\text{-[Ball]}, \Gamma\text{-[red]}]$$

The ι -state mirrors cognitive structure in the perceptual state. Furthermore, in the case of veridical perception, the cognitive bundle $^{\circ}[\Gamma\text{-[Ball]}, \Gamma\text{-[red]}]$ corresponds to—and is causally dependent on—a bundle of property instances that are spatiotemporally and causally unified. In short, we have three levels of isomorphic bundling: the bundling of worldly properties in an object; the bundling of cognitive properties in the perceptual state; and the bundling of Γ -states in the ι -property.

We shall look at the bundling of properties in the object later. We understand the cognitive bundling in the perceptual state in these terms. From the vast amount of information flooding in, a sub-cluster thereof is isolated. The information is isolated by: (i) detection of spatiotemporal relations between more basic features, like edges, colour differences, and so on; (ii) certain relations of causal dependency. Joint triggering of Γ -[*Ball*] and Γ -[*red*] occurs when the activation of these items is based in the same body of information so delineated. Hence an ι -state is marked functionally as the type of state that can be the effect of a perceptual state, such that the structure of the ι -state—its combination of certain Γ -states—mirrors structures in the perceptual state that is its cause.

Predication and Concatenation

The nucleus $^{\circ}$ [...] of an ι -property is, by its very nature, a sign of informational bundling in perception. This is its primary functional role. In turn, ι -states can be causes for other kinds of states that take us deeper into the cognitive system. The ι -property can also function as the mental precursor, or Φ -property, for a term, say *the red ball*; in other words it can enter into the causation of U's productions of *the red ball*. Because Φ -properties can combine with Γ -states to form Π -properties, ι -states can combine with Γ -states to form Π -

properties. Suppose that the perceptual state P that grounds ${}^0[D_{\text{Ball}}, D_{\text{Red}} \dots]$ also has information that activates precursors of Γ -[on the mat]. In which case:

$$P \Rightarrow {}^0[D_{\text{Ball}}, D_{\text{Red}} \dots] + \Gamma\text{-[on the mat]}.$$

This is a familiar *basing law* in the sense articulated in various sections above. See, for example, basing in the case of colour predicates in §28. We can put it in more familiar *notion*, since ${}^0[D_{\text{Ball}}, D_{\text{Red}} \dots]$ is Φ -[the red ball]:

$$P \Rightarrow \Phi\text{-[the red ball]} + \Gamma\text{-[on the mat]}.$$

The perceptual state P grounds the Π -property for: *The red ball is on the mat.*

In short, the cognitive bundling that we find in predication of empirical predicates finds its origins in, or is just a mirroring of, bundling at lower levels of perceptual processing. This does not mean that every appearance of ι -properties and Π -properties featuring them are products of perceptual states. It just means that their primary grounds are in such states. We shall see ι -states can have other causal origins.

Metaphysics and the Structure of Objects

The cognitive bundling of an ι -state finds its primary grounding in the cognitive bundling of a perceptual state tracking an object. The cognitive bundling of informational states in the case of a veridical perceptual state is causally dependent on, and mirrors, the bundling of instantiated properties in the object. The bundling in the object of property instances is of a spatiotemporal and causal kind.⁷⁷ Material objects—both whole continuents and temporal stages—are spatiotemporally and causally unified bundles of property instances. This claim is a structural claim that amounts to this: there are property instances, and in material objects they are arranged in a certain way.

This claim may sound like a claim in metaphysics. But merely to talk of properties or property instances and relations between them is not yet to *do metaphysics*. It is philosophically uncontentious to say that where there is a material object there are property

⁷⁷ Without doubt matters are more complicated. It is not a conceptually impossible to have spatially co-incident objects. Your body may be co-incident with an astral body; the astral body is colourless, but not your body.

instances that have various spatiotemporal relations to each other. What is contentious, and what signifies our entrance into the field of metaphysics, is providing theories about what properties are, what property instances are, what relations are, and what instantiation is. It is these questions whose answers we can call *theses of metaphysics*. Metaphysics is that discipline according to which such questions are urgent and demand serious replies. Surely, you may protest, these are legitimate questions. After all, *instantiates* is a perfectly meaningful term, and so there must be something that is this phenomenon of instantiation, whose nature we need to understand. But this idea is confused. The problem is the inference from '*Instantiates*' is meaningful to *There must be something in the world with a nature that corresponds to the term 'instantiates' giving it meaning*. The inference is illegitimate in the eyes of GE. It illegitimately assumes a representationalist concept of the function of *instantiates*. For GE, the language agency of instantiation-talk has no representational element. Predication with *instantiates* does not work by representation. No worldly correlate is assigned to the mental element underpinning use of *instantiates*. If that is right, having provided the analysis of the language agency of *instantiates*, there can be no theoretical basis to questions about the nature of instantiation. This is to say instantiation lacks any nature that could be theoretically investigated. It is a category error, mistaking a *non-object* for an *object*, that lies behind the investigation of its nature. But that is exactly what metaphysics does. We say more about these matters in the next chapter after we have examined property- and property-instance talk.

§47 Names

Such is the general nature of ι -states for objects in the sense of continuents, that is, physical and common sense objects. We move onto a central cognitive role for such states as the Φ -properties of names. A speaker might have a perceptual interaction with an individual, and the result is an ι -state. That ι -state may become a relatively permanent fixture in a speaker's cognitive life. We may suppose that U associates with it a term, *Smith*. A stable disposition emerges in which productions of *Smith* are caused by this state. In other words it becomes a

Φ -property for *Smith*. That means U uses *Smith* in a repertoire disposition: $RD[\iota, Smith]$. The ι -property is a dossier for an individual, referred to as *Smith*. Let us designate the ι -property so deployed by M_{Smith} . As U's informational state evolves through time, the composition of the dossier may change. U may have further interactions with the individual. Updating through perception involves a merging of information from separate perceptual episodes.

We have described a single-speaker name using activity. Name-using activities can be multi-person. Consider what happens when a speaker H acquires the use of *Smith* from U. H interacts with U, who provides testimony of an individual using *Smith*. Part of possessing a cognitive system for use of names is that U can in response to such interactions set up a dossier M_{Smith} with the structure below:

$$M_{Smith} : \quad {}^O[\Gamma-[person], \Gamma-[referred to by U], \dots]$$

H as a result possesses a new repertoire disposition with this state, M_{Smith} , as a Φ -state. Thus H now has the repertoire disposition $RD[M_{Smith}, Smith]$.

The activity of name-acquisition involves a cognitive innovation. The dossier initiated by H has the form of an ι -property. So far we have described ι -properties as states comprising a nucleus, ${}^O[\dots]$, whose characteristic causal origin is perceptual states. The ι -state instigated in H's when she acquires the name is not caused by her perceptual state. It is caused by U's earlier perceptual interaction with the referent of *Smith*. Thus H's dossier is a new cognitive deployment of an ι -state. As we saw, ι -states—states of the form ${}^O[\Gamma-[F], \dots]$ —have as their primary grounds perceptual states in the same speaker. But there can be, as we have now seen, other sufficient conditions for the production of such states. In this case, encountering speakers and extracting a state through L-processing. (So ι -states are like Π -properties, which have primary grounds, given through basing, but can also be tokened through other states.)

U and H's dossiers will contain elements that are homologues of each other. These homologues may not match perfectly—U may deploy some predicates that H does not. But H inherits hers from U. The causal path is a kind of extended artificial perceptual memory. Referential chains are the result of such graftings. This leads to a question of conditions of

reference, but as we shall see, no theory of reference emerges from this account of language agency. I shall deal with this below.⁷⁸

Proto-Referring and Names

We have defined the structure of a name-using activity for a sequence of individuals. But in fact matters are a bit more complicated. To be a name is not merely to be a term produced by a Φ -property of the form M_{Smith} . The repertoire disposition for *Smith* as used by a particular speaker is given thus, where M_{Smith} is a possible file for a name characterised by a set of Γ -states:

RD[$M_{Smith}, Smith$]: If U is directed towards producing a string caused by a state M_{Smith} , then U can utter *Smith*.

We might call U's instantiation of states like RD[$M_{Smith}, Smith$] her grasp of name-function. The cases we have looked at above are those where the Φ -property, the state M_{Smith} , is instantiated by the speaker. What causes production of *Smith* in these cases is: (i) the speaker's file M_{Smith} ; (ii) the state RD[$M_{Smith}, Smith$]; and (iii) U's being directed towards making RD[$M_{Smith}, Smith$] a cause of her utterance of *Smith*.

The conditions just described are not necessary for a name-using activity. Say that a speaker in the practice of using *Smith* comes to doubt that Smith exists. This speaker does not believe the term denotes. They dissolve their file for *Smith*. Yet they continue using the name, *Smith*. For example, they spend some time attempting to persuade others of Smith's non-existence. If so, their tokenings of *Smith* are not caused by Φ -states of the form M_{Smith} . Such speakers do not use *Smith* in referring acts—acts with intentions to denote. Nevertheless, they still perform proto-referring acts with the term. For such speakers, what lies behind production of *Smith* is their being directed towards production of a term caused by RD[$M_{Smith}, Smith$]. Lacking any mental file M_{Smith} , U cannot discharge this disposition. But the disposition can still be causally efficacious since it is a mental state, and so can combine with other mental states to produce action. There is no actual dossier M_{Smith} that U possesses.

⁷⁸ My characterisation of this basic use of names is partially inspired by Sainsbury (2005).

However, U does associate a body of information with the term. This comes in the form of a set of dispositions to perform proto-assertions with the term in sentences like: *Smith is a woman*, *Smith is a witch*, *Smith is believed by X to exist*. These are proto-asserted but have various interpretations. The first two will be fictional, the third non-fictional. The predicates deployed in these sentences have Γ -states, which as a set, correspond to a possible file.

Referring and Referring Attribution

So far we have given a structural account of agency with names. What then of reference? With a reductionist frame of mind, one might be tempted to propose the following: a name token *N* used by U refers to an object O if and only if there is a causal chain, made up of proto-referring act episodes associated with information bodies, which leads back to a perceptual state in which an object O is perceived. Reference for names is a causal matter. This reductionist line of thought does not follow from the theory of referring terms developed. We have described functional structures underlying name-using activity. These structures comprise bundled Γ -states. But the functional structures of empirical Γ -states do not fix contents—see §26, §30, and §31. Rather content-talk only comes through an interpretative stance that brings in L-processing. Information about causal connection can ground judgements of reference only in conjunction with the deliverances of interpretative judgements based in L-processing. If that is right, name-reference cannot be reduced to causal patterns. What GE offers is a necessary condition for name-reference: that such causal patterns are in place. Don't think this means that GE is incomplete and that supplementary sufficient conditions are out there somewhere. Content is functionally irreducible because interpretative judgements can depend, in ways that have no exact algorithmic representation, on appreciation of holistic factors and cognitive limits—§26 and §31. From the functional perspective, as theorists we have at most probabilistic judgements about the likelihood that a given system, a language-agent described in functional terms, will L-process another agent's use of a term *N* in a certain way.

Name Meaning?

Finally, there is the question of the meaning of a name. What is the meaning of a name, and its semantic contribution to simple sentences in which it features? Orthodoxy

provides answers such as: a denotation, a description, a sense, or a reference axiom.⁷⁹ What is GE's answer? From within the perspective of functional analysis of language agency, the question cannot be raised, since it deploys a vocabulary that cannot be used in that perspective. Rather, the question about the meaning of a name can only be raised once we adopt the interpretative perspective in which S-processing can ground use of semantic vocabulary. Within this perspective, the question, *What is the meaning of a name?* has no interesting answer. Indeed, it is doubtful that the term *meaning* really applies to names. What is the meaning of *George Bush*? The idea that names must have meanings arises from a theoretical imperative generated by the idea that there is a theoretical enterprise called *the theory of meaning*. If we are released from that imperative, the question loses its interest.

We might allow the pronouncement: '*George Bush*' means *George Bush*. But if we did, that would not vindicate, say, some Millian conception of names, according to which the meaning of a name is its denotation. In asserting this sentence we defend a disposition to L-process *George Bush* in terms of a repertoire disposition: RD[Φ_{Bush} , *George Bush*]. But we can assert equally '*Pegasus*' means *Pegasus*. In this case, we defend an L-processing disposition in relation to a repertoire disposition: RD[Φ_{P} , *Pegasus*]. Nothing in this practice then should lead us to theorise that the meaning of a name is a denotation.

§48 Reference by Descriptive Uniqueness

Names feature a form of reference that we can call *acquaintance*. Where U refers to an object O, denoting that object, then states of the object cause production of the term. But not all reference has this form. There is also what we might call *descriptive reference*. In using *the fattest man*, we seem to be able to refer to something to which we have no causal connection. We do so by unique specification: the term refers to the thing that uniquely satisfies *fattest man*. Philosophers have argued that such terms are not really referring terms because of this

⁷⁹ The classical view is associated with Russell (1905), McDowell (1977), Kripke (1980), Evans (1983), Salmon (1986), Soames (1987), and Neale (1990), the second with Frege (1892), and with Katz (1995), the third with Burge (1973). For a recent discussion of theories of referring terms in the orthodox framework, and defence of the last theory, see Sainsbury (2005).

lack of a causal connection, amongst other reasons.⁸⁰ Another reason lies in the concept of unique specification itself. We need to explain the concept of unique specification without using definite descriptions. Russell's (1905) theory of descriptions offers us such a possibility. It does so in these terms. A sentence of the form *the F is G* is translated:

At least one thing is F, and everything that is F is identical to it and it is G.

But having analysed uniqueness in these terms, we may doubt that descriptions are referring terms. Instead, they are quantified noun phrases.

There are doubts about Russell's analysis. According to it, a sentence of the form *the F is G* is true iff there is a unique F and it is G. But this looks wrong. Suppose I am told that the chair in the room is red. I then open the door in the room and find 50 chairs all red. Is the sentence, *The chair in the room is red*, without any further interpretation false? I don't think we want to judge it either true or false.⁸¹ But according to Russell it is false. Theorists have attempted to explain away these disparities in terms of a theory of conversational pragmatics—see Bach (1987), Neale (1990), and Recanati (1993). This strategy affirms that the sentence really is false, it's just that some conversational factor prevents us from asserting this. I will not evaluate the success of such approaches here. Is there an alternative analysis of uniqueness to Russell's that sticks closer to usage? There is, but in offering it, we give up the idea of providing a reductive truth-conditional analysis of sentences featuring definite descriptions. That path can only lead to infliction of damage on our patterns of use. Rather we provide an account of language agency.

GE's Analysis

GE holds onto the view that *the F* is a referring term by providing a theory of the mental element underpinning what we call referring by unique specification. From the analysis to be proposed we cannot extract informative truth-conditions for definite description-bearing sentences. All that we can say is: *The F is H* is true if and only if the F, the unique thing instantiating *F*, is H. The interesting explanatory work is done at the level of

⁸⁰ See Russell (1905), Evans (1983), Bach (1987), Neale (1990), Recanati (1993) to name just a few.

⁸¹ See Ramachandran (1993) for discussion of such cases.

the cognitive conditions that underpin use of *the F* and of evaluations of truth and falsity of sentences using *the F*. Nothing informative happens at the level of truth-conditions. Russell's theory of descriptions is the result of thinking there must be such an informative analysis.

The analysis of cognitive structure has a similar character to the style of analysis we have given of Π -properties in terms of basing. There are basing relations for Π -properties but also for Φ -states. We have already come across the latter. We saw in §46 above that in using a name, an agent can deploy an ι -property as an effect of a perceptual state. The perceptual state \mathbf{R} is like the canonical ground for tokening of an ι -property. Let ${}^{\circ}[F]$ be an ι -property with the structure ${}^{\circ}[\Gamma-[F_1], \Gamma-[F_2]...]$. We can represent this grounding as: $\mathbf{R} \Rightarrow {}^{\circ}[F]$. On the other hand, there are other grounds for ι -properties. Through testimonial interaction with another person, U can initiate a ι -property. This occurs in interpersonal name transition.

What we find in the case of definite descriptions is a functional modification of an ι -property ${}^{\circ}[F]$. I shall denote the Φ -state of a definite description *the F* by ${}^{\circ}[F]^D$. What we shall do now is functionally define this mental element. We shall do so by first characterising the primary ground for ${}^{\circ}[F]^D$. Let me first give an intuitive characterisation of that state. Suppose that U is searching through a room. In the room at some stage she spots a cat. That means, D_{Cat} is triggered—and U's system also sets up an ι -property of the form ${}^{\circ}[\text{Cat in the room}]_x$. (The subscripting just means it is assigned a particular file-status: x.) U continues in the room making perceptual sweeps. In further searches, where a state of the form ${}^{\circ}[\text{Cat in the room}]_y$ is generated—say her glimpses of a cat—U is disposed to form an identity judgement based in these new ι -states and the initial state ${}^{\circ}[\text{Cat in the room}]_x$. Cognitively, that amounts to the fact that where ${}^{\circ}[\text{Cat in the room}]_y$ is tokened, U is then disposed to inter-substitute that state in Π -properties. Such a disposition can underpin the metalinguistic dispositions that are defended in identity claims—see §27. That is, U develops a cognitive inter-substitutional disposition that we can label: $\text{IS}\{{}^{\circ}[\text{Cat in the room}]_x, {}^{\circ}[\text{Cat in the room}]_y\}$, which can ground sentences like *This cat is identical to that cat*. U perceptually sweeps the room, and is satisfied that her perceptual investigation is thorough. U's information processing state can be summed up in these terms: all perceptual triggerings of $\Gamma-[\text{Cat in the room}]$ have been linked by identity. The conditions form a canonical ground for

introduction of a definite description Φ -state, a state of the form $^{\circ}[Cat\ in\ the\ room]^D$.

Leaving the room, U is now disposed to use a description of the form: *the cat in the room*.

We can sum these conditions up in the following terms. There is a body of perceptually based information \mathbf{R}_x , such that, $\mathbf{R}_x \Rightarrow ^{\circ}[F..]_x$, and, whenever a perceptual state \mathbf{R}_y triggers an ι -property $^{\circ}[F..]_y$, then U's system is disposed to token an inter-substitutive Π -property: IS $\{^{\circ}[F..]_x, ^{\circ}[F..]_y\}$. In which case, we can capture the canonical ground for tokening a definite description Φ -state $^{\circ}[F]^D$ as a complex C-law of the form:

$$\left. \begin{array}{l} (a) \text{ There is some } \mathbf{R}_x \text{ such that } \mathbf{R}_x \Rightarrow ^{\circ}[F..]_x \\ (b) \text{ If some } \mathbf{R}_y \text{ is such that } \mathbf{R}_y \Rightarrow ^{\circ}[F..]_y \\ \text{then U tokens IS } \{^{\circ}[F..]_x, ^{\circ}[F..]_y\}. \end{array} \right\} \Rightarrow ^{\circ}[F..]^D$$

This C-law captures what is essential to the language agency underpinning *the F*. I note that in using $^{\circ}[F..]^D$ for a definite description, U may not use all the descriptive material in the ι -state $F...$. The description may be *improper*. (See Neale (1990) for an orthodox discussion of improper definite descriptions.)

We can see, intuitively, how this relates to Russell's analysis. Russell's idea was that *the F is G* is true if and only if something is F, and everything that is F is identical to it and it is G. We have provided a cognitive structure that corresponds to that part of the analysis that says: *something is F, and everything that is F is identical to it*. But there is a massive difference between Russell's approach and GE's. Russell is giving an account of truth-conditions. This is the distortion that leads Russell to attribute, as part of what is said, that there is something that is F and everything that is F is identical to it. But that gives us the wrong truth-conditions. GE on the other hand is not giving a theory of truth-conditions, but of the mental element underpinning use of *the F*. No distortion of truth-conditions results.

Mental Models, Referential and Attributive uses

We frequently use definite descriptions when we have no perceptual access to objects. A case is of this is a definite description like: *the tallest man in the world*. We may use this purely on the basis of two beliefs: (a) there is a finite set of men; and (b) men can be ranked in terms of height. But where is the perceptual state \mathbf{R}_x , that underpins the introduction of the

ι -property $^{\circ}[\textit{man taller than all men in the set}]_x$? Here is an idea: we use mental models. That means the information given above can be transformed into a simulation of a perceptual set. In this mode, individuals are in an imaginary space ranked according to some spatial relation, like *being next to*, as in: $\langle a, b, c, \text{etc.} \dots \rangle$. This imaginary state triggers $^{\circ}[\textit{man to the left of all other men}]$, which is our simulation of a state of the form: $^{\circ}[\textit{man taller than all men in the set}]$. Moreover we have the uniqueness disposition in relation to the model. That's the basis upon which we introduce *the tallest man in the world*. This is just a description of cognitive mechanism, not of what makes it right to use this description or its meaning.

Definite descriptions are meant to come in two uses: *attributive* and *referential*—Donellan (1966). In GE's terms all this distinction amounts to is that sometimes what generates production of *the F* is an ι -state functioning like those for names. Suppose U points and asserts *the man in the corner is drunk*. In L-processing U's utterance we assign not just a descriptive ι -state, $^{\circ}[\textit{man in the corner}]^D$, but also, due to our recognition that U is in perception contact with a person in the corner, a perception based ι -property. It turns out that there is no man in the corner, but U does correctly perceive a person in the corner: a woman, who is drunk. In these circumstances, we take U to be asserting something true, based in our L-processing of the definite description in terms of the perception-based ι -state.

Truth-conditions

GE's analysis of definite descriptions is that they are referring terms. They are, cognitively speaking, deeply connected with such names and demonstratives. They represent cognitive developments therefrom. But how does this relate to truth-conditions? I rejected Russell's truth-conditions for *the F is G*. There is no informative story about truth-conditions. There is an informative story about conditions of evaluation. That story goes as follows:

- (i) U judges true *the F is G* iff she accepts use of *the F* and defends $^{\circ}[F]^D + \Gamma-[G]$.
- (ii) U judges false *the F is G* iff she accepts use of *the F* and defends rejection of $^{\circ}[F]^D + \Gamma-[G]$.

A requirement of judging *true* or *false* the sentence *the F is G* is that one accepts use of the referring term *the F*. That is not a controversial assumption. In judging true or false any

subject-predicate sentence *T is G*, one must accept the use of *T*. When one asserts that *T does not exist*, one must still accept the use of *T*. Even in a negative existential, the term *T* is used. What is the condition that has to be met to accept use of a referring term *T*? This is simply that there is no problem of identity about what thing *T* is. Note: this condition does not require that *T* denote anything. Thus we are happy to use *Pegasus* even though we know it denotes nothing. Denoting something is not a precondition for legitimate use of a referring term, *pace* the absurd requirements of a Millian conception of names. Rather, *Pegasus* meets the conditions of legitimate use because there is in principle no indeterminacy about identity claims in relation to *Pegasus*. If, for example, it was unclear whether I mean *Pegasus* the flying horse or *Pegasus* the invisible talking platypus, then one cannot accept the proffered token of *Pegasus*. If I am using *Aristotle*, but cannot tell you whether I mean the philosopher or business magnate, then the term, as used by me is dropped. We refrain from evaluating my utterance of *Aristotle liked fishing* as true or false.

Consider, then, the case of definite descriptions. Suppose U says *The red desk in the room is Hegel's*. I look into the room and find fifty red desks. I cannot use the term *the red desk in the room*, because there are a whole lot of questions of identity which are indeterminate. Did my friend mean this red desk, or this one, or this one? I cannot then evaluate: *The red desk in the room is Hegel's*. This means that I must refrain from judging the sentence true or false. On the other hand, if we discussing the Holy Grail and you speculate that it is in Jerusalem, whether or not I believe in the existence of the Grail, I accept your use of the term, and I am happy to token instances of the term as well. There is no issue about determinacy. Indeed, I can assert: *The Holy Grail is not in Jerusalem. Indeed, it is not anywhere. That is because there isn't any such thing*. I am happy to assert negative existentials of the form *The Holy Grail does not exist*, whereas, in the case of Hegel's desk example, I am not prepared to assert: *the red desk does not exist*.⁸²

⁸² We shall discuss negative existential below in §53. Our analysis produces the result more or less urged by Strawson's (1950) critique of Russell's theory. But we reject three aspects of this analysis: (i) *The F is G* lacks definite truth-value if and only if there is no F or there fails to be a unique F. (ii) The sentence's resistance to truth or falsity evaluation is a truth-value gap. (iii) No thought is encoded by the sentence under conditions of there being no F or no unique F.

§49 Events and Property Instances

Let us take stock. We have provided the basic theory of those Φ -properties called ι -states; that is, the mental elements that underpin object terms. That analysis is a pure language-agency analysis of production of such terms. Truth-conditions has had no role in this analysis. So far we have looked at material-object ι -properties. These are not the only kinds. There are also events, which are denoted by terms like *Fred's walking*. Clearly events are a central posit in our conceptual schemes. It is important to distinguish them from facts, which are non-objects, as we see in chapter 8. Events are parts of the spatiotemporal world. They are dated particulars that enjoy a rich set of empirical properties. They enter into causation.

Events can be objects of perception. Event *kind* terms are, in the simplest cases, perception-based empirical terms. Perception of an event involves a perceptual state, that at some level of processing will involve triggering of precursors to generic Γ -states for predicates like *walk* or *jump*, etc. These generic Γ -states function essentially as do generic Γ -states for predicates like *dog*, *red*, or *water*. It is just the pattern of response that is different.

Events typically involve objects. Thus, walks involve agents who walk, movements, objects that move, etc. Such perceptual states will nearly always involve, as a component, perceptions of entities that participate in the event. For example, in perceiving a walk, perception of an agent will also occur. (This may not always occur. Primitive perception of basic movements may just involve perception of a movement.) The intimate relation between the event and participating entity is a spatiotemporal-causal relation. In the case of a walk, the agent will have the same space-time track as the region of the walk for the duration of the walk. That intimate relation is reflected in the cognitive bundling of activated states in the perceptual state itself.

The basic Φ -properties for event terms are perceptual-based ι -properties. Event ι -properties have the form below, for the event of Fred's walking, where we use the same ι -property structure as found for continuent terms:

$${}^{\circ}[\Gamma-[Walk], \Gamma-[of Fred] \dots]$$

Events can be identified partly through objects that participate in them. In the ι -state above we have the relational predicate Γ -state, Γ -[*of Fred*]. That state is triggered by a spatiotemporal-causal relation between Fred and the event.

Events, like continuants, can be thought of as bundles of properties, but the bundling has a different pattern. But this is an issue of detail. Both continuants and events are alike in their basic natures.

Related to events, in the same category, are property instances. They involve instantiation of properties. The concept of a property instance is not a purely technical philosophical notion. It is part of sophisticated common parlance. We have locutions of the form: *The red of this flag is different from the red of this flag*. In asserting this sentence we affirm the qualitative non-identity of two property instances. A pre-condition for the truth of such assertions is that property instances have properties. For example, the red of this flag is slightly orange. If property instances have properties, then there are truths about the identity and non-identity of property-instances.

The structure of the ι -property of a property-instance term is just like that for an event term. The structure of the ι -state for the definite description *the red of this flag* is that below:

$${}^{\circ}[\Gamma\text{-}[Red], \Gamma\text{-}[of this Flag] \dots]{}^D$$

A predicate like *of this Flag* is underpinned by a Γ -state that is triggered by a spatiotemporal and causal relation between a property-instance and continuant. So in perceiving a red flag as a red flag, there is a joint triggering of Γ -[Red] and Γ -[*of this Flag*]. That cognitive bundling can then be the ground for ι -states of the form given above.

Such ι -properties can concatenate with generic Γ -properties to form Π -properties. This is the kind of Π -property that underpins a sentence like *the red of this flag is slightly orange*.

Metaphysics

Just as in the case of continuants we may now be concerned about metaphysical questions. We have characterised events as spatiotemporally and causally unified bundles of a certain kind. But this invites questions about instantiation, properties, and so on. As we

noted in §46, a classic question of metaphysics is what is it for a property to be instantiated? That question is really assuming that properties are a kind of object that can participate in the actualisation of concrete reality. Typically metaphysics draws a blank on this question. We must just postulate some irreducible non-relation that binds a property to a bare particular, or some such proposal. Fortunately, these questions make no sense in GE. Once we have given the functional analysis of property-talk and property-instantiation talk there is no space left to ask questions about the nature of properties or instantiation. The questions disappear because, in providing the language-function of property and instantiation, no correlate in reality is invoked. We might say that the language-agency analysis leaves nothing left to enquire about. This idea is reinforced below in §54.

Chapter 7

Word and Non-Object

§50 Non-Objects and Nominalization

The last chapter dealt with object-referring terms. This chapter deals with non-object referring terms. These include terms for facts, propositions, relations and properties, and propositions. Non-object referring terms are ones whose Φ -states are not ι -properties, but other kinds of states. They have Φ -properties that are borrowed from the mental elements of predicates, relational terms, and sentences—see §45. The grammatical process that allows the transference is *nominalization*. Nominalization allows predicates, relational terms, and sentences to enter into predicate position; that is, to function as referring terms. In §45, I gave an illustration of nominalization in action in the case of predicates transformed into property terms. Here is an illustration of sentential nominalization:

In uttering *Snow is white*, U performs a proto-assertion. U's utterance is caused by a repertoire disposition of the form $RD[\Pi^{GI}, \textit{Snow is white}]$. One cannot use *Snow is white* in predications, as in *Snow is white is true*. However, we can effect a grammatical transformation of the sentence that will enable it to enter into predication by attaching *that*, as in: *that Snow is white is true*. The application of *that* renders the second sentence apt to enter into predications, and cancels any sentential syntactic status.⁸³

S and *that S* are both correlated with the mental element $RD[\Pi^{GI}, S]$, but differ in syntactic properties. *That S* can combine with predicates whereas *S* cannot. Another way of articulating the difference is that the mental element $RD[\Pi^{GI}, S]$ functions as a Φ -state—the mental precursor of a referring term—in the case of *that S*. That difference is underpinned by

⁸³ The particle *that* is a nominaliser. It is not demonstrative, *pace* Davidson's *saying that* paper (1968).

there being an extra repertoire disposition causing production of *that S*. This extra repertoire disposition is:

If U is directed towards signalling a disposition to use $RD[\Pi^{Gl}, S]$ as a Φ -state, she adds *that* to *S* to form *that S*.

Thus the mental elements associated with *S* and *that S* differ, but the difference is purely in the operations governing syntax. A speaker H interpreting *S* and *that S* will L-process both through $RD[\Pi^{Gl}, S]$, but additionally L-process *that S* using the further repertoire disposition above.

Non-Objects

Let us return to the curious distinction between objects and non-objects. Referring terms whose Φ -properties are ι -states are object-referring terms. But referring terms whose Φ -properties are borrowed from sentences and predicates are non-object terms. These terms refer to different kinds of things. But what does that mean? We shall see that this distinction is not a metaphysical one: there is no theory of objects or non-objects and their metaphysical constitution. Properties, facts, and propositions are not strange kinds of entities that somehow remain non-entities. That is an incoherent conception. Nor are we taking a deflationary position according to which properties, facts, etc, are entities, but insubstantial ones. It is not clear what *insubstantiality* as opposed to *substantiality* would amount to. Nor are we proposing any kind of fictionalism; fictional objects are still objects. Nor are we proposing that property-terms or proposition-terms refer to linguistic entities. Again, that would make them object terms, and such a view would be manifestly false. To get used to this distinction one has to renounce the material mode within the interpretative perspective and enter the formal mode—one talks of the structure of language activity—within the functional perspective. Let me give a sense of it now.

Take simple object terms such as demonstratives used through perceptual interactions. If one uses a demonstrative sincerely and successfully to refer to an object O then one is in a state ι that is caused by O—see §45. Moreover, if the perception is veridical, aspects of the structure of O, its being a bundle of certain properties, is mirrored in the

structure of ι . That means that at the level of language function, in explaining how the ι -states of an object-term T works, we can advert to the structures of the referent. In that sense, the referent and its structures can play a causal-explanatory role in relation to the language function of the term T . We are not proposing that T 's referring to something is reduced to a causal relation, or that ι -states are always caused by objects. We mean only that essential to characterising ι -states is that they are potential effects of things which turn out to be potential referents of referring terms T , which have ι -states as their mental elements.

This admittedly subtle condition does not hold for non-object terms. Consider property terms first. If one uses a term, say, *redness*, to sincerely and successfully refer to a property, then one deploys D_{Red} as a Φ -property. But it is not part of the function of D_{Red} to be, potentially, caused by the referent of *redness*. In explaining the language function of *redness* with its Φ -state, D_{Red} , the referent of *redness* makes no essential appearance in the explanatory story. One might object that in describing instances of redness that a speaker with D_{Red} may be sensitive to, we use the term *redness*. But reference to properties through property instances is not essential. We can say equally that placed in the presence of something red, the system with D_{Red} has a certain probability of being triggered. The presence of red-things causes it to trigger. Note, similarly, that we do not identify D_{Red} by reference to some set of instances of redness, even assuming there were such a set. We do not say that any speaker possessing D_{Red} will enter a perceptual state that comprises triggering of D_{Red} in any perceptual situation where a member of the set is present. (That will not work as we saw in §30.) To conclude: the referent of *redness* has no essential explanatory role in relation to the function of the mental element, D_{Red} , underpinning that term.

GE does make use of property-talk all the time in characterising mental states. It talks of Π -properties, Φ -properties, and Γ -properties. Our claim is not that in characterising language-function we shall not make use of any property terms. It is simply that in characterising the function of property terms, like *redness*, the referents of such terms play no essential explanatory role in characterising the mental elements that underpin their use.

One might object that this just shows that GE must be committed to the view that property terms lack referents; they are empty of reference or just not referring terms. This

objection is based on a kind of implicit reductionism; the thought that somehow reference will be reduced to features of language-function. But that is exactly what GE denies. GE offers no theory of reference, and in a sense, it gives no theory of referents. GE gives no theory of the referents of property terms, but it fully explains the activity behind use of property referring-terms, and talk of attribution of referent to such terms.

Matters hold similarly for *that*-clauses. The mental element, or Φ -state, of *that S* is a repertoire disposition. The state $RD[\Pi^{GI}, S]$ is not defined so that, as part of its function, it is potentially causally activated by something that is a potential referent of *that S*. In other words, the referent of *that S* makes no essential appearance at the level of language-function.

In sum, the Φ -states of object-terms are, at some level, explained by things which are potential referents of such terms, whereas the Φ -states of non-object terms are not essentially explained by things that are the potential referents such terms. We might say that in relation to the language function of mental elements of *F-ness* and *that S*, GE is nominalistic. This is the functional basis of the object/non-object distinction. How does this relate to any intuitive use of the term *object* that we might possess? Objects have natures. To think of something as an object is to think that there is real sense in asking *What is it?* or *What is its nature?*, even if the answer is that it is primitive and basic. But such questions, I suggest, make no real sense for properties and propositions qua non-objects. Just as it is meaningless to ask, *what corresponds to a predicate or sentence in the world?* it is meaningless to ask *what is the nature of a property or proposition?*. That is because property- and proposition-terms are nominalizations. This theme is best taken up below after we have dealt with more of the machinery that attends our use of non-object terms.

§51 Predication: Objects and Non-Objects

Both object and non-object terms are referring terms. The table below displays in a unified way the character of singular referring terms. Singular referring terms are terms produced in certain kinds of acts—proto-referring acts. The mental element in such acts are repertoire dispositions, $RD[\Phi, T]$ —see §44. The nature of the Φ -property—the mental element that lies

behind production of terms in referring acts—determines the character of the referring term. As indicated in the table, Φ -states for object terms are ι -properties, and Φ -states for non-object terms are derived from those of predicates and sentences:

<i>Singular Referring Terms</i>	Production of T is caused by $RD[\Phi, T]$, where Φ is a state that combines with a Γ -state (a predicate mental element).	
Objectual: Φ is an individual tracker ι .	<i>George Bush</i>	<i>The inventor of the zipper</i>
Non-Objectual: Φ is a generic Γ -state, or $RD[\Pi^{GI}, S]$.	<i>Redness, that Snow is white</i>	<i>The property common to all cats The fact mentioned by O.</i>
	Acquaintance: uses memory trace	Descriptive: the mode deploys descriptive uniqueness

Let me stress—in order to ward off a gross misunderstanding—that Φ -states are not semantic-values, references, or concepts. We are describing the faculties of the language agent from the functional perspective and not from an interpretative perspective. Hence idioms of content cannot be deployed.

How can terms correlated with mental elements associated with predicates or sentences—elements so unlike individual concepts—take on the central role of referring terms? That is, how can such terms enter into predication, combining with predicates to form simple sentences that can be asserted and be capable of truth or falsity? If our vision of predication is given by the orthodox conception, the objectual conception—see §4 and §31—then nominalizations as referring terms make no sense at all. According to the objectual conception of predication, the theory that illuminates how simple sentences function in assertion, that explains how they bear thoughts and relate to reality, is a correspondence conception. And in that conception, since nominalizations don't function by representation, the idea that these things can really enter into genuine predication is nonsensical.

Matters are otherwise in GE. GE does not analyse predication in terms of representation. Subject-predicate sentences have Π -properties made up of a Φ -property Φ , a Γ -property Γ , and concatenation $+$. They combine to form a Π -property, as in: $\Phi + \Gamma$. The act of predicating a property of a referent is that act in which a speaker defends a Π -property of

the form $\Phi + \Gamma$. All that is required of Φ and Γ for them to feature as constituents of a Π -property is that they can combine to form a Π -property, which is to say, a mental element that can be based in a perceptual state, an affective state, an agent manipulation state, a metalinguistic disposition, and so on. No other external constraint is placed on Φ and Γ . The real constraints are internal: if Φ has a certain form, Γ must have a certain form, and vice versa.

I represent the possibilities in the table below. In the first two rows we find the sentences *The ball is red* and *Redness is possessed by the ball*. The Π -properties of these sentences are essentially the same: the concatenation of Φ -[*the ball*] with D_{Red} . Both are grounded in exactly the same kind of perceptual state. What differs is how the Π -property is encoded by the syntax of the sentence. The sentence *Redness is possessed by the ball*, features a nominal, *redness*, which has a generic Γ -state D_{Red} as its mental element. The predicate, *possessed by the ball*, has an individual ι -property as its Γ -state. The mental elements of predicates can be generic or particular. *Possessed by* is simply that device with which we can form a predicate out of a referring term:

Grounding State	($\Phi+\Gamma$)-structure	Sentence
Perceptual state of a red ball	Φ -[<i>the ball</i>] + D_{Red}	<i>The ball is red</i>
Perceptual state of a red ball	D_{Red} + Φ -[<i>the ball</i>]	<i>Redness is possessed by the ball</i>
Disposition to deploy <i>Redness</i> in a nominal with generic Γ -state D_{Red} .	D_{Red} + D_{Property}	<i>Redness is a property</i>
Disposition to have positive feelings in relation to triggering of D_{Give}	D_{Give} + D_{Good}	<i>Giving is good</i>
Disposition to defend Π of $\text{RD}[\Pi^{\text{GI}}, S]$	$\text{RD}[\Pi^{\text{GI}}, S]$ + D_{True}	<i>That S is true</i>
Disposition to assert with interest S	$\text{RD}[\Pi^{\text{GI}}, S]$ + D_{Interest}	<i>That S is interesting</i>

Just as referring terms can be object and non-object-terms, predicates can be object or non-object predicates, according to whether they combine with object or non-object terms. Some predicates are ambiguous, for example, *good* or *interesting*.

Consider the third line in the table and the sentence *Redness is a property*. The predicate *property* is a formal predicate in the sense of §27. The Π -properties introduced through formal predicates are based in metalinguistic dispositions. In asserting *Φ -ness is a*

property, U defends a disposition to use Φ -ness in a predicate nominalization; that is, a referring term whose Φ -state is derived from the Γ -state of a predicate. (Compare this with the term *object* in §45.) This does not mean that *property* is a metalinguistic predicate, it does not apply to linguistic items. It applies to worldly extra-linguistic things: properties. (See again §27.)

The fourth line in our table is the case of an evaluative judgement: *Giving is good*. The primary ground for the Π -property of such sentences is a property approval state—in the sense of §37. Again, in this case we find a basing law: {Being disposed to have approval states caused by triggerings of D_{Give} } $\Rightarrow D_{\text{Give}} + D_{\text{Good}}$. The Π -property functionally defined thereby is that encoded in a sentence *Giving is good*.

The fifth line is simply a truth-attribution sentence where the bearer of truth is a proposition. The primary ground for Π -[*that S is true*] is a disposition to defend the Π -property of a sentence *S*—see §20.

The sixth row is the case of *interesting* as applied to propositions: *that S is interesting*. The primary ground for Π -[*That S is interesting*] is a disposition to assert with interest a sentence associated with the state $RD[\Pi^{\text{GI}}, S]$. Thus we have a basing relation of the form below:

$$\{\text{Being disposed to assert } S \text{ with interest}\} \Rightarrow RD[\Pi^{\text{GI}}, S] + D_{\text{Interest}}.$$

Or in other words, one of the form:

$$\{\text{Being disposed to assert } S \text{ with interest}\} \Rightarrow \Phi\text{-}[that S] + D_{\text{Interest}}.$$

The ground is metalinguistic, but the resulting statement lacks metalinguistic truth-conditions—see §36.

That, in outline, is an illustration of GE's treatment of predication that allows for object and non-object terms to combine with predicates.

§52 Modes of Reference for Non-Object Terms

I have given an account of the basic non-object referring terms. The basic phrases are ones of the form *F-ness* and *that S*. Such terms are grammatical transformations of predicate terms and sentences. Their Φ -states are the mental elements of predicates and sentences, that is, properties of the form Γ -[*F*] and $RD[\Pi^{G1},S]$. Not all non-object terms have the form of nominalizations, however. Non-object terms can also take the form of descriptive phrases. Descriptive phrases can be used with two basic modes of reference: acquaintance and descriptive-uniqueness. A phrase like *A* below could be used with an acquaintance mode:

A: Jane has *that quality* (mentioned before)

A features a kind of demonstrative, *that quality*, which may be based in a memory of a property, which was mentioned earlier. The demonstrative in *A* works by acquaintance, that is, extended artificial memory. How does this work?

Suppose U perceives a red object as a red object. That means her perceptual state triggers D_{Red} . The perceptual activation of D_{Red} can generate some degree of inclination to token *redness*. This may be purely imagined sign-production. Let us say that U contemplates a property, redness in this case, when a set of other dispositional states are triggered in U, where these states are grounds for deployment of predicates that apply to properties: predicates like, *colour*, *quality*, *spoken of by O*, etc. Thus in the case at hand, in her state of contemplation, U's NC system may also trigger D_{Colour} . Her state may also thereby generate the activation ground for Γ -[*quality*], and, in conjunction with recognition of the verbal productions of others, activation of Γ -[*is spoken of by O*]. What we find is a *bundle* of activated cognitive states, which are clustered around the activation of D_{Red} . Such a bundle of activated states constitutes contemplation of a property, in this case, redness.

Just as perception of an object, like a red ball, can generate a bundling of Γ -states of the form $^0[\Gamma$ -[*ball*], ...]—we called this an ι -state in §46—we can discern a similar kind of bundling in the case of a property contemplation. Let us call the bundling of Γ -states in a

property contemplation an θ -state. Its nucleus shall be $^{\text{No}}[..]$ —the superscript means non-object. Thus in the case of our contemplation of redness the state generated has the structure:

$$^{\text{No}}[\Gamma\text{-}[colour], \Gamma\text{-}[quality], \Gamma\text{-}[referred to by } O]].$$

This state can itself function as a Φ -state for terms like *that colour property referred to by O*.

Analogously to names, we can then allow that there are memory chains set up, in which such bundles of information enter referential chains. A sentence like *Jane has that property we spoke of yesterday* might be a case where the phrase *that property* has as its Φ -state a θ -property. Some names have θ -properties as their Φ -properties. For example, there are names for properties like *Nirvana*.

This is one mode of non-objectual reference. The other mode is descriptive specification as in a sentence like: *The property that Jane and Fred possess is greed*. This sentence features a descriptive mode of reference based on unique specification of a property. Here what we say is entirely parallel to the proposal about descriptive reference made in §48. This means that the Φ -property underpinning the phrase, *the property that Jane and Fred both possess*, is a modification of an θ -property, which is exactly like the modification that ι -properties undergo to become descriptive ι -properties. Roughly, the primary ground for using the description above will be that U is in a property contemplation state C which generates an θ -property of the form $^{\text{No}}[\textit{possessed by Jane and Fred}..]_x$, and U has the disposition that in whatever property contemplation state C_y , if that state triggers a θ -property of the form $^{\text{No}}[\textit{possessed by Jane and Fred}]_y$, U's system is disposed to token an inter-substitutive Π -property: $\text{IS} \{ ^{\text{No}}[\textit{possessed by Jane and Fred}]_x, ^{\text{No}}[\textit{possessed by Jane and Fred}]_y \}$. In short, the ground structurally is just like that given for object definite descriptions in §48. The only difference is the kind of cognitive bundles involved.

§53 Deconstructing Quine's Quantification: Existentials

We are providing a theory of the language agency of non-object referring terms. We are not providing a *Semantics* of such terms. Semantics is no where to be given. All the explanatory

action is in the cognitive characterisation of the mental precursors of acts using such referring terms. Now we need to spread the net of explanation a little further. At some stage any theory of referring terms has to relate such terms to generality. In standard semantics, generality is dealt with in terms of quantification. The most basic elements of quantification are existential and universal quantifiers. The Quinean (1948) view is that to be an object is to be the value of a variable. If we quantify over Fs, Fs are objects; components of our ontological commitment. One can say that as part of ordinary English we quantify over properties, facts, and propositions. We say: *Every property possessed by O is possessed by N*, *There is a fact that we should bare in mind*, and *Everything said by George is false*. If we allow these generalities to stand as part of our theory of what is, then we are committed to an objectual treatment of properties, facts, and propositions, contrary to what has been argued so far. We must evidently reject Quine's criterion of ontological commitment. I explain how this is tenable by offering a language-agency analysis of *there is*-sentences and universal noun phrases, *every F*. The result implies that we can quantify over non-objects. Being a Quinean naturalist is hard, since to quantify over properties, propositions, or facts, commits one to the existence of these things as entities. Being a global expressivist naturalist is easier, since no such commitment is incurred.

'There is' and 'There exists'

Consider *there is/are* locutions. These locutions take both plural and singular indefinite descriptions as their linguistic arguments: *there is a dog*, *there are dogs*.⁸⁴ We can represent the form of these sentences as *There is/are --N*, where *--N* is a singular or plural indefinite, such as, *a blue cat*, *beer in the fridge*, *some facts you should know about*, etc. In standard semantics such sentences are analysed in terms of existential quantification. *There is an F* is true if and only if some entity O satisfies *x is F*. GE, of course, is not into meaning-analysis. Truth-conditions play no explanatory role in relation to language function. GE's tack is a different one: it explains language agency in relation to *there is* without reference to truth-conditions, extensions, or satisfaction—the tools of Semantics.

⁸⁴ *The there is/are*-locution can also take names and definite descriptions, as in *There is George*. *He might help us*. The analysis provided is consistent with this fact. It does not treat *there is an F* as an idiomatic lump—unlike classical treatments.

There is/are-statements come in two forms. One is an existential form, in which *There is/are --N* has the force of *There exist(s) –N*. The other is an existentially neutral form. An example of such a use is: *There are some things that exist and others that do not*. The classical view is that there is no existentially neutral form of *there is/are*. Quine is the authority here—(1948). There are some renegades who argue for an expansion of our ontology into the reaches of the non-existent. They agree with Quine about ontological commitment, but just want to separate ontology from the existent.⁸⁵ In contrast, GE is quite capable of allowing for both forms, but does not need this idea of expanding ontology into the non-existent since it rejects Quine’s idea of quantificational ontological commitment. I now show how this is so.

The language agency behind both kinds of *there is/are* is basically the same, but let us begin first with existential *there is/are*. To provide this analysis we need to focus for a moment on referring acts and proto-referring acts—in the sense of §44. A referring term is any term *T* whose production is caused by a repertoire $RD[\Psi, T]$, where Ψ is a Φ -state. A proto-referring act is the production of *T* caused by $RD[\Psi, T]$, where *U* may or may not instantiate Ψ . A sincere referring act is an act in which *U*’s term *T* is caused by $RD[\Psi, T]$ and *U*’s instantiation of Ψ . The Φ -property Ψ is a bundle of Γ -states—see §46-8 and §52. So it will have the form $^O[\Gamma-[F], \dots]$ or $^{No}[\Gamma-[F], \dots]$. A class of such Φ -properties now need special notice. Consider terms like *the thing referred to by ‘Pegasus’*. In this case, the Φ -property Ψ has the form:

$$^O[\Gamma-[thing], \Gamma-[referred to by ‘Pegasus’]].$$

The predicates *thing* and *referred to by ‘Pegasus’* are formal predicates connected with dispositions to associate terms with proto-referring acts. Thus an assertion of (i) below simply involves defending a disposition to use *Pegasus* in a proto-referring act. An assertion of (ii) involves defending a disposition to L-process a token of *Pegasus* using a proto-referring act performed with (other tokens of) *Pegasus*:

⁸⁵ Philosophers have recently developed such views about existence and being are Parsons (1980), Routley (1982), and Zalta (1998).

(i) *Pegasus is a thing.*

(ii) *Pegasus is referred to by 'Pegasus'.*

Thus in using the term *the thing referred to by 'Pegasus'*, U instantiates a Φ -state, whose cognitive structure is associated with nothing more than dispositions to deploy terms in proto-referring acts. Let us call such terms *formal referring acts*, and the Φ -properties of such terms purely formal Φ -properties. Such terms are existentially neutral. Whether I believe Pegasus exists, or I am agnostic, I can use the term: *the thing referred to by 'Pegasus'* in a sincere referring act.

By *informal referring acts* I mean referring acts with Φ -properties containing Γ -states that are not purely formal, in the sense just defined. In asserting a *there is*-sentence, U defends a disposition to deploy the descriptive material in the scope of *there is* in sincere informal referring acts. For example, in asserting *There is beer in the fridge*, U defends a disposition to use *the beer in the fridge* in a sincere referring act. That means, U instantiates a repertoire disposition, $RD[\Psi, \textit{the beer in the fridge}]$, where Ψ is a Φ -state that she instantiates, which is not purely formal. Ψ may be a ι -state based in perception of beer in the fridge. Consequently, U is disposed to make some simple subject-predicate assertions, such as *The beer in the fridge is cold*. In order to perform a simple assertion it is not required that the speaker U have any concept of existential commitment or denotation.

That is the existential use of *there is*. When U deploys *there is* in the existentially neutral way, U defends a disposition to use the descriptive material in the scope of *there is* in proto-referring acts. U performs a proto-referring act with a term *T* if and only if U utters *T* and that utterance is caused by a repertoire disposition $RD[\Psi, T]$, where Ψ is some possible Φ -state that U may or may not instantiate. An existentially neutral use of the *there is*-construction is: *There is a thing referred to by 'Pegasus'*. Asserting this sentence involves defending a disposition to use *the thing referred to by 'Pegasus'* in a proto-referring act. Indeed, we are also happy to use it in a purely formal referring act, but not in a sincere informal referring act.

Sentences such as *The thing referred to by 'Pegasus' is Pegasus* are existentially neutral. In asserting this sentence U defends an inter-substitutive disposition in the sense of §27. U defends the state: IS[*The thing referred to by 'Pegasus', Pegasus*], a disposition to inter-substitute speech-acts of the first kind with those of the second kind in all non-opaque contexts. In asserting the negative existential *The thing referred to by 'Pegasus' does not exist* we simply defend rejection of a disposition to deploy the given term in an informal sincere referring act. That's all there is to negative existentials. They are not metalinguistic statements, despite the fact that their Π -properties are metalinguistic—see §27.

That is it. We sum up the results in the table below. Let *the N* simply be the definite description derived from $--N$. The general thesis is:

There is/There are	Π -Property Specification
Existential: <i>There is/are</i> $-N$	Disposition to use <i>the N</i> in sincere informal referring acts, viz., production of <i>the N</i> caused by RD[Ψ , <i>the N</i>] and instantiation of Ψ , where Ψ is informal.
Neutral: <i>There is/are</i> $-N$	Disposition to use <i>the N</i> in a proto-referring act, viz., production of <i>the N</i> caused by RD[Ψ , <i>the N</i>].

This analysis assumes that existential statements carry certain uniqueness implications. To assert *There is (exists) an F* is to commit oneself to performing *the F* in sincere informal referring acts, and correct use of *the F* implies uniqueness. Clearly the analysis as it stands will not correspond to the classical truth-conditions for existentials, according to which more than one thing can be satisfy x is F . I think that this uniqueness implication is a fact about use of indefinites. But I cannot argue for that here—see Barker (2004). We can modify the account to bring it into line with the classical treatment. The modification is that in making the existential assertion U defends use of *the F* or *Fs* in sincere informal referring acts, as in terms like *The cat or cats*. U may deploy this locution where she is ignorant of how many things satisfy *cat(s)*.

Indefinite descriptions in *there is*-statements can be objectual. The descriptive term F in *There is an F* can be that for an object kind, such as *hippo*. However, nothing prevents the descriptive term F being a descriptor of a non-object kind, property, fact, or proposition. Hence, sentences like *There is a property possessed by O but not by N*, do not, in being asserted, commit us to properties as objects. In asserting this sentence U defends a disposition

to deploy *The property possessed by O but not by N* in sincere informal, referring acts. U also defends a related disposition to assert simple sentences like *The property possessed by O but not by N*____. We have seen above how such assertions function—§52. There is nothing ontological in the narrow sense about existential generality, *pace* Quine (1948). There is no imposition of an object interpretation of property-, proposition-, or fact-talk.

In this discussion we have encountered two distinctions. We have a distinction between existential and existentially neutral uses of *there is/are*, and a distinction between objectual and non-objectual uses. These are not the same distinction. Non-objects are not non-existent objects. Non-object talk is not an existentially neutral mode of speech.

§54 Deconstructing Quine's Quantification: Universals

Consider now universal noun phrases of the form *every F*. A sentence of the form *Every F is G* is one whose Π -property underpins the inference form of universal instantiation:

$$\textit{Every F is G. T is F.} \vdash \textit{T is G.}$$

Universal instantiation has two forms: one where the premise *T is F* is asserted and one where *T is F* is supposed. According to GE, the first case is that where the sentence *T is F* is produced in a proto-assertion performed in an assertion. In which case, U is committed to asserting *T is G*. In the second case, *T is F* is produced in a proto-assertion performed in a supposition. In which case, U is committed to performing a proto-assertion of *T is G* in the scope of that supposition, or, as I shall say, *conditionally*. The Π -property of *Every F is G* is then:

Every	Π -Property Specification
<i>Every F is G</i>	Where U is disposed to perform a proto-assertion of <i>T is F</i> , either as assertion or supposition, U is disposed to perform a proto-assertion of <i>T is G</i> , either as assertion or conditional.

I emphasise that what is being specified here is the structure of language agency using universal sentences. We are not specifying truth-conditions. We can derive the usual story we

have been giving about truth-evaluation in terms of **ID-True**—§17. U judges true an assertion of *every F is G* if and only if she sincerely defends a disposition of the kind above, which is that in terms of which she L-processes *every F is G*. There is no embedding problem that we encounter here. In embedded uses of *Every F is G*, U performs a proto-assertion with the sentence—see §15. Thus no problem arises for explaining the language agency of sentences like: *If Every F is G, then no F is not G*.

We should not confuse the present proposal with a substitutional quantification approach (SQ). It isn't SQ. First, SQ gives the truth-conditions of *Every F is G* in terms of substitution cases. That makes SQ metalinguistic, but universal statements are not metalinguistic statements, that is, statements about substitution instances.⁸⁶ GE's theory is not a specification of truth-conditions, so it is not open to this charge.

Secondly, SQ has a problem with non-denumerable domains. The concern is that there are not enough names for all the objects in the domain of quantification. A response is to define truth-conditions in terms of indefinite extensions of the language—see Evans (1977). That means that substitution cases would be names *T* in our existing language or some extension of it. But, of course, this would still be a worry for non-denumerable domains since any such set of names would only ever be a denumerable infinity of names.

Although these are legitimate concerns about SQ they are not a problem for GE. Even if the substitution terms *T* for *Every F is G* were restricted to some finite set of names, or the potential denumerable infinity of names given by extensions of our current set of names, it would not follow that the sentence would simply be about the things picked out by this set of names. The universal claim is not about substitution instances.

Secondly, it isn't just names that can be used in universal instantiation. In suppositional reasoning, we can use a generic term for an arbitrary real, as in:

Suppose *r* is a real number. Then *r* may not be recursive.

⁸⁶ Van Inwagen (1981) claims not to be able to understand substitutional quantification for this reason; it has metalinguistic truth-conditions but universal statements are not meant to be metalinguistic, that is, about language.

How exactly do generic suppositions function? To answer this question we need to understand suppositions. To suppose that T is F is to stipulate the *sayability* of T is F for an argumentative purpose. In GE's terms, this means that in supposing T is F , U proto-asserts T is F , where the purpose of allowing production of T is F is to determine the logical consequences of allowing that production. The proto-assertions that are allowed as consequences of the supposition are speech-acts that are conditional on the supposition.

Recall that proto-assertions are utterances of sentences caused by repertoire dispositions—see §14. So utterance of the sentence T is *real number*, where T refers to some particular thing, is caused by the underlying repertoire disposition of the form below, where Φ_N is the mental element underpinning use of T :

R1: Where U is directed towards producing a sentence signalling $(\Phi_N + \Gamma\text{-}[real])^{GI}$ she can produce a sentence of the form T is a real number.

This repertoire disposition is particular in the sense that $\Phi_N + \Gamma\text{-}[real]$ is a determinate Π -property; U could make an assertion with the sentence T is *real*. But not all repertoire dispositions are particular. Lying behind *R1* is a more general repertoire-disposition. This has the form:

R2: Where U is directed towards producing a sentence signalling $(\Phi_x + \Gamma\text{-}[real])^{GI}$ she can produce a sentence of the form x is a real number.

R2 is the disposition to produce sentences of a certain form x is a real number given a ground-indicating disposition of the form $(\Phi_x + \Gamma\text{-}[real])^{GI}$, where Φ_x is an as yet unspecified Φ -state. *R2* lies behind *R1* in the sense that as U acquires new Φ -states, she generates new particular repertoire dispositions as a result of *R2*.

The sentence r is a real number is a sentence of the form x is a real number. The repertoire disposition lying behind production of r is a real number could just be *R2*. Of course, so caused, such a sentence cannot be asserted, as no determinate reference has been assigned to r . But that does not prevent a certain utterance of r is a real number being caused

by *R2* and appearing in embedded context. One context in which it can appear is a suppositional one.

In a generic supposition, *suppose r is real*, the proto-assertion *U* performs is one of this kind. *U* stipulates the sayability of *r is a real number*. The supposition then amounts to nothing more than this: one is given permission to say *r is a real number*, and to derive other sayable sentences of the form *.... r*, if conceptual principles will allow the derivation. The predicate *real number* is analytically connected to the predicate *corresponds to an infinitely long decimal expansion—which may or may not be recursively specifiable*. If so, we can carry out the suppositional reasoning above. In this suppositional reasoning, the term *r* is empty of content. The term *r* so used is interpretable as employing a pure concept of a thing. This does not mean that *r* refers to some generic object. It lacks a Φ -state so cannot be used to denote something.⁸⁷ With the supposition, it gains content in this sense: we can proto-assert that *r* is a real number.

Non-Objects and Universals

The current treatment of *every* does not require that we only quantify over objects. In *Every F is G* , the *F*s can be non-objects. It may be that number terms are non-object terms. Generally, in *Every F is G* , the predicate *F* may apply to properties, propositions, and facts. An interesting case is Leibnitz's law:

L: If $A = B$, then every property possessed by A is possessed by B and vice versa.

In asserting the generality in ***L***, *U* defends the state:

Where *U* performs a proto-assertion of the form *T is a property possessed by A* , *U* is disposed to perform a proto-assertion of *T is a property possessed by B* , and vice versa.

In this case, what takes up residence with *T* are terms that are nominalizations, ones whose Φ -properties are Γ -properties—terms like *redness*, etc. We have already explained how

⁸⁷ See Fine (1985), who introduces the idea of arbitrary objects. For Fine, *r*, in our suppositional reasoning, refers, according to Fine, to an arbitrary object. An arbitrary object is an abstract entity that ranges over concrete, or lower level abstract entities.

assertions of *Redness is possessed by A* work—see §51. If so, there can be quantification over non-objects. Quantification is not a vehicle of ontological commitment.

The analysis of generality offered by GE is not one in terms of truth-conditions. What we have provided, rather, is a description of the speech-act and cognitive activity that goes into the production of basic sentences of quantification. The key idea is that generality does not bring ontological commitment to entities in a domain of discourse. Evidently, such an approach needs to be developed to cater for issues of scope. However, these matters cannot be dealt with here.⁸⁸

§55 Propositions as Non-Entities

So much for general structures. I have mainly illustrated the idea of non-object terms and non-objects through property-terms and properties. But let us dwell for a while upon those creatures of darkness: propositions. Philosophers have attempted to formulate many theories about what propositions are. I will not explore these theories in detail here.⁸⁹ But, GE's claim is that such theories are based in category error. The assumption is that propositions are objects. That is false. Propositions are those things we refer to with *that*-clauses. But *that*-clauses are sentential nominals—see §45. Propositions are then non-objects. The predicate *is an object* was analysed in §45. So in asserting that propositions are non-objects, we reject the application of the predicate *object* to them, which means we reject a disposition to use propositional referring terms in referring acts whose Φ -states are ι -properties.

Propositions without Nature

This approach to propositions makes our relation to them unproblematic. We do not have to ask how natural beings have capacities to grasp abstract entities like Fregean senses or sets of worlds thought of as vast concrete but spatiotemporally and causally isolated

⁸⁸ See Barker (2004) for a related speech-act treatment of scope issues.

⁸⁹ Some views hold that propositions are extra-pragmatic entities, like sets of worlds—Stalnaker (198), Lewis (1986)—states of affairs Barwise and Perry (1983)—Fregean (1892) senses, etc. Or that they are pragmatic entities like assertions, assertion-types, or utterances—Barker (2004). GE denies that propositions—the bearers of truth—are assertions or proto-assertions. Indeed, it looks like a category error to assign truth or falsity to an *act*—see Macfarlane (2005). We talk of assertions being true all the time, just as we talk of statements being true. But when we speak of assertions being true, what we really mean is *what was asserted* being true.

systems. Generally, in GE, there are no abstract *entities*—properties, relations, facts, and propositions are non-entities. To grasp a proposition is to grasp the meaning of a sentence. That means being able to produce a sentence whose meaning is characterisable through a referring term of the form *that P*. That means being able to entertain that P, and potentially believe or disbelieve that P. These are all folk-semantic platitudes, and there is no deepening of these explanations. The real explanations come with the functional analysis of language agency, where we can provide an analysis of meaning attributions to sentences and understanding-attribution. I provide these explanations in chapter 10.

It might be objected that propositions must have natures because we must think of them as having a compositional character of some kind. For example, the proposition that snow is white is about whiteness and snow. So must we not think of the proposition as somehow made up of snow and whiteness? Similarly, we might think that logically complex propositions contain constituents. The proposition that P and Q must somehow involve the propositions that P and that Q. In both cases we do not have to understand *aboutness* and logical complexity in terms of the constitution of an object. In producing a *that*-clause of the form *Snow is white*, U must perform proto-acts with *Snow* and *white*. In producing *that P and Q*, U produces a proto-assertion corresponding to *P* and *Q*, and *that*-clauses of the form *that P* and *that Q*.⁹⁰ The idea that thoughts are about the world does not require that we think of them as somehow made up of worldly things such as snow and whiteness.⁹¹ In making this claim we are not proposing that propositions are intermediaries between our minds and the world. They are not. Both views, that they are made up of worldly things, and that they are intermediaries, are theories about the nature of propositions, and GE denies all such views. It does not deny them dogmatically, but rather because they all involve the idea that propositions are objects, and GE denies the view that they are objects.

⁹⁰ GE's theory of truth is minimal and nominalistic. Horwichian minimalism is quite ontically profligate in comparison, since it accepts propositions qua entities and a truth-property. The present view is entirely distinct from Schiffer's (2003) view that propositions are *pleonastic entities*. Schiffer seems to think that our assertion of sentences of the form *The proposition that P is true iff P*, entitles us to quantification over propositions.

⁹¹ In his famous exchange with Frege about Mont Blanc, Russell took this line (Frege 1950, p. 169).

The Unity of Thought

The non-objects of which GE speaks are not the kinds of non-objects postulated by Frege (1892a). Frege forged the distinction between concept and object. Concepts are incomplete entities in need of supplementation. Concept terms are incomplete expressions, that is, terms that are entirely distinct from referring terms. Frege thinks he needs this distinction to explain how thoughts are not mere collections of entities. If a concept were an object, then the thought encoded by *T is F* would be a mere pair of two entities. Thoughts are truth-apt, but pairs of objects do not seem to be truth-apt. So the thought cannot be such a pair. Adding more entities to the pair will not bring into being the unity that is the thought, since we will just generate triples, quadruples of entities—more non-thoughts. Frege's solution is to reject the initial assumption that the concept is an object. The constituents of the thought encoded by *T is F* are an object, referred to by *T*, and a non-object, or incomplete object, that is referred to by the rest of the sentence having removed the term *T*. The rest of the sentence is not itself a term. We can gesture towards the rest of the sentence with the notion *...is F*. The dots indicate that this expression is to be completed. It is an abstraction from a whole sentence. The concept has an analogous incompleteness. It is what is left of the thought when we abstract out the object.

The paradox is that concepts as Frege defines them cannot be referred to by singular terms. If a concept could be named, say through a term *C*, then the mere presence of the terms *T* and *C* would suffice for encoding the thought. But they do not. So concepts are essentially unnameable. But if that is correct, what are we doing when we talk about concepts? As Frege put it pithily, the concept of a horse is not a concept. That is because *the concept of a horse*, being a singular term, has the wrong structure to be able to refer to a concept. Only incomplete expressions can do that.

It is not obvious that Frege can escape from the paradox.⁹² It is not obvious to what extent this is a general problem for the analysis of thought. Whatever the answers are, in GE there can be no analogous problem. It is an artifact of an objectual view of predication.

⁹² See Dummett (1981), Priest (1996), and Wright (1998).

Take thoughts first. Thoughts are not objects, and so we cannot talk of constitution. But surely you might object that we can talk of the thought that *T is F* and the pair *T* and *being F*. How do they differ? This question asks for a theory of constitution. But in fact the question is illegitimate. Having understood how the phrases concerned function—*that T is F* and the plural referring term *T and being F*—there is no further question to ask about constitution. The first phrase, *that T is F*, functions as a sentential nominalization—which we have already described. The second, *T and being F*, is a plural referring term, whose constituent terms are an object term *T* and a predicate nominalization, *being F*. The two phrases are structurally distinct. We are not disposed to inter-substitute these terms in all non-opaque contexts. *U* could be disposed to assert *It is true that T is F*. But *U* is not disposed to replace *that T is F* by *T and being F*. The result is not even grammatical. So *U* is not disposed to assert the identity of the proposition that *T is F* and the pair *T* and *being F*.

The problem of unity can also be raised in relation to sentences themselves. Sentences are objects—patterns of marks. The sentence *T is F* features the marks *T* placed in a line with *is F*. The sentence comprises two marks *T* and *is F* in a relation of spatial contiguity. That is what a sentence is. As an account of the constitution of the sentence this might appear inadequate. One might say that we have described the spatial structure of a sentence, but there is more to the structure of a thing than its spatial, temporal, and physical structure. What is this extra ingredient? Some may be tempted to say, in the case of the sentence, that there is a relation of spatial contiguity holding between ‘*T*’ and ‘*is F*’. But how do the relation and these two linguistic marks come together? In answering this question we are in danger of being led down a well-known path—the third man argument—in which we shall discover an infinite metaphysical structure. Or we shall end up postulating unsaturated *gluons*—in Priest’s (1995) sense—that is, entities that are somehow non-entities binding our relational universal spatial contiguity ‘*T*’ and ‘*is F*’.

The right response is to hold the question at the beginning. In saying that the sentence is ‘*T*’ spatially next to ‘*is F*’, or ‘*T*’ in a relation of spatial contiguity to ‘*is F*’, we do not have to see the relation spatial contiguity as a constituent of the sentence. There is no account of what the relation of spatial contiguity is. As a non-object—but not in Frege’s sense—it is a

category error to think that it can be part of the constitution of something. Only objects are parts. Having given the account of how terms for relations function, there is no question left about what relations are. If so, there is no question about the metaphysical structure of the sentence. This is not saying that its metaphysical structure is primitive. It is rather that the idea of metaphysical structure is confused.

We trace the error back to the objectual view of predication—see §4 and §31. This is the idea that in explaining the content of a subject-predicate sentence like '*T* is spatially contiguous to *F*' we shall find some worldly correlate, a state of affairs, which, obtaining, will explain its truth. Embracing that idea, we are then forced to take seriously questions about what the conditions in the world are that make sentences true, since they are part of our explanation of the meaning of sentences. This leads then to metaphysical investigation of the kind illustrated above.

There is nothing wrong with the idea that things in the world make sentences true. We shall explain the language function behind such talk in the next chapter. What is wrong—and this is the error of the objectual view of predication—is the idea that relations of making true are part of the explanation of language function. On the contrary, correspondence relations, being semantic, have no role in explaining language function.⁹³

⁹³ I have been talking as if every *that*-clause with a well formed declarative sentence can be used to pick out a proposition, that is, a truth-apt thing. But we do not have to say this. (Matters here refer back to our discussion of direction of fit—§9 and §15.) *Belief* is an ambiguous term. It can mean the state of belief or what is believed. The term *what is believed* refers to a proposition. Hence we think of belief-states as kinds of propositional attitudes. There is a tendency in the literature to talk of desires as propositional attitudes as well. If desires can have propositions as objects, why can't they be truth-apt? The answer is that they do not have propositions as their objects. This is not to deny that we can use *that*-clauses to refer to what is desired. I can say: *Jane desires that Fred comes tomorrow*. Is this not a proposition? No. Propositions are truth-apt, but this is not. Why not? Because in producing the *that*-clause *that Fred comes tomorrow* we are associating the proto-assertion of *Fred comes tomorrow* with a non-assertoric purpose. In desiring that Fred will come, U has a disposition to try to bring about circumstances that could be described as Fred's coming tomorrow. This state is not a disposition to defend the Π -property of *Fred comes tomorrow*. When U consciously expresses her desire to herself, she tokens *Fred comes tomorrow* as a proto-assertion, using it to signal to herself her state of desire. Call such acts *optative acts*. Optative acts are a form of illocutionary act. Like all illocutionary acts there is a proto-act lying behind them. Call this a *proto-optative act*. In attributing a desire to someone by asserting *Jane desires that Fred come tomorrow*, U performs the *that*-clause content sentence, *Fred comes tomorrow*, in a proto-optative act.

§56 *Objects, Non-Objects, and Natures*

The real distinction between objects and non-objects cannot be analysed by giving a theory of objects, and then a theory of non-objects, and showing how they are different. That would be to make the real distinction between objects and non-objects a metaphysical distinction. The distinction between object and non-object is not a metaphysical distinction. It is empty of any nature. Rather this distinction is illuminated—though not truth-conditionally analysed—through study of the language agency and of object and non-object talk.

The structural difference between object-terms and non-object terms is this—§45:

- (a) The kinds of things that are the potential referents of object-terms play an essential explanatory role in characterising the Φ -states of these terms;
- (b) The kinds of things that are the potential referents of non-object-terms play no essential explanatory role in characterising the Φ -states of these terms.

The reason for (b) is that the Φ -states of non-object terms are borrowed from those of predicates and sentences, and the mental elements of predicates and sentences are not assigned worldly correlates as part of the explanation of their functional natures. If we are fully aware of this feature of language agency, then we cannot seriously pursue questions about the nature of properties or propositions since there is nothing to discover.

The thing about objects is that we can always ask what their natures are. What is it to *have a nature*? One might say this: there is something that it is to be that thing. *Something* here introduces quantification over properties. To have a nature is to instantiate properties that make something a particular kind of thing. Or, in other words, if something has a nature it is essential to the being of that thing that it has properties that individuate it and distinguish it from other things. Any way of thinking about a thing with a nature will always require a mental element that treats the thing as a subject of predication. The necessary functional backbone of thinking about a thing is tokening a Φ -state. So, any way of thinking about a thing with a nature will require using a Φ -state that has predication built into it. The kinds of Φ -states that meet this condition are ι -properties. They are cognitive bundles of Γ -states, predicate mental-elements, bound by concatenation—see §46. Concatenation is the mental

element behind predication. So ι -properties have predication, or the cognitive precursor thereto built into them. Thus, any way of thinking about a thing with a nature will always require an ι -property. We can now propose a language-agency analysis of *x has a nature*. This predicate is a formal predicate in the sense of §26-27. To say *O has a nature* is to defend a disposition to use *O* in a proto-referring act with an ι -state, the mental element of an object-term. This hypothesis explains why we intuitively link talk of objects with talk of things that have natures. Metaphysics, of course, pursues questions about the natures of things like propositions and properties—non-objects from the point of view of GE. If one has any residual feeling that there is something peculiar about these questions, it is, I conjecture, the cognitive echoes of the language agency of *nature*.

Although objects have physical natures they lack metaphysical natures. I just tried to clarify this idea with the case of sentences qua physical tokens—see §53. In claiming that objects are bundles of property-instances that are spatiotemporally and causally unified, metaphysics is compelled to ask: what are properties? What is instantiation? But GE cannot allow such questions since they ultimately presuppose objectualism about predication. There is no theory of what a property is. Nor a theory of what it is for a property to be instantiated.⁹⁴

Perhaps there might be some doubt about this last point. We make claims such as *Red is instantiated here*. What is the structure of the Γ -state that underpins use of the verb *instantiate*? Our task is to characterise a functional structure in ways that are familiar in terms of *basing*—see §16. Let D_{Inst} be the Γ -property underpinning *instantiates*. That state has a certain combinatorial capacity. It can combine with the Φ -property of a predicate nominalization and a Φ -state of a term *X* for a spatiotemporal region, as in: Φ -[Redness] + D_{Inst} + Φ -[*X*]. Its canonical ground is the activation of a generic Γ -state, say D_{Red} , in conjunction with activation of spatial relation detectors. The system of spatial relation detectors is a hugely complex system that cognitively relates perceptual triggerings of generic Γ -states or *on-states*. In seeing a red patch next to a blue patch, U's system cognitively relates

⁹⁴ Why can't numbers or other purely abstract things be objects? There is a very large issue here that I will not be able to tackle in this book. My conjecture is that they are all non-objects. Number terms are based in nominalizations. I cannot argue for this view here.

an on-state of D_{Red} to an on-state of D_{Blue} within the spatial relation detection system. Note that we are not saying what space is or what it is for a red patch to be next to a blue patch. We are just theorising, speculatively, about the structure of spatial-relation detection systems. Any specification of a spatial position will deploy this system. So let us suppose that a perceptual episode involves the triggering in conjunction of Γ -[Red] and the mental element underpinning use of *here*. We can represent this as $\{\Gamma$ -[Red]-On, Φ -(*here*)-On $\}$. We have then the following C-law:

$$\{\Gamma$$
-[Red]-On, Φ -(*here*)-On $\} \Rightarrow \Phi$ -[Redness] + D_{Inst} + Φ -(*here*).

Such C-laws give the basic functional character of D_{Inst} . We can now analysis the referring term *instantiation* as a nominalization of *instantiates*, where the mechanism of nominalization is as we have described it above. Having provided this functional analysis of the basic language agency of *instantiates* and *instantiation*, there are no theoretical questions left about the nature of *instantiates*. Or rather, we should see the classical metaphysical conundrum, *What is the nature of instantiation?* as based in a category error.

Chapter 8

Truth-Making and Facts as Non-Objects

§57 Truth-Making and its Discontents

Any adequate theory of human language activity needs to explain why the T-schema, **T** below, is intuitively correct:

T: $\langle S \rangle$ is true if and only if S .⁹⁵

For all sentences S we are capable of comprehending, we are disposed to assert corresponding instances of **T**. GE explains this disposition—§20.

The T-schema indicates a correspondence between truth and being—between truths and how things are in relation to the subject matter of the truths. But another kind of schema, less discussed, but equally important to our understanding of truth, indicates another kind of relation between truth and being. We might speak of a *truth-making schema*—or **TM-schema** below, which is restricted to truth-bearers $\langle S \rangle$ that are true:

TM: $\langle S \rangle$ is true because S .

TM-sentences have the same seemingly transparent rightness that T-sentences do, for example: $\langle \textit{Snow is white} \rangle$ is true because *snow is white*; $\langle \textit{There are no unicorns} \rangle$ is true because *there are no unicorns*, and so on. Related schemata that feature equivalent sentence-types—which are again restricted to true propositions $\langle S \rangle$:

TM1: $\langle S \rangle$ is true in virtue of the fact that S .

⁹⁵ In what follows I use the notation $\langle . \rangle$ as a nominalization device. That is, $\langle S \rangle$ has the same function as *that S*. Hence we can say $\langle S \rangle$ is true for *that S is true*. This notion makes the presentation of ideas a little simpler. I note that the T-schema **T** above is not strictly speaking assertable. It does not take into account *conventional implicature* whose presence in a language undermines the disquotational properties of truth. See Barker (2003).

TM2: The fact that *S* makes $\langle S \rangle$ true.

A theory ought to explain the intuitive rightness of **TM** and **TM1/2** and also explain the intuitive wrongness of their converses. We are not disposed to assert instances of *S because it is true that S*, or *S in virtue of its being true that S*, as in: *Snow is white in virtue of the truth of \langle Snow is white \rangle* , or *What makes it the case that snow is white is its being true that snow is white*.

TM1 and **TM2** bring in reference to facts. Fact-talk is about as common and unproblematic as truth-talk. **F** below captures one of our central commitments about facts:

F: It is a fact that *S* if and only if *S*.

Where one asserts *S* one is disposed to assert *There is a fact that S*, and vice versa. Assertions of *It is a fact that S* and $\langle S \rangle$ is true are closely related, but not obviously identical in nature. For example our assertion of **TM2** and denial of its converse indicates a difference.

T, **TM** and its variants, and **F**, sum up some of the core conceptual commitments of truth and its close conceptual entourage. Just as it would be deeply revisionist to dispute **T**, it is revisionist to dispute **TM** and **F**. What I want to show now is that orthodox semantics makes heavy weather out of preserving these commitments. The reason for this is that orthodox semantics has a tendency to reify facts. That reification is the result of its objectual view of predication, and relatedly, its embracing of Quine's quantificational principle of ontological commitment. In contrast, GE has no problem in formulating a theory of talk of facts that preserves all these commitments. This chapter demonstrates that although GE makes no use of representation and correspondence in its account of language agency, GE still respects our talk of truth's fixation by being and thoughts potential to correspond to fact.

§58 Reifying Facts

Many philosophers hold views inconsistent with the commitments just described. If we accept the identity theory of truth—the truth that *P* is simply the fact that *P*—truth is essentially primitive, so **TM** cannot be valid. But I think identity theories of truth must be

false for that reason, and indeed, are so for other reasons.⁹⁶ More interesting are those philosophers rightly impressed by the idea that truth is grounded in being. But they argue that for truth to be grounded in being, being must be seen as a plenitude of facts, and facts must be entities. The following argument—drawn from Rodriguez-Pereyra (2005)—gives the idea. How do **TM**-sentences, such as *<Rosy is red> is true because Rosy is red*, function? ‘Because’ is a sentential connective. So one might see it as signifying an operation on propositions. But *A because B* cannot convey merely that an entailment holds between *B* and *A*, nor a counterfactual relationship, since entailment and counterfactual dependency goes both ways in the case of *<P> is true* and *<P>*. What we need to acknowledge is that **TM**-statements are explanatory statements; they can be used to answer why-questions. Explanations only function as explanations because they indicate relations of asymmetric determination holding between things in the world. So **TM**-sentences like *<P> is true because P* must be about a relation, call it *grounding*—not to be confused with the grounding we have spoken of in earlier chapters—which connects things corresponding to *<P> is true* and *<P>*. The obvious thought is that these things are facts. But here a principle about relations and entities kicks in, which is obvious to Rodriguez-Pereyra:

REL: Relations relate entities.

Evidently, **REL** is entirely at home in the world of the objectual construal of predication. It is a special case of the principle that properties apply to entities. The ground for that principle is the objectual treatment of predication. Predicates apply to entities, but properties correspond to predicates, so properties are instantiated by entities. Given **REL**, facts must be entities.

It might seem that the ontic understanding of facts delivered by **REL** vindicates our commitments to the **T**-, **TM**- and **F**-schemas. But that is far from the case. The ontic conception commits us to a panoply of logically complex entities: negative, disjunctive,

⁹⁶ One form of the identity theory is that truths are analysed as worldly states of affairs—Hornsby (1997). What of false thoughts? If false thoughts are non-identical to reality, then thoughts that are true and thoughts that are false differ in their intrinsic natures. If false thoughts are parts of reality as negative states of affairs—we accept an identity theory of falsity—then we cannot understand negation as a compositional operation on truth-bearers. Another form is this: we analyse facts as truths—Dodd (2000). But then how things are, the facts, would be fixed (and explained) by what thoughts are true. But surely that gets the explanatory priority wrong.

universal facts, and that offends intuitions honed by Occam's razor. Austerity programs ensue to minimize the set of facts we need to fix truth. Armstrong (2004) argues that we cut down our fact-commitment to positive atomic facts and one totality fact. There are two problems with Armstrong's program. One is that we deny the validity of the **TM**-schema and the **F**-schema. Unproblematic **TM**- and **F**-sentences, below, cannot really be assertable, because there are no individual negative facts corresponding to negative truths in Armstrong's program:

<There are no unicorns> is true because there are no unicorns.

It is a fact that there are no unicorns if and only if there are no unicorns.

Without doubt, Armstrong would claim he is not concerned with ordinary conceptual commitments, since he is doing metaphysics. But this position is a tenuous one. It implies that most of our fact-talk is massively in error. We need then to explain why natural-language users systematically embrace this massive error, and how they manage to be systematic despite their error. Moreover, the imputed error does not look like one; the fact that there are no unicorns is surely a fact. Alternatively, there is some analysis of **TM**-statements that reveals their legitimacy as an ontologically benign mode of discourse. But that undermines the whole metaphysically loaded interpretation of facts.

Secondly, even if we could live with Armstrong's program, totality facts are still negative facts and worldly, metaphysical negativity impresses us with its deep strangeness—as Molnar (2000) vividly shows.

We may hope to escape the clutches of the ontic conception of truth-making by granting that facts are entities, but by somehow downgrading their metaphysical weight. Two weight-reduction methods are minimalism and fictionalism. But they fail because although they make us feel happy about facts, they do not do justice to their truth-making role. Minimalism about facts would propose that all there is to our concept of fact is given by our disposition to assert all instances of the fact-schema **F**. But then there is little to distinguish fact from truth, given that the **T**-schema fixes truth, and so **TM**-sentences fail to be certified. Fictionalism about facts seems to have similar problems to minimalism. We might fashion

the idea that talk about facts and truth is fictional. But what grounds fictional talk about truth and fact is talk the corresponding worldly talk. That is, commitment to the fictions *It is a fact that P* and *It is true that P* is just indirect commitment to P. But it is then hard to see how fact can explain truth, since both depend equally on how things are.⁹⁷

We might resist truth-making altogether, or claim that only some sentences, simple atomic sentences, have truth-makers.⁹⁸ But the problem with this is, again, **TM** and **F** go by the wayside—we would find ourselves denying claims like those above—and the truth of logically complex propositions is left unexplained.⁹⁹

It might appear that there is no way of making sense of what look like central commitments we have about truth, summed up in **TM** and **F**. That suggests that our concept of truth is somehow deeply in error. But how could commitments of such an apparently simple and transparent kind involve such incoherence? I suggest that the appearance of incoherence is generated by tacit acceptance of the objectual understanding of predication (§3 and §31). This leads us to reify the relata of relations and issues directly in acceptance of **REL**, which is the cause of all our woe. GE provides that framework in which we escape from the objectual treatment of predication, and consequently from **REL**.

§59 *Facts Freed: Not all Relations Relate Entities*

The idea that facts are entities delivered by **REL** runs in the face of what we have already said about facts. As proposed in chapter 7, referring terms like *the fact that S* are sentential nominals. Facts are not entities in any sense at all; it is a conceptual error to apply the predicates *object to fact*. (Strictly speaking, the claim that facts are not entities, is not a part

⁹⁷ Fictionalists might propose that **TM**-statements like: *<P> is true because P* are just fictions in turn. But what non-fictional facts ground such claims?

⁹⁸ See Simons (2005) as an instance of someone who denies that negative truths need truth-makers. Fear of the negative leads Mumford (2006) to deny that there are any negative truths. Instead of negative truths, there are falsities: *<not-P> is true* should be replaced in our talk by *<P> is false*. This only works if there are no false-makers. But there is no false-making only if falsehood is primitive, and that requires accepting an identity theory of falsity.

⁹⁹ Lewis (2003) offers a theory according to which truth-makers of *<Rosy is red>* is Rosy qua red. Rosy qua red is Rosy taken under a certain counterpart theoretic constraint. The problem with this approach is that it does not capture asymmetry. If Rosy qua red had not existed, *<Rosy is red>* would not have been true. But similarly if *<Rosy is red>* had not been true, Rosy qua red would not have existed.

of GE. Rather GE explains why, with a certain clarity of mind, we are compelled to assert it.) If so, **REL** must be false.

Following Rodriguez-Pereyra, we cannot accept that *because* in **TM**-statements merely indicates an entailment between S and $\langle S \rangle$ is true. That won't get the vital asymmetry we require; namely, that $\langle S \rangle$ is true because S , and not S because $\langle S \rangle$ is true. But there is an option close to the entailment view that does get us the asymmetry we require. Entailment has its syntactic correlate in derivation. Derivations involve introduction and elimination rules. Think now of the metaphor behind *making* as it appears in *truth-making*. In *making* we construct something. Introduction rules are linked to construction in the sense that they reveal the canonical grounds for use of a logical constant. Let us liberalise the idea of introduction-rule to that of any inference whose premises are canonical grounds for the application of a concept, be that concept a logical constant or non-logical concept. Behind introduction rules are what we might call introduction C-laws. If $S_1, \dots, S_n \vdash R$ is an introduction rule for U then there is a C-law $\Pi_1, \dots, \Pi_n \Rightarrow \Psi$, where Π_1, \dots, Π_n and Ψ are the Π -properties of $S_1, S_2 \dots$ and R respectively, such that either:

- (a) $\Pi_1, \dots, \Pi_n \Rightarrow \Psi$ is a basing law.
- (b) $\Pi_1, \dots, \Pi_n \Rightarrow \Psi$ is such that in order to instantiate Ψ one must process elements in Π_1, \dots, Π_n , but not vice versa.

An example of (a) is the basing law for the truth-predicate—see §20:

$$[\text{Disposition to defend the } \Pi\text{-property of } S] \Rightarrow \Phi\text{-}(N) + D_{\text{True}}.$$

Because of this basing law, $S \vdash N$ is true functions as an introduction rule for U . An example of (b) is the C-law introducing the Π -property for conjunction—see §22:

$$\Pi_1, \Pi_2 \Rightarrow C[\Pi_1, \Pi_2]$$

The cognitive asymmetry here is that to process $C[\Pi_1, \Pi_2]$, the NC system needs to process Π_1, Π_2 , but to process Π_1 and Π_2 , the system does not have to process $C[\Pi_1, \Pi_2]$. For this

reason $S_1, S_2 \vdash S_1 \text{ and } S_2$ is an introduction rule. Let us note that neither of the following inference forms are introduction rules for U, assuming normal functioning of her NC system:

$$N \text{ is true} \vdash S \quad S_1 \text{ and } S_2 \vdash S_1$$

Both correspond to C-laws that fail to meet either condition (a) or (b) above.

The term *because* can have causal and non-causal predicates. In TM-sentences it has a non-causal reading. Our proposal is that, on that reading, *A because B* is used to express a disposition to derive *A* from *B* using, what are for the speaker, introduction rules. Or in other words, it is a disposition to cognitively derive the Π -property of *A* from that of *B* deploying only introduction C-laws. The hypotheses in play are then:

Because	Π -Property Identical to or based in:
<i>P because S₁, S₂...</i>	(i) Disposition to defend: Π -[<i>P</i>], Π -[<i>S₁</i>], Π -[<i>S₂</i>]. (ii) Disposition to derive <i>P</i> from { <i>S₁</i> , <i>S₂</i> ...}, using what for U are valid-introduction rules.
<i><R> is true because S₁, S₂...</i>	(i) Disposition to defend: Π -[<i><R> is true</i>], Π -[<i>S₁</i>], Π -[<i>S₂</i>]. (ii) Disposition to derive <i><R> is true</i> from { <i>S₁</i> , <i>S₂</i> ...}, using what for U are valid-introduction rules.

I submit that with this analysis of the language agency of *because*, we explain our disposition to assert instances of TM-sentences, *<S> is true because S*, but deny their converses, *S because <S> is true*. This follows since for us, given the cognitive criteria specified above, the introduction and elimination rules for *true* are: $S \vdash \langle S \rangle \text{ is true}$ and $\langle S \rangle \text{ is true} \vdash S$. The present hypothesis also explains our intuitions about more interesting forms of truth-making claims, such as *<P \vee Q> is true because P*, *< $\neg\neg$ Q> is true because Q*, etc. We say more about these below.¹⁰⁰

The whole orientation of the present analysis is to supply an explanation of the compelling nature of truth-making principles like TM-, TM1-, TM2-schemata and the F-

¹⁰⁰ On the view being presented, truth-making is a doctrine about truth, not a doctrine that can be merely expressed succinctly using the truth-predicate *pace* Lewis (2001). Lewis thinks that the central claim behind *<P>* is made true by something, is simply that the existence of the entity entails *<P>*. Thus one can sum up all that is interesting in claims about truth-making without using the truth-predicate. But this way of looking at things does not come to grips with the asymmetric dependence of truth on being.

schema. (This is in parity with our treatment of logical inference in §21-22 and folk semantic principles in §76.) To continue with this analysis we need to analyse our talk of facts.

§60 *Fact-Talk and Nominalization*

Facts are non-entities in GE's sense of the term—§45. It is a conceptual error to apply the predicate *object* to any fact. Facts are not true propositions. It is a conceptual error to apply the predicates *true* or *proposition* to any fact. We gain illumination about these claims not by providing a theory of facts, objects, and propositions, but by analysing the language agency of the terms *that S*, *the proposition that S*, and *the fact that S*. We proposed in §50, that *that S* is a sentential nominal; it is a referring term whose Φ -state corresponds to the mental element of a proto-assertion, that is, a repertoire disposition $RD[\Pi^{GI}, S]$. This means that in producing *that S* one performs a proto-assertion of *S*, but with a grammatical modification allowing it to function as a subject term in predications. In using *the proposition that S*, U performs basically the same act. In using *the truth that P*, U refers to a proposition—U performs the same referring act she performs with the proposition that P—but U implicitly asserts the truth of the proposition by using *truth*. Exactly the same thing goes on with: *the proposition that P, which is true*.

What then of the complex noun phrase *the fact that S*? Two data need explaining. The first is that, in ordinary thought, facts are not propositions. Facts make propositions true. That a proposition is true is not a brute fact. It needs explaining. That explanation is given by how things are, that is, the facts. So, the mental precursor of *the fact that S* must be cognitively distinct from that of *the truth that S*. The second datum is that like *the truth that S*, use of *the fact that S* in referring acts generates a commitment to *S*. So in asserting *The fact that Winter was coming did not concern Napoleon* we are committed to asserting *Winter was coming*. A theory that explains the facticity of *the fact that S*, and its distinctness from *the truth that S*, is that the Φ -property of *the fact that S* is just Π^{GI} . In other words, in performing a referring act with *the fact that S*, U performs an assertion—U signals her possession of a ground-signalling disposition corresponding to *that S*, which is Π^{GI} —but the verbal sign embodying

this assertion has the grammar of a referring term. In other words, in uttering *the fact that S* and *the truth that S*, U performs proto-assertions of *S* in both cases, and defends Π in both cases, it is just that the Φ -property of *the fact that S* is the disposition to defend Π , whereas the Φ -property of *the truth that S* is just the repertoire disposition $RD[\Pi^{GI}, S]$. We sum up these differences as follows:

Noun Phrases	Φ -state
<i>That S</i>	$RD[\Pi^{GI}, S]$
<i>The truth that S</i>	$RD[\Pi^{GI}, S]$, (and U signals through <i>truth</i> a disposition to defend Π).
<i>The fact that S</i>	Π^{GI}

Referring acts using *the fact* comprise assertions—the fact-locution refers to what the assertion is about. But not all uses of *the fact that S* are referring acts. We can also use *the fact that S* in pure proto-referring acts. Such a use is found in the conditional, *If winter had been coming, the fact that it was would not have bothered Napoleon*. Here U deploys the fact-locution in a pure proto-referring act, and there is no factive commitment.

On this view, facts are not thoughts, propositions, or sentence meanings. Nor are they sentences. Nor are they states of affairs, qua world-parts, which is to say, entities constituted by objects, properties and relation in combination.¹⁰¹ Nor are they fictions. Facts are non-objects. Like propositions, there is no metaphysical question about their constitution. To pose such questions is to misunderstand the nature of fact-talk. As we proposed in §45, *object* is a formal predicate associated with a Γ -state based in a metalinguistic disposition. Likewise, *proposition* and *fact* are formal predicates based in another kind of metalinguistic disposition. The Π -property specifications are given thus:

¹⁰¹ Strawson (1950a) affirms that facts are not entities, but does not really explain what that amounts to. He suggests that they are not in the world and are mere pale shadows of sentences. Neither of these claims follow from the view being developed here, and in themselves they are but metaphorical gestures at a theory.

Object/Fact/Proposition	Π -Property Specification
<i>N is an object</i>	Disposition to deploy <i>N</i> in a proto-referring act with individual ι -property.
<i>N is a fact</i>	Disposition to deploy <i>N</i> in a sentential nominal act with Φ -state Π^{GI} .
<i>N is a proposition</i>	Disposition to deploy <i>N</i> in a sentential nominal act with Φ -state of the form $RD[\Pi^{GI}, S]$.

As with all formal predicates—see §27—the Π -properties concerned have a metalinguistic dimension, but the predicates themselves are not metalinguistic. The claim made using these predicates are not claims about language. Furthermore, this analysis has no anti-realist implications. To assert that *O* is an object or fact does not entail that *O* is being referred to, or could potentially be referred to. As no theory of facts is being offered, no such claim about facts in relation to our acts of referring could follow. The theory above predicts that we are inclined to assert that facts and propositions are not entities. That is, we are disposed to reject a disposition to deploy *the fact that S* in a proto-referring act whose Φ -state is an ι -property. Here goes: facts and propositions are not entities. Although both facts and true propositions are non-objects, they are distinct from each other—*pace* Dodd (2000). Distinct speech-act structures are assigned to *the fact that S* and *the truth that S* proposition. The first has as its Φ -state Π^{GI} . One might say, a referring act using *the fact that S* is just an assertion with the syntax of a referring term imposed on it. In referring with *the truth that S*, *U* performs a proto-assertion with the syntax of a referring term imposed on it, but also signals a disposition to defend a Π -property with the structure: Φ -[*that S*] + D_{True} . The difference between referring acts using *the fact that S* and *the truth that S* is where the nominalization occurs. In one case it is applied to an asserted sentence, in the other to a proto-asserted sentence. One might respond: *So is this the great difference between facts and true propositions, a grammatical twist?* The reply to this is that we are not giving a theory of the difference between facts and true-propositions. There is no theory of the difference.

Locutions and Predications

With this theory of talk of facts and truths we can explain certain rudimentary linguistic data. First, we cannot predicate truth—in the sense of the property—of facts. Take:

The fact that snow is white is true. This is slightly odd. The reason is that is that truth-bearers are propositions but facts are not. Facts are not sayables. We cannot assert: *Granny said the fact that snow is white.* The verb *say* takes sentential nominals whose Φ -properties are proto-assertions. The speech-act performed by *the fact* is not a sentential nominal—though it has a constituent that is. Just as facts cannot be said to be true, truths cannot be said to be facts. So, we have the defectiveness of: *The truth that snow is white is a fact.* The defectiveness of this sentence is explained by the table above. We are not disposed to use *the truth that snow is white* in an assertive nominalization, whose F-property is Π^{GI} . If so, truth cannot be fact.

Facts: Predication and Quantification

Facts are non-objects. Yet we apply predicates to non-objects. That view can only be maintained by renouncing the non-objectual view of predication—§3 and §31—according to which predication is constituted by object-property instantiation. Casting aside that view as we did in §51, predications of the form below are seen as totally unproblematic:

FP: The fact that Bob went to Berlin is concerning.

In asserting **FP**, U defends the Π -property based in:

Ψ : Being apprehensive as a result of contemplating the assertion that Bob is in Berlin.

Thus, in terms of **ID-True**—§17—an audience H judges U's utterance of **FP** true just in case H judges that she, H, defends Ψ in her own case. Defending Ψ requires that U actually be prepared to assert that Bob is in Berlin. Assertions of the form *The fact that S is F* are about the fact that S, which is to say, about what assertion of S is about. Thus **FP** is about the fact that Bob went to Berlin. But, of course, the Π -property specification makes no reference to the fact: aboutness in assertions does not go with the content of the Π -property specification.¹⁰²

Similarly, we have no problem with quantify over facts. We can assert statements like: *There are facts you should be aware of* or *Every fact makes true a truth*. But we have

¹⁰² Dodd (2000) makes much of such predications for his argument that facts are objects, for him, true propositions. If one works with the resources of Semantics, then this line of argument has cogency.

already argued in §53-4, that generality carries no ontological commitment in the sense that the things quantified over are entities. Quine is wrong.

A Plenitude of Facts

In GE, for every truth there is a fact. Thus for arbitrarily logically complex truths, there are arbitrarily logically complex facts. There are negative, universal, disjunctive, etc. facts. In terms of an ontic concept of truth, such commitments would be a form of ontological profligacy, which Occam's razor would threaten. But in GE, facts are not entities. There is nothing problematic about logically complex facts of all kinds. If this is the case, nothing stands in the way of explaining our disposition to assert the validity of the **TM**-schema. For every *S* that is judged true we can assert *<S> is true in virtue of the fact that S*. Assertion of the validity of the **F**-schema is similarly secured.

The Slingshot

Facts are non-objects, but that does not mean that we cannot talk of the existence of facts and the identity of facts.¹⁰³ We can say that the fact that Fred hit the ball is the same fact that the ball was hit by Fred. There must be then principles of fact identity. But fashioning such principles seems to be fraught with danger. The Slingshot argument purports to demonstrate that from relatively anodyne principles of fact identity we can derive outrageous conclusions:

L1: Logically equivalent sentences correspond to the same facts;

L2: If two true sentences *S* and *R* differ only by substitution of co-referring expressions, then *S* and *R* correspond to the same fact.

It can easily be shown that any two true sentences correspond to the same fact. Here is the argument—see Davidson (1969) and Neale (1995):

(a) Snow is white.

(b) [Socrates = the thing that is Socrates and Snow is white]

¹⁰³ In GE there is nothing about the language agency of the identity predicate that requires that the words flanking the identity sign have to be object-referring terms—see §27. They do not even have to be referring terms, as we shall see in §75.

(c) [Socrates = the thing that is Socrates and Grass is green]

(d) Grass is green.

According to the up holders of the slingshot, in this form, (a) is logically equivalent to (b), (c) to (d)—so these pairs correspond to the same fact. (b) differs from (c) only by substitution of co-referring expressions. If so, (a) and (d) correspond to the same fact. But we would not want to say that. In particular we would not want to say that what makes (a) true is what makes (d) true.

The argument is dubious in relation to L1. Why would we accept that logically equivalent sentences refer to the same fact? On that view *Snow is white and either grass is green or not* corresponds to the same fact that *Snow is white* does. Only someone who accepts the thesis that logically equivalent sentences mean the same would be inclined to accept L1. But it is not obviously a folk semantic principle that logically equivalent sentences mean the same. The Π -properties defended in (a) and (b) are distinct.

That still leaves us with a concern about the identity conditions of facts. I propose that two true sentences correspond to the same fact if and only if they are logically equivalent and about the same things. *About the same things* needs clarification. Basically, if U asserts that S and R are about the same thing, U defends a disposition to L-process S and R using the same basic Π -properties or constituents thereof; that is, Φ - and Γ -states. Distinct Π -property constituents are needed to L-process (a) and (b) above—*Snow is white* and *Socrates = the thing that is Socrates and Snow is white*. That is the basis of the judgement that the sentences (a) and (b) are not about the same thing.

§61 Truth-Making and Analytic Recipes

We now return to the issue of how **TM**- and related statements function. We can now apply our theory of fact-talk. My suggestion above was that the Π -properties of **TM**-statements are based in disposition to derive $\langle S \rangle$ is true from S, (along with a disposition to defend Π -[$\langle S \rangle$ is true] and Π -[S]). We now apply a similar kind of analysis to sentences of the form: *The fact that R makes true $\langle P \rangle$* , and *$\langle P \rangle$ is true in virtue of the fact that R*. They are likewise

simply statements involving defence of exactly the same kind of Π -property. The only difference between these statements—**TM1**- and **TM2**-statements, and **TM**-statements—is syntactic. This syntactic difference is that we use the fact-locution. It functions exactly as we described it above. So in *the fact that R makes true <P>*, U (a) uses *the fact that R* in a referring act and so defends Π -[R], and (b) defends a disposition to derive *<P> is true* from R using what for U are introduction-rules. We can analyse, syntactically, the predicate, *makes true <P>* as a two-place relational term, *makes true*, combining with a referring term for a proposition: *<P>*. However, the relational term, *makes true*, contrary to **REL**, does not relate entities. It relates non-entities. We are familiar with the things related: they are facts and propositions. We have already commented on their non-objectual status.

We can see truth-making as a species of a genus of *making*, called *analytic making*. Other candidates include *making the case*.¹⁰⁴ So Fred's being unmarried and male makes it the case that he is a bachelor. Jane's sharing memories with an earlier temporal person stage makes it the case that she is identical with that earlier being. The fact that *P* makes it the case that either *P* or *Q*. This suggests a distinction between analytic and *synthetic making*. Synthetic making, one might conjecture, includes causation and physical necessitation, both of which can feature in explanations. I proposed briefly in §33 that the language agency underpinning *cause* is to be understood in terms of recipes of manipulation. We can see the language agency of *analytic making* as a kind of manipulation theory, the recipes being defined by what speakers take to be introduction rules.

Formal Truth-making

Let us take our canonical form for truth-making statements to be that below where $f[S]$ abbreviates *the fact that S*:

$f[S]$ *makes true* *<P>*

We shall allow that more than one fact can contribute to the truth-making of a truth-bearer, *collective truth-making*, as in:

¹⁰⁴ We may want to analyse *makes true* into *makes* and *true*. We could see something like adverbial modification occurring here. The verb *makes* can be adverbially modified by *the case* to form *makes the case*.

$f[S], f[R], \dots$ makes true $\langle P \rangle$

So far we have only dealt with simple kinds of truth-making statements, like: $\langle S \rangle$ is true in virtue of the fact that S . There are less trivial sounding cases of the form $\langle S \rangle$ is true in virtue of the fact that R , where $\langle S \rangle$ and $\langle R \rangle$ are distinct in content. What interests me now is the question: for logically complex S , what facts, apart from $f[S]$ can be said to make $\langle S \rangle$ true? I shall deal with this question by going through cases: negations, conjunctions, disjunctions, existentials, and universals.

§62 Truth-Making and Logically Complex Truths

If we assimilate the standard introduction and elimination rules for conjunction and disjunction, in combination with the introduction rule for truth, then we ought, assuming normal functioning of NC systems, to assert:

$f[P]$ makes true $\langle \neg\neg P \rangle$
 $f[P], f[Q]$ make true $\langle P \ \& \ Q \rangle$
 $f[P]$ makes true $\langle P \vee Q \rangle$
 $f[T \text{ is } G], f[T \text{ is } F]$ makes true $\langle \text{Some } F \text{ is } G \rangle$.

Intuitively, these seem right if we consult our sense of explanatory order. For example, disjunctive truths are true in virtue of their disjuncts. A truth of double negation is true in virtue of the simple fact. On the other hand, we exclude cases like:

$f[\neg\neg P]$ makes true $\langle P \rangle$
 $f[P \ \& \ Q]$ makes true $\langle P \rangle$
 $f[P], f[Q]$ make true $\langle P \vee Q \rangle$
 $f[P \vee Q], f[P \rightarrow R], f[Q \rightarrow R]$ makes true $\langle R \rangle$.

And again, intuitively, that seems right. The truth of an atomic thought is not secured by a conjunctive fact. Similarly, $\langle P \vee Q \rangle$ is not true in virtue of both $f[P]$ and $f[Q]$ working together.

Universals

Universal truths, it might seem, present us with a special problem. In terms of a Fitch-style natural deduction system, the introduction rule for universals is:

$$\text{U-Intro} \quad \left| \begin{array}{l} \left| F\alpha \right. \\ \left| G\alpha \right. \\ \text{Every } F \text{ is } G \end{array} \right.$$

In this rule, α is an *eigen* variable or arbitrary object term. U-Intro involves a sub-proof in which we suppose that an arbitrary object α is F and derive a conclusion that it is G —we examined the speech-act structure of these proofs in §54. How are we to understand truth-making of universals in terms of this rule? The answer is that it is not the proof that is the truth-maker. It is the premises for the proof that correspond to the truth-makers. What U depends in asserting a truth-making claim is a derivation from premises, that is, assertions, using introduction rules, to the truth of a claim. U-intro, unlike, say, disjunction introduction, does not specify premises, in the sense of propositions, but a proof—derivation—of some kind. Obviously, what we want are the premises that could support application of the U-Intro, and not the derivation itself. The question now is what these premises are.

There are two cases that we have to consider in answering this question. There is the case of non-accidental universals, true by virtue of necessitation of some kind, and accidental generalities, true by virtue of brute facts.

Case 1: Where the universal *Every F is G* is a non-accidental generality, then there is a connection between F and G . That can either be analytic, nomological, or some fancy metaphysical necessity. In which case, if we accept that there are such facts of necessity, then the following is a truth-maker:

$$f[F \text{ necessitates } G] \text{ makes true } \langle \text{Every } F \text{ is } G \rangle$$

Now is this OK, since the assertion of F *necessitates* G can function as a premise for the proof. In this case, we might think, rightly, that the modal connections between F and G do

make $\langle \text{Every } F \text{ is } G \rangle$ true. This, in the context of GE requires some analysis of statements of the form: $A(\text{Being } F \text{ necessitates being } G)$. But that is not a problem as such.

Case 2: Where *Every F is G* is an accidental truth. One might wonder in this case, how U-Intro will be applied? What is the minimal information about contingent facts required to apply U-Intro? We answer this by considering what premises we need in order to carry out a suppositional proof that begins with suppose $F\alpha$ and ends with $G\alpha$. The supposition is that some arbitrary α has F . The answer is simple. The premises we require are the following:

$\{T_1, T_2, T_3, \dots, T_n\}$ are all the Fs
 T_1 is G, T_2 is G, T_3 is G, \dots, T_n is G.

If that is correct, then the following is our basic truth-making claim about universals:

$f[T_1 \text{ is } F], f[T_1 \text{ is } G], f[T_2 \text{ is } F], f[T_2 \text{ is } G], \dots, f[\text{Every } F \text{ is in the class } \{T_1, T_2, T_3, \dots, T_n\}]$ makes true $\langle \text{Every } F \text{ is } G \rangle$

This result entails that universals need facts of totality as part of their truth-makers. This is more or less what Armstrong (2005) argues, and indeed, it seems intuitively correct. We have derived a conservative result, but at least we get a principled reason for explaining an intuition shared by many people.

§63 Formal and Non-Formal Truth-Making

I think we can see how GE's theory of truth-making statements captures intuitive instances of formal truth-making. How does this analysis work in cases where $f[S_1], f[S_2] \dots$ make true $\langle R \rangle$ but $\langle P \rangle$ is not logically connected to $\langle S_1 \rangle$ and $\langle S_2 \rangle$, but has some other connection? Two instances of truth-making falling under this category are the cases below:

$f[\text{Fred is an unmarried man}]$ makes true $\langle \text{Fred is a bachelor} \rangle$
 $f[X \text{ is scarlet}]$ makes true $\langle X \text{ is red} \rangle$

Both these cases can be explained on the theory. We have the following introduction rules:

Fred is an unmarried man \vdash Fred is a bachelor

X is scarlet \vdash X is red

The first is unproblematic, since being an unmarried man is the canonical ground for being a bachelor. The second is unproblematic, if we allow Π -[*X is scarlet*] \Rightarrow Π -[*X is red*] to be a basing law.

We may worry that the vagaries of definition could get in the way here. Suppose your concept of *brother* is derived from *male sibling*. My concept of *sibling* is disjunctive: *either brother or sister*. In which case you will accept as an introduction rule: *Fred is a male sibling* \vdash *Fred is a brother*. I will accept as an introduction rule: *Fred is a brother* \vdash *Fred is a sibling*. You will accept the first truth-making claim, I will accept the second:

The fact that Fred is a male sibling makes true \langle Fred is a brother \rangle

The fact that Fred is a brother makes true \langle Fred is a sibling \rangle

Which is right? The answer is that there is no objective fact about which is right. There does not have to be.

Non-Formal Truth-Making: Synthetic

A posteriori theoretical identities, in contrast, represent a real challenge. Consider (i) below, which Armstrong (2004) takes to refute a reduction of truth-making to entailment:

(i) f [*X is H₂O*] makes true \langle *X is water* \rangle .

(i) looks right, but how can we explain our affirmation of it? As I shall argue, (i) is not a pure instance of truth-making in the sense that the *making* is purely analytic. The making in (i) is partially synthetic. That is because physical necessitation is being appeal to.

There is no analytic reduction of the concept of water to that of H₂O. Still, we can think of the concept of water as being captured in the phrase: *the underlying stuff that is causing watery appearances*. If so we have the following introduction rule:

X is the underlying stuff causing watery appearances. $\vdash X$ is water.

Given acceptance of that rule, we are obliged to assert (ii):

(ii) $f[X$ is the underlying stuff causing watery appearances] *makes true*
 $\langle X$ is water \rangle .

We now have to find the connection between X 's being H_2O and X 's being the underlying stuff causing watery appearances here. There is no analytic connection. There is, however, a non-analytic connection, given the physical necessitation:

For all x and y , if x is H_2O at y , then necessarily x causes watery appearances at y .

If so we have the following explanatory deduction:

X is H_2O

X is around here.

For all x , if x is H_2O at y , then x causes watery appearances at y .

X causes water appearances around here.

Thus, X 's being H_2O physically necessitates that X causes watery appearances. In which case, we are committed to the non-analytic, that is synthetic, claim of making (iii):

(iii) $f[X$ is H_2O] makes it the case that X is the underlying stuff causing watery appearances here.

We can now use (ii) and (iii) above to derive (i), by transitivity of making. This is a cross-modal or cross-categorical making. (i) is not a pure instance of analytic making; it has a synthetic element. This is not an objection, but merely a refinement in the notion of *making*.

§64 *Facts and Being without Metaphysics*

That completes the treatment of facts and truth-making.¹⁰⁵ The task has been to explain our intuitive acceptance of the core principles of our concept of truth and fact encapsulated in the schemata: **TM**, **TM1-2**, and **F**. What we have done is explain the intuitive appeal of instances of these schemata, rather than justify our assertion of them. We explain why people are inclined to assert them. Ought we to assert them? To assert an *ought*-sentence is to express approval of agents undertaking actions and disapproval of their not doing so—§39. Assertion of the sentence *O ought to assert instances of TM* involves defending the approval of defending the Π -property of **TM**-instances and disapproval of not doing so. We have these motivations since the Π -properties of **TM**-sentences will be parts of any normally functioning system. I submit that we have these motivations. So we assert these *ought*-sentences. A speaker ought to assert **TM**-sentences. They are justified.

It may be wondered if this somewhat trivial condition about introduction-rules is the basis for the great principle, that truth is fixed by being. We have explained the letter of the idea, but, it appears, we have diminished its spirit. We have diminished its spirit by deflating explanatory connection or grounding. If this means metaphysical impulses are thwarted, that is no problem, since they are based on semantic misconceptions. There is no metaphysics of facts, truth, or truth-making.

The result is a certain lightness about facts and being. For any domain *D* of sentences we are serious about asserting, we should not hesitate to talk of facts for that domain. Where does this leave us? GE is a theory which allows that for every truth, there is a fact, or a number of facts that make it true. There are physical facts and facts of value. There are

¹⁰⁵ There is the question of the so called *entailment principle*—Armstrong (2004)—which relates truth-making to entailment:

EP: Where $\langle P \rangle$ entails $\langle Q \rangle$, then if T makes true $\langle P \rangle$, then T makes true $\langle Q \rangle$.

The principle needs revision but Armstrong does not know how. GE's account of truth-making talk will underpin assertion of the following restricted principle:

EP*: Where $\langle P \rangle$ entails $\langle Q \rangle$, and *Q* is derivable from *P* by introduction rules only, then if T makes true $\langle P \rangle$, then T makes true $\langle Q \rangle$.

EP* is more or less trivial given given GE's analysis. Another matter is whether objects make true truths. One might think that Socrates makes true \langle Socrates exists \rangle . But there is apparently no room for that given GE's analysis. Rather it is the fact of Socrates's existence that makes true \langle Socrates exists \rangle . Given the non-objectual treatment of facts, we can treat this as a spoils for the victor case.

mental facts and, as we shall see, facts of meaning. What makes one's statement, say, about identity true? The answer is a fact of identity. What makes it true that murder is wrong? The answer is the fact that murder is wrong. Moral facts are no more light weight than physical facts. We cannot say that facts about colour have some kind of metaphysical weight not possessed by value facts. Negative facts are no more or less real than positive facts. Does the admission of negativity mean commitment to dubious entities? Is it an assault on common sense, and good metaphysical taste, or Occam's razor? The answer is not at all. Facts are not entities. Occam's razor does not apply to them since they are not entities. They lack any metaphysical nature since they lack any nature at all. Common sense is preserved: surely it is a fact that there are no dinosaurs left on Earth. Pick your negative facts!

This lightness about facts does not imply that facts are linguistic, or that our cognitive-linguistic activities somehow bring facts into existence, or that facts are pale shadows of sentences or thoughts. The heroic enterprise embarked on in works of metaphysics, such as Armstrong's (2004) recent work, is, from the lights of GE, deeply confused. But that is the semantic assumption of metaphysics: the implicit belief in a representational theory of meaning. With the dissolving of Semantics, much of what we call metaphysics dissolves as well. Being is not an object, with a general nature, that must be uncovered. It lacks a metaphysical nature. Having described the structure of the home language, and done natural science, we have no further explanatory task to carry out in relation to questions of the form of reality. (I say more about this outrageous claim in chapter 11.)

Being without Nature

One way of understanding the claim that being is not an object is just to consider the nature of the referring term *being*. This term is a nominalization of the verb *be*. The use of the verb is the predicative use, as in $O \text{ is } F$, for whatever O and F . The mental element underpinning *being*, which functions as a Φ -property, is that cognitive element that underpins all predication: concatenation $+$. In brief: $\Phi\text{-}[being] = +$. When we think of the problem of being, we think about that general worldly phenomenon that is correlated with all predication. That worldly correlate is referred to by the term *being*, that is, a term whose Φ -property is the

general concatenative function $+$. This does not mean that the term *being* refers to this concatenative function. That is absurd and does not follow from the present claim. The present claim is simply that the mental element of the term *being* is $+$. As for the question, what does *being* refer to? That can only be answered in a banal way: *being* refers to being. Being is itself without nature.

Terms like *existence* are to be understood similarly. This term is a nominalization of *exist*. The verb *exist* is that verb we use to defend dispositions to use referring terms in simple assertions—see §53. So the Φ -property underpinning use of *existence* is something like the mental relation: *disposition to use ... in simple assertions*. This kind of mental state is related closely to predication and the concatenation function. Hence, *existence* and *being* seem almost to be synonyms, and to refer to the same kind of thing.

The search for theories of the nature of existence and being, and indeed reality, that is, the kinds of theories sought by metaphysicians, is a vain pursuit. Just as there are no metaphysical theories of properties, propositions, and facts, there are no metaphysical theories of being and existence. These comments will be amplified in chapter 11, when issues of realism and metaphysics are examined.

Chapter 9

Objective and Non-Objective Truth

§65 Objective Truth

GE's ability to account for our talk of facts and our acceptance of the unrestricted validity of truth-making schemata—see §57—may not dispel an impression that the framework of GE leaves unexplained an essential feature of our understanding of language in relation to reality. This is the possibility of objective truth. The impression is this: GE is merely expressive, and so cannot capture the requisite sense of the mind-independent fixing of truth. Perhaps GE can give an account of why we are inclined to assert instances of *<S> is true in virtue of the fact that S*, and so on. But this leaves us up in the air about how things have to be for there to be an objective fact of the matter about whether *S* or not. That is, it leaves up in the air objectivity of fact and objectivity of truth. Objective truth requires the existence of a determinate, structured, mind-independent reality that our thoughts can represent or fail to represent. Representation of determinate reality has no role to play within the theory of language agency offered in GE. So how is GE consistent with the possibility of objective truth?

This attack on GE makes an assumption: that we need a theory of objectivity and objective truth. That without some theory of objectivity, our talk of the objective will be shown to be empty, baseless, without justification. In other words, the objection assumes that objectivity has a nature, and it is because we can realise this nature potentially in our interaction with the world that objective truth is possible for our thoughts. The assumption is false. Objectivity has no nature that can be uncovered by a theory, and certainly lacks a nature grounded in representation. Objectivity does not need to have a nature for our talk of objectivity to be perfectly fine, and indeed for us to obtain objective truth. The claim of *no nature* is simply a fancy way of saying that having provided a language-agency analysis of

our talk of objectivity, there is nothing left to say about what objectivity is. We exhaust it theoretically.

In approaching the question of the language agency behind the predicate *objective*, we shall first look at domains of talk in which, apparently, objective truth is not a possibility. The domain in which truth seems to go non-objective is the domain in which we discover the phenomenon of *faultless disagreement*. Two parties can disagree about a proposition that S —one party asserts S , the other denies S —but neither is in error. Below I examine and discard some well-known approaches for dealing with this phenomenon, including relativism about truth. I then provide GE's treatment. GE's analysis is based in a refinement of the theory of assertion. Assertion can come in two forms, which I label *objective* and *non-objective*. Equipped with that theory of the structure of assertion, we can then provide a theory of the language agency of *objective*.

§66 *Faultless Disagreement*

Say that Smith asserts that P and Jones asserts that $\neg P$. Smith and Jones disagree. But does it follow that one of them must have made a mistake? It certainly does if the following principles are correct—here $\langle P \rangle$ denotes the proposition that P and T is the truth-predicate:

- (a) The T-schema is valid: $P \leftrightarrow T\langle P \rangle$;
- (b) Excluded Middle is valid: $P \vee \neg P$;
- (c) Truth is a norm of assertion: it is incorrect to assert P where it is not true that P ;
- (d) Claims with the content that P are truth-apt.

Here is the proof. Given excluded middle, $P \vee \neg P$, suppose P . Then by the T-schema and logic, $\neg T\langle \neg P \rangle$. Jones then asserts something not true, and so has made an error. Suppose $\neg P$. Then by the T-schema and logic, $\neg T\langle P \rangle$, and so Smith has asserted something not true and so has made an error. Conclusion: either Smith or Jones has made an error. The error may not involve any incorrect reasoning, it is just the wrongness of falsity. Even if excusable as an instance of bad epistemic luck, it's still error. It follows from (a)-(d) that there can be no

faultless disagreement—cases in which speakers’ assertoric commitments are contradictory but neither assertor is wrong. Nevertheless, some areas of discourse, it seems, allow faultless disagreement. Statements of taste are a well-known case.¹⁰⁶ Smith asserts **H** below, savouring a steamy haggis, and Jones denies **H**, shrinking from the plate of offal before her:

H : *Haggis is tasty.*

Assuming that both make their assertions sincerely, and clear-headedly, in the light of their food preferences is either Smith or Jones in error here? We are inclined to say no: faultless disagreement is possible in cases of taste. Thus we have (*e*):

(*e*) There are regions of discourse that allow faultless disagreement: parties can disagree about P, but neither is in error.

The principles (*a*)-(*d*) and (*e*) are jointly inconsistent. So we have a kind of paradox. There are at least four well-known ways of escaping the conundrum. First, deny there can be faultlessness: in any disagreement, some party is in error—give up (*e*). Second, go non-cognitivist and somehow deny that there are real truth-apt assertions here—give up (*d*). Third, go content-relativist and deny that there really is disagreement here, since the sentence **H** is different in content when mouthed by the respective parties—this again involves giving up (*e*). Fourth, deny the objectivity of truth by embracing relativism about truth, claiming that truth is relative to speakers or standards, and that the norm of assertion is rephrased in terms of relative truth—give up (*a*) and (*c*) as they stand.¹⁰⁷

In what follows I argue against all of these positions, and urge adoption of a new approach, that naturally emerges from GE. In terms of GE, the solution to the paradox is to give up the idea that truth is essentially objective without embracing relativism about truth. This is achieved by the idea that assertion comes in two modes: *objective* and *non-objective*. Assertion is the act in which speakers defend Π -properties. But our relation to Π -properties has this complication: some Π -properties are such that if we instantiate them we feel

¹⁰⁶ I take this setting up of the problem from Kölbel (2002).

¹⁰⁷ We might also ditch the T-schema and non-contradiction. I won’t consider these here.

compelled to defend them. Defence is non-optional or binding. In contrast, some Π -properties are such that if we instantiate them our defence of them is optional. We may, under certain circumstances, renounce their defence, even though we keep the Π -properties. Assertions involving defence of the first kind of Π -properties are objective, whereas assertions involving defence of the second kind of properties are non-objective. Both objective and non-objective assertions are assertions, it is just that assertions in the non-objective mode have a certain modal feature: speakers producing those assertions can renounce the performance of such assertions under certain circumstances of disagreement. They can do so without acquiring or losing grounds for the Π -properties of these assertions. They just renounce their defence of the Π -properties.

Faultless disagreement is analysed in these terms. (Details are given below.) Assertions of taste are instances of non-objective assertion. From the non-objective mode, Smith and Jones are committed to their respective assertions, which contradict each other. But from the objective mode, they are not committed to these assertions, and indeed, are not disposed to make any evaluation at all about H , including affirming the excluded middle or bivalence of H . From the point of view of the objective mode, there is no fault, since bivalence cannot be asserted. The solution to the paradox of faultless disagreement, as set out above, is that we have a fluctuating attitude to bivalence depending on the mode of assertion in which we are approaching a sentence. So *(d)* is no longer affirmed unequivocally. This approach invokes a kind of relativity, but not a relativity of truth. The relativity of mode of assertion does not imply relativity of truth, as we shall see.

I now I lay out my main objections to the four standard responses to the paradox of faultless disagreement for taste-statements. I then offer GE's alternative. Having put the structure of the approach in place we can then provide a language-agency analysis of talk about objectivity.

§67 Realism, Non-Cognitivism, Content Relativity, and Relativity of Truth

There are currently four well-known approaches to solving the paradox of faultless disagreement. These are: realism, non-cognitivism, content relativity and truth-relativity. I will not argue in detail against these views, but give some general indications of why they are far from promising treatments of the problem.

Realism

Call *realism* the view that all our cases of alleged faultless disagreement are cases in which there is error. Either Smith or Jones has made a mistake. A view about taste predicates that would accord with this view is a response-dependent semantics: tastiness is that property which normal individuals under normal circumstances find to satisfy their gustatory preferences. The problem with realism is that it leaves most people in error, and few people with the wherewithal to make judgements of taste. As a speaker contemplating *H*, Smith would have to ask *am I normal?* Who knows? Maybe Jones is, maybe not. Speakers are frequently very confident in their taste judgements, but if they are really making assertions about the reactions of normal speakers, then clearly they must be blind to the semantics of taste-statements. Such wide-spread semantic blindness is untenable.

Non-Cognitivist Expressivism

The non-cognitivist expressivist solution looks to the assumption we have made so far that *H* is truth-apt. We might deny this, and claim furthermore, that no assertion is made with *H*. There is no norm of assertion being followed since Smith and Jones do not make bona fide assertions, and thus no conflict with the T-schema. On what grounds would we deny that *H* is truth-apt? Well, we take up a non-cognitivist theory in relation to statements of taste. According to non-cognitivists, statements like *H* are mere expressions of gustatory preference; they are neither genuine assertions nor manifestations of belief.

We briefly discussed classical expressivism in §7 and found it to be a dubious doctrine. It is unfruitfully radical—it denies the obvious fact that statements of taste are assertions, truth-apt, and embed in logical compounds and manifest beliefs. At the same time, it is fruitlessly conservative in its theoretical outlook—it accepts the conservative view of

orthodox semantics in relation to truth-aptness, propositional content, and belief. I therefore do not think this position is worth considering further in relation to faultless disagreement.¹⁰⁸

Content Relativism

The content relativity line attempts to preserve objectivity by holding that the content of *H* in Smith's mouth is different to its content in Jones's. But the problem with this approach is that it takes away the disagreement and treats interlocutors as speaking at cross-purposes. Take the dialogue:

D: Smith: Haggis is tasty.

Jones: That's false. It is not tasty.

If the content relativist is right, **D** involves Smith and Jones making reports about their respective states. **D** comes out as:

Smith: Given my preferences, I am committed to liking haggis.

Jones: False. Given my preferences, I am committed to disliking haggis.

One can have a variation of this in terms of taste-standards if one wants, as in:

Smith: Given my taste standards, haggis is tasty.

Jones: False. Given my taste standards, haggis is not tasty.

Either way, the result is a complete distortion of the dialogue **D**.

A content relativist might attempt to repair matters by proposing that the preferences or taste standards in operation in **D** do not have to be those peculiar to Smith or Jones, but, say, shared preferences or standards. But it is not clear what shared preferences or standards would be. Furthermore, introduction of a common set of preferences or standards would remove any faultless disagreement, since, presumably, there would be a fact of the matter about what holds taste-wise relative to the standards and preferences. Moreover, that such conversations are debates about conversationally shared preferences or standards does not

¹⁰⁸ See Kölbel (2002) for further elaboration of this non-cognitivist orientation.

accord with the fact that observers, not included in the conversation, will not judge Smith or Jones right by consulting the shared preferences or standards of Smith and Jones's conversation. The observer will judge by their own preferences and standards.

Relativism about Truth

The fourth and most interesting solution is relativism about truth—recently defended by Kölbel (2002) and MacFarlane (2005)—touted by many to be the solution closest to common sense. On this view, the faultless disagreement of Smith and Jones in relation to **H** results from the fact that truth in this context is relative to taste standards or perspectives of taste. **H** is true for Smith and false for Jones. How do we understand relative truth? According to MacFarlane (2005) we should understand propositions not simply as functions from worlds to truth-values—which is the standard idea—but as functions from worlds and *standards* onto truth-values. Thus the factors that enter into the evaluation of a proposition of taste are not simply a world, but that world and standards of taste. The same world can be associated with more than one standard of taste. Thus the actual world as a context of assessment contains both Smith's and Jones's standards. The proposition that Haggis is tasty is true relative to one, but false relative to the other.

In terms of the five mutually inconsistent principles (a) to (e)—§66 above—we are dropping (a) and (c). The T-schema takes on a relativised form—*<P> is true relative to x iff P relative to x*, where *x* is the standard or taste perspective. Also the principle (c), the thesis that truth is a norm of assertion, needs to be modified by reference to relative truth. What a speaker commits herself to in asserting *S* is truth relative to the world in which she makes the assertion and her standards. There is then no reason to think that Smith or Jones has made an error, since each asserts relative to his/her own perspective.

Does relative truth solve the problems of faultless disagreement? I don't think so. There are several related problems:

1. Relativism as articulated seems to remove the disagreement from faultless disagreement just as content relativism does, but in a more subtle way. We can see on MacFarlane's account that Smith and Jones both use the sentence **H** to encode the same proposition: a function from worlds and standards of taste onto truth-values. Smith asserts it,

Jones asserts its negation. That might look like disagreement. But it is not. First, consider the case in which we think propositions are simply functions from worlds onto truth-values. Then two speakers, asserting P and asserting $\neg P$, only disagree if they are in the same world, qua context of assessment. The relativist—as MacFarlane views them—sees propositions as functions from worlds and standards onto truth-values. Where Smith asserts H , and Jones the negation of H , the only way there is genuine disagreement is if they are making the assertions relative to the same world and standards, but they are not. In other words, all Smith's assertoric commitments are relative to his standards of taste. All Jones's assertoric commitments are relative standards of taste.

This does not strike me as being a mere artifact of MacFarlane's particular semantic orientation to relativism. Any relativist approach will encounter the same problem. They must relativise the T-schema and the norms of assertion to speakers' perspectives or standards of taste. But that is enough to seal each speaker in her own assertoric universe, closing off the possibility of conflict. How then could we explain Smith and Jones reaching a point in debate when they agree to disagree? There is not so much a moment of agreeing to disagree at some point in the debate as a point in which they moving from an assumption of absolute truth—they were assuming that a non-relative truth-predicate was in operation—to one in which it is recognized that truth is relative, and so there is no disagreement.

2. We might want to treat taste-discourse, amongst others, in relativistic terms. But there are other areas of discourse that we want to treat absolutely: most statements about the physical world, some areas of morals, mathematics, and so on. The bearers of truth in the latter case, it would seem, should be treated as functions from worlds, conditions, onto truth-values, and not as functions from worlds and perspectives onto truth. If so, a disunity infects the treatment of truth-bearers—sometimes they are one kind of function, sometimes another—and the truth-predicate—sometimes one-place, sometimes two-place. Or we globalize the relativity: all statements are relative to perspectives, its just that for certain classes of statements, there is agreement across perspectives—and must be such agreement. But here *perspective* must be a multifarious matter, include elements that allow gustatory,

moral, aesthetic, and epistemic orientations. For our absolute statements, perspective does no work. And why should relativity in certain isolated areas produce relativity across the board?

3. What are standards of taste or perspectives for the relativist anyway? Let us say that taste standards are just preferences. How are statements true relative to preferences? One answer is by virtue of propositions being functions from worlds and preferences onto truth-values. But this talk of functions is surely only our means of modelling a semantic reality that cannot be identified with such functions, which after all, are just sets. So truth-bearers are sets of ordered n-tuples of worlds, preferences, and truth-values, or—cutting out the truth-values—ordered n-tuples of worlds and preferences. But again, this looks like formal construction providing a tool for modelling, rather than the reality itself.

§68 Two Modes of Assertion and Belief: Objective and Non-Objective

We have examined the four main approaches to phenomena of faultless disagreement. All seem to have serious difficulties. I don't suppose I have shown that they cannot be made to work. What I show now is that from GE there naturally emerges a more promising option with very little conceptual stress. This is provided by GE and the two-mode view of assertion outlined in the introduction.

Assertion is an act requiring two features of cognitive life. It requires a substructure, which I have called the NC system—a network of Π -properties and other states in C-law relations—and a superstructure whose assertive component involves acts of signalling dispositions to signal C-law relations amongst Π -properties. Belief is not at the level of the NC system, it is not a Π -property, or a grounded Π -property. Both belief and assertion are tied to inner or outer voicings of the architecture of Π -properties in C-law relations. Deception occurs when speakers have the superstructure without relevant parts of the substructure. Call the opposite situation *renunciation*: speakers have the substructure but lack the superstructure. They have Π -properties and C-law grounding relations, but refrain from voicing, even to themselves, the architecture of the substructure.

Here is an example: I like Vegemite. It has a very salty, yeasty flavour and is a pitch black, spread for bread. I have the gustatory preference state for Vegemite, which is grounded in a preference for other salty products. I am disposed to assert *I like Vegemite*; that is, to make reports about this liking. In making such reports, I defend a Π -property based in perceptual states. These perceptual states involve capacities to introspectively self-monitor my own preference states. I have no general problems with asserting *I like Vegemite*, and defending the perceptual state about my taste preference. In contrast, I am not so readily disposed to assert *Vegemite is tasty*, thereby defending the taste preference state itself. In the face of the salt-phobe, and cognisant of the rich inner life of her taste preferences, I will refrain from publicly and vigorously asserting *Vegemite is tasty*. That is, I am tentative about publicly defending my preference state. But more than that, introspecting, and keeping the salt-phobe firmly in my mind, I may find compunctions in voicing the state even to myself. As much as I try to elicit an inner voicing, no voicing emerges. I feel intuitively, that, in the light of the salt-phobe, I cannot maintain the disposition to defend my liking of Vegemite. It fades away. It appears an empty exercise to voice, even internally to myself, my preference.

On the other hand, under different circumstances, enjoying Vegemite-smearred toast with a fellow salt-file, intermittent voicings of the form *This is tasty* take place. There are no compunctions about defending the preference state here.

Here then is the key idea. Our readiness to voice grounding relations, at least for some kinds of Π -properties, can fluctuate. Put the speaker in one context of conversation, and she lacks any disposition to voice grounds for Π . That is renunciation. Put her in another, and there is no renunciation, but voicing of grounds for Π . Has the speaker changed her mind in the meantime? Not in any conventional sense. Changing one's mind occurs when, within the NC-system, there is alteration in what Π -properties are instantiated. A speaker gets new evidence, or improves her processing of the Π -property states that underpin her evidence-based beliefs. There is change in the substructure—the NC system. But the fluctuation phenomenon we are describing is not this. All that is changing is the speaker's decision to voice grounding relations, that is, defend the Π -properties concerned.

Let us say that where speakers fluctuate in voicing their grounds that the Π -properties for those speakers are *optional*. The pre-conditions for optionality are given thus: let us call a Π -property type a set of Π -property homologues in members of a population that underpin assertions of a sentence S . For example, a Vegemite taste preference is a type realised as homologues in U1, U2, etc, and encoded by *Vegemite is tasty*. In which case, a Π -property type Ψ for a population is optional if and only if:

- (a) Some members in the population instantiate Ψ , whereas others do not. Furthermore, disagreements with regard to asserting S frequently end in deadlocks. Opposing parties cannot persuade each other to change their minds;
- (b) Speakers do not view disagreement in asserting S as a danger to the cognitive, physical, or psychological life of themselves or others;
- (c) Nevertheless, a limited practice of assertion of S or its negation is allowed. But there is a degree of intolerance in regard to such practices; audiences will accept dispute about assertion of S or its negation up to a certain point. If argument about positions is too protracted, those participating are potentially subject to disapproval.

Let us call *non-objective assertion* the kind of assertive practice allowed in (c). It will become apparent later why we call such assertion *non-objective*.

I propose that taste-assertions are paradigm instances of non-objective assertion. My Vegemite example satisfies the conditions (a) to (c) above. (a) There is diversity in taste preference and lack of effective procedures for persuading others to change their taste preferences where there difference in taste. (b) We do not see others with different taste preferences as disadvantaged or in danger, and so, do not disapprove of their activities or states. (c) We allow the practice of taste-statement making in which such preferences are defended or rejected. Yet, if there is dispute, we expect people, at some stage of debate, to renounce their dispositions to voice grounding architecture. If someone insists that Vegemite is tasty, refusing to back down, and signals sincere disapproval of dissenters, we are inclined to judge that person as irrational. That is, we express a specific kind of disapproval of them.

Why do we allow assertive practices in the non-objective mode, the *NO*-mode?

Presumably there is a range of social and pragmatic reasons for allowing *NO*-mode assertions. Discourse about taste in food is a useful way of providing information of a certain kind. Asserting the tastiness of Vegemite can induce someone to try it, and find it likeable in turn. Who knows? Others may share our preferences, and by introduction to such food items, enrich their lives, and so on.

Not all Π -properties have this feature of optionality. Let us call Π -properties *non-optional* or *binding* if, given that we possess them, we are disposed always to defend them. Here grounding relations are always voiced. The voicing may not be public. I may believe that $2 + 2 = 4$, and the regime may deny it and forbid public assertion. Yet I still believe it. I still internally defend the Π -property. I cannot conceive of not doing so. I may walk in a forest and encounter an animal that I take to be a rock wallaby. I am disposed to defend Π -[*There is a rock wallaby here*]. If someone disputes that there are such wallabies in the area, in the course of an exchange, I do not renounce voicing of my state. I may cease to produce public assertion—maintaining good social relations can require that. But, unless someone convinces me I am in error through changing my Π -properties structures, my private voicing remains. There is no optionality: no fluctuation phenomenon. Certain kinds of Π -properties are not optional. They are binding, which means, having grounding states for them, we are always disposed to voice those grounds, even if only privately. Call assertion based purely upon such Π -properties, assertion in the *objective-mode* or *O-mode*.

The Relativity of Assertion

One might be puzzled. Given that the Π -properties defended from within a *NO*-mode are non-binding, how can utterances of sentences like *Vegemite is tasty* really be assertions? Won't they be insincere? No. I really do possess the gustatory preference state Σ , when I assert *Vegemite is tasty*. To sincerely defend Σ is to possess Σ , and to have a ground-indicating disposition in relation to it. I have the preference state Σ , and that preference state really can be grounded by instantiating other Π -properties actually accepted. We can say the same of the person who asserts that Vegemite is not tasty. The difference with the *O*-mode assertions is that some of the grounds of *NO*-mode assertions are non-binding: in the face of

deep conflict, we are obligated to renounce them: we cease making assertions and indeed having the beliefs. The fact that the assertion of *Vegemite is tasty* is based on optional grounds is a modal fact about the assertion. Were conflict of a certain kind to arise, Σ would cease to be defended. That does not imply that Σ is actually not defended.

The two modes of assertion, *NO* and *NO*-modes, both involve the identical activity of defending Π -properties. This means that assertability or believability is a relative matter. With my salt-free friends, I believe that Vegemite is tasty, but in renouncing in a bitter debate over Vegemite, and conceding that it is only a matter of taste, I cannot at the moment be said to *believe* that Vegemite is tasty. In conceding that there is no real issue about the tastiness of Vegemite, one is conceding that one cannot clear-headily assert with conviction: *Vegemite is tasty*. There is a form of Moorean paradoxicality in saying *I cannot properly assert that Vegemite is tasty, but I believe it is*. One cannot properly believe that it is tasty. Under such circumstances of renunciation, we do say things like this: *It is only a matter of taste. Vegemite is tasty for me, but I cannot insist*. The *for-me* locution is just that device we use to convey our disposition to (non-objectively) assert that Vegemite is tasty, and also to convey that one does not believe it objectively.

§69 Faultless Disagreement and the Relativity of Assertion and Belief

How does this relate to faultless disagreement? We can now understand what's going on in the dispute between Smith and Jones. Smith and Jones launch into their assertions of taste. Smith asserts that haggis is tasty, but Jones denies it. They disagree. Smith and Jones disagree in attitude. Asserting **H**, or any statement of taste, involves speakers defending a gustatory preference state Ψ . Smith utters **H** defending Ψ . Jones utters the negation of **H**, which means that Jones defends rejection of Ψ . Jones's rejection of Ψ is constituted by the fact that (i) he is cognitively constrained from tokening Ψ , and (ii) the ground for that constraint state is itself a Π -property. (See the discussion of rejection in §23 and §37.) Given his preferences, and the urge to defend them, Smith must say that Jones is wrong. And given his preferences and the elan to defend them, Jones must say Smith is wrong.

This is the disagreement part. The faultless part is derives from the fact that Smith and Jones produce their assertions within the non-objective mode, or *NO*-mode. Taste-preferences and rejections of them are optional grounds. If called on to adjudicate between Smith and Jones we would simply say, if we were rational, that there is no objective fact of the matter. It is just a matter of taste. In making this judgement we take up the position of assertion within the *O*-mode, objective assertion. In that stance, we renounce our defence of such preference states. Perhaps we in fact have a preference for haggis. But we do not voice the architecture of the grounding we have underpinning that state. Indeed, if we are sincere in our stance, we will not even say privately to ourselves, (**H**) *Haggis is tasty*, and so Jones is wrong, Smith is right. We do not voice even to ourselves. From this stance, we cannot say that either party is right or either party is wrong. If so we cannot attribute fault. Neither has made an error. In short, Smith and Jones disagree (in the *NO*-mode) but there is no fault, since, from the *O*-mode, we renounce all assertions featuring use of **H**.

Indeed, the only fault we could attribute would be if either Smith and Jones or both persisted in their assertion in the face of dispute. In which case, the fault would not be a fault about fact, in the sense that something false was asserted, but a fault about the nature of the subject matter, viz, that anything objectively true or false could be asserted.

Bivalence

From this *O*-mode perspective, **H** and its negation are judged as lacking cognitive content. **H** is not associated with a Π -property that can enter into grounding relations for assertion. Therefore, **H** cannot be believed since belief is a disposition to defend a Π -property. Nor can **H** be denied. From within this mode, then, there is no question of asserting or denying **H**. There is also no question of defence of a compound Π -property, such as those defended in assertion of $(P \vee \neg P)$, $\neg(P \ \& \ \neg P)$, $(T\langle P \rangle \vee \neg T\langle P \rangle)$, where $P = \textit{Haggis is tasty}$. In the *O*-mode, we are not assigning a third truth-value to **H**. It is rather enforced abstinence from assigning any value at all to **H**. If so, from the perspective of the *O*-mode, as bivalence fails for **H**, there is no objective fact at stake; no objective fact attends the debate in relation to **H**. In which case, one cannot speak of either party being wrong.

Furthermore, the set of T-sentences differs in the *O*-mode and in the *NO*-mode. There is a larger class in the *NO*-mode. In the *O*-mode and the *NO*-mode, the set of sentences with cognitive content differs. So the set of assertable T-sentences differs. For example:

T : It is true that Haggis is tasty if and only if Haggis is tasty

is not assertable in the *O*-mode but assertable in the *NO*-mode. In reaching that stage in their debate where Smith and Jones acknowledge that there is no objective fact of the matter, they are retreating to the *O*-mode. In that position, no judgements about ***H***'s truth are permitted. It would be odd, in denying ***H*** has cognitive content, to assert ***T***. (Or at least, that is the claim.)

Old-Time Emotivism

There is, then, an echo of the old classical non-cognitivism about taste statements, according to which taste-statements fail to encode propositions, manifest real beliefs and so are without truth-value. But there is a profound difference between the present view and that the non-cognitivist view. The classical non-cognitivist view would hold that ***H*** just ought not and cannot be asserted, because its mental content element is just a preference and preferences are not beliefs or truth-apt, but that is not the position of the present analysis. Although, from the point of view of the *O*-mode, we must refrain from asserting or denying ***H***. Nevertheless, this is not so simply because ***H*** is used to express an attitude. It is because in the *O*-mode, the attitude Ψ cannot be defended or rejected. The other difference is that this judgement only holds from the point of view of the *O*-mode. Within the *NO*-mode some speaker may defend Ψ , and assert ***H***, and some other speaker may reject Ψ and deny ***H***.

Relativism?

Does this relativism about assertion imply relativism about truth? It does not. Relativism about truth requires that truth is a two-place predicate: *x is true relative to y*. But there is no two place-predicate required to make sense of modes of assertion and the structure of renunciation as we have described it. Of course, there are locutions that might, to the untrained eye, appear to bespeak relativism of truth. In asking Smith the question *Is ***H*** true?* couldn't Smith say, *For me ***H*** is true?* Indeed, Smith could so answer, but what is behind this locution: *For x, ***H*** is true?* Nothing more than that Smith is disposed to judge ***H*** true in the

non-objective mode. Smith and Jones have different gustatory tastes. That means relative to the *NO*-mode, the gustatory preferences possessed by Smith commit her to asserting *H*, and the gustatory preferences possessed by Jones commit Jones to asserting *H*'s negation. But nothing in that requires relativity of truth, just relativity of assertion. We find the same old one-place truth-predicate, it is just that the practice of assertion can have two forms: an objective form and a non-objective form.

The term *objective* is that predicate we use to defend a disposition to make an assertion in the *O*-mode. Thus assertion in the *O*-mode is directed towards objective truth, but assertion in the *NO*-mode is only directed towards truth. This might make it look as if truth has two forms: an objective and non-objective form. Suppose one thought truth was a property of either correspondence, coherence, or the minimal property that Horwichian minimalism assigns to *is true*. Then, if we assume that Smith is perfectly sincere in asserting *H* in non-objective mode, making no error, and perfectly sincere and correct in refraining from asserting *H* in objective-mode, then *is true* must mean something different in each case: it must pick out a different property. This shows that one cannot think of truth in these terms if one adopts the dual mode view of assertion. That's fine. In GE we think of the truth-predicate differently. The view of truth required is an expressive view. The function of the truth-predicate, in the case in which we attribute truth to a known sentence *S* that we are capable of L-processing, is to defend the Π -property of *S*—see §17.

None of this means that *NO*-mode assertion—for example, the practice of assertions of taste—is a practice involving error. Assertions in the *NO*-mode are not objectively true or false. In the *O*-mode, one can assert neither their truth nor their falsity. Error only arises if we mistake the nature of the practice, thinking of it in purely *O*-mode terms.

In sum, the present analysis explains faultless disagreement in the case of taste through a dual-mode theory of assertion. This produces a kind of relativity of assertability, but no relativity of truth. Unlike our other four options canvassed in §67, the present expressivist treatment seems to explain faultless disagreement in the case of taste. Moreover, it is an entirely natural development in the explanatory system of GE.

§70 Putnam's Conceptual Relativity and Epistemic Modals

Are statements of taste the only places in which we find non-objective assertion? Clearly a whole range of value-sentences are domains in which faultless disagreement can arise. Are there candidates for objective assertion outside of the evaluative? Places to look are domains for which theorists have defended relativist thesis about truth. Two cases suggest themselves. One is Putnam's (1987) conceptual relativity thesis. We shall see in a moment that it fits our model of NO-mode assertion, so contra Putnam, does not involve any relativism about truth. The other is sentences with *epistemic modals*. A range of philosophers have argued that they exhibit relativity of truth—Egan, Hawthorn, and Weatherson (2004) and Macfarlane (200+). I argue that epistemic modals are not cases of relative truth, but nor are they cases of the NO-mode of assertion.

Putnam's Conceptual Relativity

Putnam (1987) raises the issue of what he calls conceptual relativity. This is the idea that there are alternative conceptual schemes for describing some domain of reality, and there is no objective fact of the matter about which one is the correct one. For example, Putnam suggests that there are two basic, but distinct, conceptual schemes about the structure of spacetime. In one conceptual scheme, spacetime is made up of points as concrete entities, in the other, points are abstractions, and non-zero extensions are the only concrete entities. Putnam argues that there is no objective fact of the matter about which one is right. How would one decide? Putnam sees this as the basis for claims of relative truth. Thus, *Points are concrete entities* is true relative to one scheme, false relative to another.

In terms of GE we can make sense of Putnam's categorial scheme relativity, though not of his relativism about truth.¹⁰⁹ The sentence *Points are concrete objects* has as a constituent *concrete object*. The term *object* is a formal predicate in the sense of §27. In asserting *T is a concrete object*, U defends a disposition to deploy *concrete object* in a

¹⁰⁹ One objection to this proposal focuses on the problem of the common truth-bearer: does the sentence *Points are concrete objects* mean the same thing in the different schemes? One might doubt that it does since the predicate *concrete object* has distinct extensions within the distinct schemes. So how do we have the same truth-bearer? (See Cox 1997). This is not a problem for GE since extensions do not fix meaning.

referring term whose Φ -state is an ι -property—see §45. On the other hand in asserting *T is an abstract entity*, I propose, U is defending a metalinguistic disposition to deploy *T* in a referring term that is based in nominalization—§50. In GE, talk of abstract objects is largely talk of certain non-objects, reference to which is introduced through nominalizations.

It could be that our language agency with the term *point* leaves it undetermined whether we treat *point* as an object term or non-object term. If that is right, then we can have two possible practices of non-objective assertion. That one in which we assert, non-objectively, *Points are concrete*, and that one in which we assert, *Points are abstract*. In a debate between two parties we can imagine renunciation. If disputing parties are really aware of how the discourse in these regions of talk functions, they will not pursue the matter as a question of objective truth.

But does this support relativism about truth? That is, do we say that *Points are concrete* is true relative to one scheme and not true relative to another? I think not. This is because Putnam's idea of schemes is a bogus one. We find two speakers U_1 and U_2 . U_1 is inclined, for whatever theoretical or pragmatic reason, to non-objectively assert *Points are concrete*. For U_1 , points are concrete. All that means is that U_1 is non-objectively disposed to assert *Points are concrete*. Similar comments apply to U_2 . Nothing here means that truth is relative. There is no two-place predicate of truth.

What leads Putnam erroneously to the relativism is the nonsensical image of carving up reality in two alternative ways: through one conceptual scheme—the concrete point scheme—or the other—the abstract point scheme. Predicates do not function in any way that makes such an idea intelligible.¹¹⁰

Epistemic Modals

Epistemic modals are modal auxiliaries like *must* or *may*, used to convey information about epistemic states. Ignorant of the exact whereabouts of Bob, but believing he has a tendency to go to Berlin, U asserts: *Bob may be in Berlin*. In this case, U is conveying, in some sense, the possibility, relative to what she believes, that Bob is in Berlin. Epistemic

¹¹⁰ Here GE is in perfect agreement with Davidson's (1974) classic attack on conceptual schemes.

possibility is not mere consistency with what is known or firmly believed. We often express epistemic possibility about P and $not-P$, where we know that one of the set $\{P, not-P\}$ must logically follow from our beliefs. We just cannot work out which because of our computational limits. Another feature of epistemic modals is that although they seem to convey something about the speaker's epistemic state, they are not assertions about that state or about states of related individuals, *pace* De Rose (1991). *May-Q* uttered by U is not equivalent to *P is epistemically possible for U*. Consider the dialogue:

Smith: Bob may be in Berlin.

Jones: Bob cannot be in Berlin. He said he had a conference in Spain.

Jones judges Smith's assertion false—*cannot* is the standard negation of *may*. But Jones is not denying that Bob's being in Berlin is epistemically possible for Smith. Jones can grant that it is epistemically possible for Smith that Bob is in Berlin. Nevertheless, Jones judges what Smith says is false. The basis for his falsity claim is what he thinks about the question of Bob's whereabouts. One might attempt to preserve the idea that epistemic modals are used to make claims about epistemic states by proposing that in the dialogue above, Smith is making a claim about the shared epistemic state of Smith and Jones, and, say, others included in the conversation. But this will not work either. A third party, U, totally outside the conversation will not judge what is said by Smith or Jones true or false, on the basis of what she believes about what is epistemically possible for Smith and Jones as a collective. Rather, U will make that judgement on what she believes about the whereabouts of Bob.

Epistemic modals have then a kind of evaluative relativity. A speaker H judges *May-Q* as uttered by U true based on what H believes about Q , not on what H believes about what U's state is. This kind of relativity has led some theorists to propose that *May-Q* is only open to a kind of relative truth—see Egan, Hawthorne, and Weatherson (2004) and Macfarlane (200+). That is, *May-Q* is true relative to H if and only if Q is epistemically possible for H. But a relativistic thesis is not attractive. It fails for two related reasons:

Firstly, it takes away the incompatibility that is present between Smith's assertion and Jones's. Imagine that you are Jones, confronting Smith asserting *Bob may be in Berlin*. You

assert *He cannot be in Berlin* as a refutation of what Smith asserts. But if the sentences only have relative truth-conditions, your sentence cannot do that. Each is true within its own universe. Of course, you have denied the proposition that Smith asserts, but since only relative truth is available, your stance is not incompatible with Smith's. You are not refuting Smith, although it seems you are. Secondly, the relativist view cannot allow for the fact that there is a sense in which Jones's view is the right one in the light of the facts: Bob is in Spain. So Jones's assertion is right, whereas Smith's is wrong. In making our own judgement as onlookers, based on the knowledge that Bob is in Spain, we side with Jones. But all that means is that relative to our state *Bob may be in Berlin* is false. There is no question of any party having the right view on the sentence. But that seems wrong. There is a sense in which Jones's assertion is a better reflection of the reality. It is an improvement on Smith's.

I think this indicates that the relativist view cannot be right. If so, both the non-relativist truth-conditional analysis and the relative-truth-conditional analyses fail to account for epistemic modals. So what kind of analysis is there? The answer is: *not one in terms of truth-conditions*. What works is a routine application of GE. The evaluation relativity of *May-Q* noted above is simply the evaluative relativity found in all assertions: the intersubjective evaluative dimension recorded in **ID-True**—§17. Thus, in asserting *May-Q*, U defends the epistemic possibility for herself of *Q*. H judged U's assertion true just in case H is disposed to sincerely defend the epistemic possibility in her own case.

What is this epistemic possibility? This really boils down to the question of what the Π -property of *May-Q* is. I suggest there are two aspects summed up below:

Epistemic May	Π -Property Specification
<i>May-Q</i>	(a) \neg D-token: Rejection of Π -[<i>Q</i>], where no state Ψ generates this state. (b) Some disposition to token Π -[<i>Q</i>].

The first aspect (a) involves lack of a disposition to reject Π -[*not-Q*]. But note: this state is not generated by any Π -property Ψ . U does not instantiate a C-law of the form:

L: $\Psi \Rightarrow \neg$ D-token: Rejection of Π -[*Q*].

If U did, her state would amount to rejection of rejection of Π -[Q], since rejection is simply that state which is not being disposed to token some Π -property Σ grounded in some Π -property Ψ —see §23. L is the kind of grounding that underpins assertion of double negation *not-not- Q* . It must be then that for *May- Q* , U's lack of a disposition to reject Π -[Q] is not grounded in any evidence Ψ . So what then grounds U's defence of (*a*)? Nothing beyond the state itself does that. U just finds through introspection that she is not disposed to reject Π -[Q]. Self-monitoring is not evidence—see §21. U's defence of the state (*a*) is just (*a*). It is circular. The assertion is dogmatic but of a virtuous kind—see §21—that same kind of grounding as we find in assertions about basic logical laws such as excluded middle. The second aspect (*b*) is that U has some disposition to token Π -[Q]. It may be minimal, but U assigns a non-zero probability to Q . On the other hand, the element (*b*) can be grounded in evidence. Thus U can reason: *It is Sunday. So Bob may be in Berlin*. In this case, U's positive degree of belief that Bob is in Berlin is grounded in her belief that it is Sunday (along with other background beliefs).

There are puzzling features of epistemic *may* that we still need to comment on. The first is a kind of faultless disagreement. In the dialogue above, Jones is right, and not Smith. But Smith's assertion was not made upon any incorrect evidence or faulty reasoning. Smith does not believe anything false apart from *Bob may be in Berlin*. That might seem strange. But it isn't if we remember that epistemic modals are asserted in a dogmatic way. They are not asserted *on evidence*. (Strictly speaking it is the (*a*)-component of their Π -properties that is not based in any Π -property, as noted above.) This puzzle is related to another strange feature of epistemic modals. Typically we think assertion goes by high subjective probability. This means that if we think there is a 50% probability that what we are saying is false, we cannot assert it. But this seems to be contravened by epistemic modals. For me it is 50% likely that the coin landed heads. So, it is 50% likely that *The coin may have landed heads* is false. Yet, I assert, *The coin may have landed heads*, despite the fact that it only has 50% probability of being true.

The way out of the paradox is to realise that epistemic modals cannot be assigned subjective probabilities. Only assertions fully based in evidence can be. There is no

subjective probability of *May-Q*. But if not all sentences can be assigned subjective probabilities, not all assertion goes by high subjective probability—pace Jackson (1987). In which case, there is no paradox.

§71 Objectivity and Representation

In explaining faultless disagreement I have suggested that the predicate *is objective* is employed to express attitudes of bindingness in relation to Π -properties. In asserting '*S*' is *objectively true*, U defends a disposition to assert *S* in the objective mode. One can assert *It is true that haggis is tasty*—one may do so when one speaks of Haggis's tastiness in the non-objective mode. But one ought not to assert *It is objectively true that it is tasty*. A paradigm form in which we convey objectivity for a domain is through assertions like: *There is an objective fact of the matter about whether S or not*. In making such a claim, U defends a disposition to assert, in the *O*-mode, *either S or Not-S*. These comments capture the basic language agency underpinning the predicate *objective*. In that analysis, there is no attempt to reduce facts of objectivity to discourse features, that is, reduce objectivity to facts about discourse. Such an impression would be to treat language-agency analysis as analytic reduction. It isn't so. Still, there are some who might think we need a theory of objectivity. They will ask: what constitutes objectivity? From the point of view of GE, that question taken in a certain way, is a confused question.

There is a tendency to think of objectivity in terms of reality constraining discourse. This is the view that objectivity is constituted by sentences in a certain domain *D* having representational content where the reality represented is determinate. More explicitly:

Objectivity: Sentences in a domain *D* are open to objective truth in virtue of the fact that:

(i) Sentences in domain *D* have representational content, representing possible states of affairs featuring certain objects, relations, and properties;

(ii) The domain of reality featuring these objects, properties and relations is determinate in structure.

Taken in one way this is not objectionable for a proponent of GE. *Objectivity* is fine, if by *representational content of sentences* we just mean that they have propositional content, and by *reality having determinate structure* we just mean that for any given property (in the subject domain of *D*) and any given object, either the object instantiates the property or not. Take that way, GE can make perfect sense of our inclination to assert *Objectivity*. GE makes perfect sense of our claims that sentences have propositional content. GE allows that for some domains we assert bivalence or excluded middle. Moreover, it allows that we can express these claims of excluded middle in terms of objects instantiating properties.

But there is another way of taking *Objectivity*. This is that we need to explain the function of the sentences involved in *D* in terms of representation. That representation has an explanatory role in the account of language function. This is to adopt an objectual view of predication—§1-2, §31, and §56. That leads to a theoretical need to assign object-correlates to predicates. We are then inclined to say that it is because predicates in domain *D* are assigned extensions or universal-like entities, and the reality described by sentences in *D* is constituted by objects and sets or instantiated universals, that objectivity is possible for these sentences in domain *D*. But then our theory of objectivity becomes a substantial theoretical enterprise: we need to understand the nature of predicate relations to extensions or universals, to understand the metaphysical structure of reality corresponding to *D*.¹¹¹

GE does not go down this path. It rejects the assumption that representation has an explanatory role in language function. And, saving our talk of objectivity does not require that it does. So the fact that GE does not allow representation to play any explanatory role in the account of language function does not remove its ability to explain the language agency of *objective* or explain our intuitive inclination to assert *Objectivity*. Reality makes our thoughts true or false. GE is perfectly consistent with our making this claim, and indeed explains our disposition to assert it—see chapter 8. There are objective facts of the matter in some domains of talk. GE is perfectly consistent with our making this claim, and indeed

¹¹¹ Wright (1992) and his notion of cognitive command seems to invoke representation as a component of the function of regions of talk to explain objectivity. One can see it in Dummett's (1959, 1963) equation of realism with bivalence, and his idea that we need to grasp some substantial truth-conditions for bivalence to be asserted for verification transcendent propositions.

explains our disposition to assert it—that just resides in our treating Π -properties for sentences in D as non-optional. There is no suggestion here that our mental states are the truth-makers of instances of bivalence. Expressivism about logically complex sentences does not entail that they are true in virtue of our mental states—see §19 and §23. Rather, an instance of bivalence, *either 'S' is true or 'S' is not true*, is made true by the fact that either S or not S . (See chapter 8 and the discussion of truth-making.)

One might be worried that we have not yet justified our assertion of bivalence in any domain. But note that appeal to representation as a component of language function for certain domains of talk will not help us with this issue. The question of justification is really the question: what justifies our belief that reality is determinate in some domain? This is not a question about our belief in our thoughts representing reality, but about our belief that reality is determinate. Expressivism is not better or worse off than any theory about the status of such basic epistemic questions.

Explaining the Non-Optional

A question remains: What is the natural explanation of why speakers treat some class of Π -properties as non-optional whereas others are optional? Here again, no appeal to representation qua component of language function is required. Take Π -properties that underpin the use of empirical predicates such as colour or natural-kind predicates. These tend to be non-optional for us. Why? The explanation is a natural one about our cognitive evolution. The explanation has two parts. (i) Generally, objects are determinate with respect to their colour properties. Leaving aside vagueness, something is usually red or not red, and so on. (ii) Because of (i), there is a payoff to be had in defending the non-optionalness of the Π -properties of colour-attribution sentences. Defending the non-optionalness of such Π -properties means that, when we instantiate them, we defend them without renunciation in the face of dispute. We also assert bivalence for the corresponding assertions. This tends to promote investigation, which often pays off in new perceptual episodes that reinforce defence or rejection of colour Π -properties or lead to the tokening of new (colour) Π -properties. Of course, defence of non-optionalness also occurs for cases in which no possibility of further perceptual encounters is likely. Say, when we speak of the deep past. Certainly, no direct

perceptual encounters are available, but some indirect ones, through causal traces, are possible.

This explanation does not require that we view language function, or the function of the NC system, in *semantic* terms; that is, as relating to the world in some representational way. The embedding is causal. On the other hand we have invoked bivalence, or the material mode correlate of it: determinacy. But that is fine. This is part of our conception of the natural world. We are not falling back into the earlier view of providing a theory of what constitutes objectivity through an appeal to a theory of representation as part of a theory of language agency.

The kind of explanation offered has agents perceptually bumping into parts of reality. But some classes of assertion, evaluative assertions, do not use Γ -states based in perception, they are based in *affect*—see chapter 5. Our explanation of why we tend to treat some Π -properties in the evaluative domain as non-optional cannot take this form. Rather, in the case of evaluative predicates, the answer for why some Π -properties are non-optional has to be an explanation of the role such attitudes play in enhancing human well-being.¹¹²

In the evaluative domain the borderline between the optional and the binding is not absolute. Speakers may differ on where they draw the distinction. Some may treat certain facts about aesthetic beauty as objective, whereas others may not. Who is right? GE offers no theory in terms of language agency that decides. Rather, it offers a theory of why it is that for one domain of Π -properties we tend to find optionality, whereas in another we do not.

¹¹² See Blackburn (1993, 1998) for discussion of similar matters in the context of his quasi-realism.

Chapter 10

Meaning without Nature

§72 Meaning-Theory and Language Agency Revisited

The principle argument, outlined in §1, was basically that a theory of thought's relation to reality must be a theory of the home language. A theory of the home language is not a theory of meaning, because meaning attribution is essentially interpretative, and so presupposes a home language. So a theory of thought's relation to reality cannot be a theory of meaning. Rather, the theory of the home language is a theory of language agency. It is an analysis of the home-speaker, the interpreting and interpretable symbol-using agent, and its activities, seen not from an interpretative perspective but from the only other available, a functional perspective. GE is such a theory. In the course of this book I have attempted to flesh out this argument and the form of GE. In GE, meaning-theory is replaced by a functional analysis of meaning-talk. But we have not yet focussed on meaning-talk. This chapter does so.

I proceed as follows. In §73, I provide a theory of talk of content based in that form of simulation called *S-processing* and its particular species called *L-processing*. That leads naturally into a theory of talk of normativity in relation to meaning in §74. I provide in §75 an explication of the idea that meanings are non-objects in their own sui generis category. In §76, I examine the statue of folk semantics. There is no theory of meaning qua formal normative science that explains the relation of thought, language, and reality, but, there is a theory of meaning in this sense: we, as folk, have a set of commitments about meaning. I look at some theorems of folk semantics, and offer a language-agency explanation of why speakers find theorems of folk semantics deeply intuitive. I then explain why folk semantics, unlike folk physics in relation to physics, cannot be developed into a formal normative science of

meaning. I offer reflections on justification of interpretative judgements in §77, and dissolution of meaning scepticism in §78.

§73 *Talk about Content*

According to GE, attributions of meaning are assertions in which we defend Π -properties based in dispositions to L-process verbal productions—see §17. L-processing is that cognitive activity which goes on when U's language-system simulates features of *locutionary* speech-acts; be they sentences, referring terms or predicates. In each case, U simulates homologues of repertoire dispositions. These are $RD[\Pi^{Gl}, S]$ for declarative sentences, $RD[\Phi, T]$ for referring terms, $RD[\Gamma, F]$ for predicates, and so on. Claims of meaning are assertions whose Π -properties are based in dispositions to L-process sentences or sub-sentential units. So where $RD[\xi, x]$ is a repertoire disposition U associates with the word(s) x :

Means	Π -property identical to or based in:
'E' means x	Disposition to deploy $RD[\xi, x]$ in L-processing tokenings of 'E'.

Let us note again that we are not proposing that meanings are the mental elements $RD[\xi, x]$ in terms of which U L-processes strings. Such properties are integral to the production of meaning-attribution sentences but they are not meanings. Basing—in our technical sense of §16—is an essential part of the account. It is possible to assert that a word means something but not be disposed to L-process the sentence because one does not comprehend the language. Thus, U, a non-Spanish speaker might assert: '*Yo tengo hambre*' means *I am hungry*. But U is not disposed to L-process the sentence in Spanish: *Yo tengo hambre*.

Sentence-meanings, as kinds of linguistic meanings, are considered types. Where U asserts '*J'ai faim*' means *I am hungry*, the use of the first person *I* is not deployed to refer to the speaker. The sentence *I am hungry* is not used to encode a propositional content. Rather, if anything, it encodes a proposition type. This typehood has its correlate at the language function level in the fact that U utters *I am hungry* in a proto-assertion, and the constituent term, *I*, in a mere proto-referring act, in which *I* refers to no-one. What causes production of *I*

is simply the repertoire disposition that underpins use of the first person pronoun, which we can represent as: $RD[\Phi_{\text{self}}, I]$. U's token of Φ_{self} is not a causal factor behind production of this token of I . U is not using I to refer, but merely to proto-refer.

We can distinguish between sentence- and word-meaning and *speaker-meaning*. U can assert *By 'Yo tengo hambre' Gonzalo means that he, himself, is hungry*. Here, properly speaking, U deploys a *that*-clause in the content specification, whereas it is not clear that one ought to use these for sentence-meaning specifications. The S-processing goes beyond the merely locutionary L-processing to processing of a full-fledged illocutionary act. In S-processing Gonzalo's use of the first person Spanish pronoun *yo*, U deploys the term *he, himself*, in combination with the mental element Φ -[*he, himself*]. This mental element is the homologue state that a speaker deploys in simulating another's use of the first person pronoun.

Functional Reduction

Meaning attribution to an agent O is one thing. Functional specification of the language agency of O is another thing. The first occurs within the interpretative stance, the second within the functional stance—§1-2. What is the relation between these two modes of description, and the facts they describe? Meaning-attributions always involve sign-production simulation, and that form of cognitive activity is entirely distinct from the character of the cognitive activity underpinning assignment of cognitive states. For that reason, there is no reduction of meaning-facts to functional facts. We cannot accept anything like:

By 'E' a speaker y means x iff y 's production of 'E' has functional structure F .

Of course, I have proposed that some words have functional essences underpinning them: take negation. Although *not* has a functional essence underpinning it, we can't we assert:

A term 'N' used by O is negation iff O deploys 'N' in a sentence ... N ... to reject Π -[.....]

This is not correct. Any interpretation of a term as negation still depends on the assumption that the sentences to which O attaches the term are meaningful. But this requires an

interpretative judgement about those expressions. If we could make no sense of the sentences to which *O* attached the term, then even if *N* had this functional structure, we would have to doubt that it is negation.

If this is right, how do GE's claims about language agency relate to claims about meaning? The answer is that GE provides necessary conditions for meaning. GE accepts:

If 'E' means *n* then production of 'E' has functional structure *F*.

If the expression 'E' has a functional essence—the way negation has—we can be quite specific about the form of *F*.¹¹³ If 'E' lacks a functional essence—say, it is an empirical term like *red*—then matters are different. We can take, say, all the visually unimpaired speakers who use 'E' meaning red. The mental elements underpinning their utterances will be Γ -states. Then we can say of this class as proposed in §28-9: (a) It will have perception-based primary states where the perception is of the same mode; (b) The physical spectrum of conditions that trigger these perceptual states will fall within a certain cluster on a spectrum; (c) These Γ -states will be functionally related to other Γ -states, which may interact in quite diverse ways within the NC network. These conditions will not give us a functional essence. Some structural features will be common across the set, but not sufficient. But these conditions are very likely to produce homologues in the sense of §31. Two functional states Γ_x and Γ_y possessed by speakers U_x and U_y are homologues if and only if were U_x and U_y to encounter each other then, *ceteris paribus*, they would very likely develop a robust disposition towards mutual L-processing of each other's states.

GE then describes certain kinds of necessary conditions for strings of words 'E' to mean *n*, for some meaning *n*. It also provides a functional analysis of the language agency underpinning meaning-attribution sentences. These might seem but fragments of what we seek in explicating meaning. Are we not missing something? We are not. Do not think of GE as providing a theory of meaning at all. Think of it as offering, potentially, a highly complex

¹¹³ This does not imply that to use 'not' to mean not all a speaker has to do is associate 'not' with this underlying functional essence. That is because a necessary component of meaning not by 'not' is that one combines 'not' with meaningful sentences. To judge that the sentences with which 'not' combines are meaningful is a broad-based interpretative judgement.

functional analysis of individual signalling systems or language agents. In doing that it provides a theory of how these systems interact, which is partly an account of mutual S-processing. That will involve, amongst other things, a largely probabilistic analysis of outcomes of S-processing, that is, the likelihood of certain S-processing episodes in which one system simulates another using certain homologue states. Because of its holistic character and sensitivity to cognitive resources, we cannot map out, in any simple functionally definable way, conditions for mutual S-processing. It has to be probabilistic. That is why meaning is irreducible to functional description.

Public Languages

The issue of word- and sentence-meaning is connected to the question of public languages. Public language has no explanatory role in the theory of language agency.¹¹⁴ *English speaker* or *Spanish speaker* is not a functionally reducible predicate. Can we not define English abstractly as a certain set of (homologue) repertoire dispositions? That is doubtful. English is vague, both diachronically and synchronically, so it is certainly false that there is a unique set of such dispositions. Rather than look for a theory of what public languages are, I propose that we need to analyse the language agency behind terms like *English*. In this case we are concerned with a Γ -state, D_{English} . This element can combine with elements underpinning verb phrases and demonstratives, as in: *Fred is speaking English*, or *This is English*. The canonical ground for triggering of D_{English} is L-processing activity of a certain kind. The paradigmatic case of a speaker who possesses a state D_{English} will be someone who possesses a set of repertoire dispositions and corresponding L-processing dispositions. It is the triggering of the L-processing systems in this set that trigger D_{English} , and underpin issuing of a demonstrative referring term *English*.

We have not provided here a theory about what constitutes a language. GE has no such theory. Like meaning, there is no theory of language in this sense. Evidently, we can say a lot of things about a thing called English. Books are written about English syntax. But that is consistent with the present view. There is a vast overlap in structure between homologue

¹¹⁴ Here there is close agreement with Davidson and his views about common public languages (1994a).

repertoire dispositions of speaker, who form a cluster of English speakers. Nevertheless, there is no reductive analysis of what it is to be an English speaker. Not all of us fit the paradigm. The paradigm is a normative ideal.

Linguistic Meaning and Rules

The language system is not characterised in terms of assimilation of rules, grammatical or semantic. In GE, tacit knowledge of rules has no part in the explanation of language agency. Rule is a normative concept—it is a component of folk semantics. GE certainly does not attempt to reduce rules to regularities of the speaker's language system. The repertoire dispositions I have described in earlier chapters—see §14 and §44—are not rules. They are simply dispositions.

The concept of a rule is related to the concept of linguistic meaning. The sentence *It is a rule in French to use 'rouge' for red things* amounts to the same as *'Rouge' in French means red*. In asserting a sentence like *In L, 'E' means n*, U defends a Π -property based in a disposition to L-process tokens of 'E' identified as belonging to L, using the repertoire disposition $RD[\xi, n]$, that U correlates with n . For ambiguous words, U asserts that *In L, 'E' can mean n*. In such assertions, U defends a Π -property based in a disposition to process tokenings of 'E' identified as belonging to L using the repertoire disposition $RD[\xi, n]$, under some circumstances.

Conceptual and Non-Conceptual Content

GE has room for the distinction between conceptual and non-conceptual content. We might want to affirm that perceptual states have representational content, but that it is not conceptual.¹¹⁵ In terms of GE, non-conceptual content is content possessed by elements of the NC system that have not reached that level of refinement where they are immediate precursors to predicates, such as Γ -states. Thus cognitive scientists may propose that there are cognitive structures involved in perception. They may attempt functional characterisations of these states. But beyond such functional characterisations it may be useful to assign representational content to them. *Representation* is a term in that cluster of terms that are all

¹¹⁵ This is a distinction that has generated some controversy. See Tye (1997).

ultimately rooted in S-processing. We cannot L-process the perceptual states spoken of by the scientists; they are beyond the native ken of our dispositions to L-process. But like all terms of meaning, the Γ -state underpinning *represents* is based in dispositions to L-process. That means it can be applied on grounds that are non-primary, and purely theoretical. Is this backsliding on the part of GE? No. We are not proposing that representational content in such contexts is functionally reducible. Just as we assign reference to each others' utterances, we may assign representational content to mental states of certain kinds. This plays no fundamental explanatory role in an account of language function and cognition. Nevertheless, it can be a useful way of talking.

Assertion that P

Meaning attributions extend to illocutionary act status: identification of an utterance as an assertion or an order, and so on. GE, from its perspective of *Wissenschaft*, provides a structural analysis of assertion in terms of defence Π -properties; that is, signalling disposition to indicate grounds for a Π -property. We have been using this all through the book. This description of assertion is not one that speakers ordinarily have access to. This is theoretical functional analysis of the act of assertion. Ordinary speakers have no theory of Π -properties so their ability to use the term *assert* does not depend on their detecting Π -property states in others. Furthermore, GE does not analyse in NC-system terms what goes on when U asserts *O asserts that P*, for some *P*. To assert *O asserts that P*, is to take on an interpretative perspective since we must interpret O's utterances, and so one in which we L-process O's utterance.

What, then, is the structural analysis of assertions of sentences of the form *O asserted that P* as produced by non-theoreticians? To be able to assert is to be able to produce sentences manifesting a ground-indicating disposition in relation to a Π -property. Speakers can be assertors without any grasp of the character of the underlying system that allows them to assert. Yet, many speakers are aware that they make assertions, that they have beliefs, and that they make judgements. If so, speakers must have a capacity to perceptually register their own inner states. They are able to track their instantiation of Π -properties and ground-indicating dispositions, and repertoire dispositions. Call this a kind of introspective capacity;

a capacity to introspectively track their own states without necessarily being able to conceptualise those states except through categories like *feeling of conviction*, and related notions. Let us suppose that in the case of assertion, U has an introspective capacity labelled D_{Assert} . The triggering of D_{Assert} is just the detection of her own directedness towards defending a certain Π -property and production of a sentence. In short, such states are the primary grounds for introduction of a Π -property of the form: $\Phi-[O] + D_{\text{Assert}} + \Phi-[that S]$. That is, we find a C-law of the form below, where U L-processes S in terms of Ψ :

$$\{\text{Triggering of } D_{\text{Assert}} \text{ by an assertion defending } \Psi\} \Rightarrow \Phi-[O] + D_{\text{Assert}} + \Phi-[that S]$$

Attribution to others is simply the tokening of $\Phi-[O] + D_{\text{Assert}} + \Phi-[that S]$ through another primary ground, whose source is perception of outward behaviour, conversational context, gesture, voice tone, combined with the deliverances of L-processing. You may wonder: how does U make the transference from the first person to the third person? The answer is that no such transference is made. The speaker does not first acquire an ability to use *asserts* based on monitoring of their own inner activity, and then, by analogy, applies that term to others. Rather, there are simply two primary grounds for the state $\Phi-[O] + D_{\text{Assert}} + \Phi-[that S]$; inward activity and outward activity. Both are primary grounds for the production of the state. One might likewise see psychological verbs like *pain* in exactly the same way. There are two primary grounds: one internal, one external.

Understanding

GE is not a theory of understanding. Understanding is a category of folk psychology. GE gives a structural characterisation of what goes on when there is *an event of understanding*: a speaker L-processes speech of another. But there is no reduction of understanding events to events of L-processing, since claims of understanding are connected to attributions of content. We require then, as we did for assertion, an analysis of the language activity that underpins use of the term *understands*. As with all cognitive folk concepts, its deployment depends on introspective capacities. When U L-processes utterance by O, she has a capacity to introspect that she is L-processing. That processing is manifested to her by feelings of a certain kind. Such feelings are one kind of primary ground for Π -

properties with the mental element $D_{\text{Understand}}$. There are other primary grounds. Again we find the same structure of basing described above.

§74 Normativity and Meaning

I have provided a basic sketch of structures underpinning talk of meaning. We might now ask: what of the normative dimensions of meaning? Meaning is meant to be normative. This is one of the reasons why people take meaning to be functionally irreducible. It certainly seems that meaning-claims are closely allied with claims about what speakers ought to do, and what is correct. Thus we have sentences like:

N: ‘Rouge’ means red, and this is red, so it is right to call this *rouge*.

Claims of rightness or correctness are claims in which we defend Π -properties based in affect—see chapter 5. In asserting *It is right to call this ‘rouge’* U defends a state of approval of calling *rouge* the object picked out by *this* and disapproval of not calling it *rouge*, (when called upon to call it something). U implicates that she is disposed herself to call this *rouge*.¹¹⁶ The statement *N* could be held up as evidence that meaning is intrinsically normative. The champion of intrinsic normativity will say that the correctness claim, *It is correct to call this ‘rouge’*, is entailed by *‘Rouge’ means red and this is red*. Is this right? GE can go either way. I consider both options now.

Within GE, the claim that meaning is intrinsically normative can only be maintained if meaning claims have an affective component. If such claims lack an affective component, then claims of meaning are not normative. According to the analysis we have given, the Π -property defended in assertions of *‘Rouge’ means red* does not include an affective aspect. If we stick to that analysis, we cannot treat *N* as an entailment as it stands. There must be a hidden premise. The sentence, *It is right to call red-things ‘rouge’* would suffice. The hidden premise is used to express U’s general approval state of calling red-things *rouge*. Why does U

¹¹⁶ In sincerely asserting *It is right to call this ‘rouge’*, U must really be disposed to defend the Π -property of *This is called ‘rouge’*. U cannot simply have an inclination to do so, but, because she believes she is unreliable in using *rouge*, is hesitant to make the judgment.

have this state? It is just an approval state that she has. We do not have to say that meaning is intrinsically normative for her to develop such an approval state. Perhaps most of us would approve of using *rouge* in the way we use *red*, given we are disposed to L-process *rouge* in terms of the state assigned to *red*. But perhaps we cannot rule out the indifferent individual. They have the dispositions to L-process, but no affective states in relation to the behaviours.

If, on the other hand, meaning claims are normative, then in producing the sentence '*Rouge*' means *red*, U expresses an affective state of some kind. We might propose that in asserting '*Rouge*' means *red*, U not only defends a Π -property based in L-processing, she also defends a disposition to approve of certain uses of *rouge*, uses of *rouge* that coincide with her dispositions to use *red*. If this is the Π -property defended in asserting '*Rouge*' means *red*, then we can explain the intuition that *N* is an entailment. The explanation is that corresponding to *N* there will be an intrinsic C-law, the kind that underpins intuitions of entailment—see §22. This C-law will link Π -properties with an affective component to Π -properties with the same affective component.

What, then, is the showdown between those who want meaning to be intrinsically normative and those who do not, according to GE? There is not much in it. We can go either way. Clearly, we have all sorts of affective attitudes to aspects of language activity. This close association may be taken as showing that in making assertions with *means* we are defending affective states, or it may be taken to show that we are defending states to which we have certain attitudes. Deciding this matter probably requires a more fine-grained language-agency analysis. For GE it is not a crucial matter.

Normativity and Extensions

What makes normative claims about meaning true? The answer is nothing about the structure of language agents and their affective states. What makes them true are normative facts, that is, facts of correctness. Thus what makes it true that in French *rouge* means red is the fact that in French *rouge* means red. This is a fact as good as any other. Of course, facts are non-objects, and we refer to facts by phrases based in nominalization—see chapter 8. But that does not mean that facts are not real and do not exist. We are just released from any need to analyse their natures. They lack natures.

Normative claims are not made true by extensions being fixed in advance by O's grasping a meaning. That is a fantasy of representationalist semantics. If that is right, one source of sceptical challenge has been removed: we have no need to give an account of how a finite fallible mind acquainted with a finite number of instances of a correct application of a predicate can fix in advance all applications of the predicate. We shall expand upon this matter in §78 below.

Correct English

We have a normative attitude to public shared languages. We think that users ought to speak in a certain way. Some ways of speaking are correct, others incorrect. It could be part of someone's idiolect to call cats 'bats'. We say they are not using the term 'bat' correctly insofar as we identify them as an English speaker. Relative to the idiolect, the speaker's usage is perfectly acceptable, but not relative to the designated public language users. In failing to conform, we judge what they do to be incorrect; in conforming, we judge them to be correct. How do we analyse such judgements of correctness?¹¹⁷ Take *JC* below:

JC: In asserting 'This is a cat' in this context, O spoke correct English.

JC is equivalent to a conjunction: *O said 'This is a cat', he spoke English and its utterance was correct*. The Π -property of *JC*, divides into three as follows:

- (i) Π -[U asserted 'This is a cat'];
- (ii) a disposition to process U's utterance using Ψ , utilising L-processing dispositions labelled *English*;
- (iii) a disposition approve of U's action.

The first Π -property, (i), is that of an assertion attribution, which I dealt with in §76. The second component is the L-processing disposition, and the third is the normative component.

¹¹⁷ Gauker (2002) thinks that semantics is not merely the recording of actual practice or dispositions thereto, but an essentially normative business—what speakers ought to do. His argument depends on the point that we do not think of ourselves merely as idiolectic speakers, but rather think of ourselves as constrained by the obligation to conform to standards. As a case for an essentially normative approach to meaning, his argument is right. But all this means for the purveyor of GE is that folk semantics has normative implications. But the philosophical explication of language, thought and reality is not the refinement of folk semantics. It is rather the characterisation of language agency that we have given here.

§75 *Meaning without Ontology: Meanings as Non-Objects*

Here we have been talking in some depth about meaning talk. But we have not said much about meanings. So what are they? Quine sought refuge from postulation of meanings since, by his ontological commitment principle, quantification over them would bring ontological commitment, and ontological commitment to meanings, and the attendant metaphysics was not something he was prepared to countenance. Others are prepared to countenance the metaphysics, and so, face up to what they take to be serious questions about the nature of these curious things. What kind of abstract entities are meanings, and how do they relate to our epistemic capacities? Such is the burden of orthodox semantic thinking. But it is not a burden carried by GE. For GE, Quine's ontological criterion contains a category error. Quantification is not essentially linked to ontological commitment—see §53-4. Meanings are non-objects. What kinds of non-objects are they? My answer will be that they are *sui generis*. They are not objects, events, properties, relations, or facts. This might appear a horrifying prospect to some! Not only are you saying that meanings are non-objects, but that they fall outside of all respectable categories. My interlocutor is speaking from a metaphysical point of view in which the goal is to provide theories of things, and a neat unified description of the world's metaphysical nature. But GE's views about the categorical status of meanings are not voiced from within a theory of what they are. There is no metaphysics of meaning. What we do now is explore the basis for my oracular claim that meanings are *sui generis* things, then show why it should not disturb. We begin with syntax.

Meaning-statements have an interesting syntactic property that we have so far let pass without comment. In a sentence of the form '*E*' means *n*, the grammatical status of the expression referred to by '*E*' and that of *n* are the same. This expression *n* then can be a sentence, predicate, logical particle, noun phrase, and so on. Thus the following are acceptable:¹¹⁸ '*Snow is white*' means *snow is white*, *This token of 'red'* means *red*, '*Plus*' means *plus*, '*If*' means *if*, and '*Leave!*' means *leave!*. This is a curious fact, since

¹¹⁸ It might be objected that the case of sentences '*E*' means *P*, should really have the form '*E*' means *that P*. However, the use of *means* in *means that P* looks like another use of *means*. This is the use in which reason-relations are expressed—§21—as in, *Fred's being late means that the party will be cancelled*.

‘*__ means __*’ is a verb linking what one would naturally have thought were referring terms. But none of the expressions *n*, in the list of meaning-sentences just given, look like referring terms. This curious fact is exemplified further by related constructions, as in *the meaning of ‘if’ is if* and *If is what ‘if’ means*. This can only lead one to ponder upon what the logical syntax of such sentences is.

Within orthodox semantics, one would have to say that the expression *n* in a meaning-sentence ‘*E’ means n* is, despite appearances, somehow used to refer to its own meaning. So in ‘*If’ means if*, the particle *if* functions as a referring term. Somehow *if* in such contexts refers to its own meaning. This view is not very attractive. That would make the sentence, *The meaning of ‘if’ is if*, a tautology—as in *The meaning of ‘if’ is the meaning of ‘if’*—when it does not quite seem to be that. Furthermore, it requires a whole hidden referential apparatus.¹¹⁹

Within GE we don’t have to, and should not, take this path. According to GE, in asserting a sentence ‘*E’ means n*, U defends a Π -property based in a disposition to L-process ‘E’ in terms of a mental element she associates with *n*. U employs the string *n* to display the kind of meaningful unit whose mental element is used by her to L-processes ‘E’. The displaying is simply this: U utters *n* in a proto-act. (Proto-acts are generally characterised in §14 and §44.) So, in ‘*If’ means if*, the second *if* is performed in a proto-act, which is the production of *if* caused by the repertoire disposition that animates her use of *if*, $RD[\omega, if]$, where ω is the mental element that underpins deployment of *if* in assertions of conditionals. In uttering ‘*if’ means if*, U performs *if* in a mere proto-act. U signals that she lacks the state ω ; U is not using *if* in a conditional. Or, to take another example, in using ‘*not’ means not*, the unquoted *not* is used in a proto-act but *not* to perform a negation. In both cases, *if* and *not* are performed in mere proto-acts.

The terms *if*, *not*, and so forth, do not function, then, in meaning-attribution sentences as referring terms, but merely terms displayed in proto-acts. It might still seem strange that we talk about meanings in sentences like ‘*E’ means n*, without referring to a meaning. The

¹¹⁹ Another idea—following Davidson (1969)—might be a paratactic analysis, according to which ‘*E’ means n* comes out as ‘*E’ means the same as this: n*.’

verb *means* does all the work. This is not to say that we do not sometimes use referring terms to refer to meanings.

Take constructions of the form: *the meaning of 'if'*. Such phrases are definite descriptions. Definite descriptions depend upon general terms. To understand how *the meaning of 'if'* functions we need to understand how *meaning* functions. *Meaning* can function as a mass term, *This has meaning*, or a count noun, *This has a meaning*. In both cases we find that the term *meaning* is underpinned by a Γ -state, D_{Meaning} , which has certain conditions of triggering. Broadly speaking, D_{Meaning} is triggered under conditions in which U L-processes words. The mass term use is that where the L-processing is not differentiated. The count noun reading, which is what we have with *a meaning*, is triggered by the individual L-processing of a significant unit, that is, a string with some overall mental element attached to it.

Just as we can have states of contemplating properties—see §52—we can have states of contemplating a meaning. U contemplates a meaning by L-processing a phrase, say, *if*, and maintaining a state that generates the triggering of predicate precursors, such as Γ -[*meaning*] as well as Γ -[*meaning of 'if'*]. These triggered states are bundled cognitively. The bundling is analogous to that of ι -states—see §46—and θ -states—see §52. This cognitive bundling is the source of a Φ -state of the form: $^{\text{No}}[\Gamma$ -[*meaning of 'if'*]...]. There are then further cognitive uniqueness conditions that can be met—see §48—allowing for the introduction of a definite description Φ -state upon which production of *the meaning of 'if'* is grounded. The result is a definite description that can be used to refer to a meaning.

What goes on when we assert *If is the meaning of 'if'*? Such sentences are identity-sentences. In asserting such sentences, U defends dispositions to inter-substitute in non-opaque contexts the speech-acts performed, or tokened, in the production of *If is the meaning of 'if'*—see §27. These acts are the proto-act performed with *if* and the referring act performed with *the meaning of 'if'*. That might seem odd: one is a pure proto-act, the other is a referring term. So how can they be inter-substituted? One of the expressions is a referring term, and one is not. My advice to the reader is not to be drawn in by the idea that the *is* of identity has to be flanked by referring terms. Effectively the commitment of *If is the meaning*

of 'if' is that in all non-opaque contexts in which a mere proto-act performed by *if* appears, a referring act performed with *the meaning of 'if'* can appear and do the same job. Taking the identity-sentence *If is the meaning of 'if'*, we can derive the following identity statements:

I₁: *The meaning of 'if' is the meaning of 'if'*;

I₂: *If is if*.

The orthodox semantic theorist will find I₁ acceptable but I₂ an abomination. Certainly, no artificial language of the kind that formal theorists deal with would allow I₂. The identity sign, for formal systems, cannot be flanked by non-referring terms. But if one does not come with a certain theoretical prejudice and expectation, I₂ is perfectly acceptable. After all, *if* is the meaning of 'if', which is *if*, and so *if* is *if*. But that amounts to asserting I₂. Utterance of I₂ is just a defence of a disposition to inter-substitution of *if* for *if*.

Another seeming paradox is found below. The sentence I₃ is perfectly acceptable:

I₃: *The meaning of 'if' is the meaning of 'the meaning of 'if''*;

I₄: *If is the meaning of 'the meaning of 'if''*.

Is it not an implication of the original meaning claim, *If is the meaning of 'if'*, that we can derive the clearly false I₄ from I₃, by substituting *if* for *the meaning of 'if'* in I₃? It is not. Note carefully what the original meaning-claim *If is the meaning of 'if'* commits us to. To repeat, it commits us to this: a mere proto-act performed with *if* can be inter-substituted with a referring act performed with *the meaning of 'if'*. Now note that the substitution from I₃ to I₄ does not involve a substitution of the required kind. The first token of *the meaning of 'if'* in I₃ is not used in a referring act, it is used in a mere proto-referring act. Recall in claims of the form *n is the meaning of 'E'* the string *n* is always used in a mere proto-act. Our original statement *If is the meaning of 'if'*, licences substitution of the mere proto-act performed with *if* for the referring act using *the meaning of 'if'*. It does not license substitution of the mere proto-act performed with *if* by a mere proto-referring act performed by *the meaning of 'if'*. But that's what we find in I₃.

Quantification over Meanings

We talk about meaning through explicit meaning statements, like '*E*' means *n*, and through referring terms like *the meaning of 'E'*. We can also talk about meaning by *quantifying* over meanings. We can say: *There is something that 'if' means: if*. It is wrong to call this second-order quantification since it is not restricted to predicate position. It is a kind of polymorphous quantification. Such quantification can be made sense of in GE by applying the structures described in §53-4. In asserting *There is something that 'if' means*, U defends a disposition to use a phrase of the form *the meaning of 'if'* in simple sentences, such as *the meaning of 'if' is if*.

Similarly, in asserting *Every meaning of 'bank' is determinate*, U defends a disposition to the effect that where U is disposed to perform a proto-assertion of *T is a meaning of 'bank'*, either as an assertion or a supposition, U is disposed to perform a proto-assertion of *T is determinate*, either as an assertion or a proto-assertion in the scope of a supposition—see §54. There are no restrictions on the character of the terms *T* that may appear in sentences of this form. Thus an instance of the universal statement above is: *River-bank is a meaning of 'bank' and it is determinate*. In this sentence the predicate *river-bank* is a substitution-case for the term *T*. It is used in a mere proto-act.

Here then is a rich array of quantificational and referential forms concerning meanings. But none of this brings ontological commitment. Meanings are non-objects. And as non-objects they lack, we might say, inherent nature. There can be no question, either physical or metaphysical, about what their natures are.

But what category of being do they fall within: are they facts, properties, or relations? The answer is none of the above. Take properties. We talk about meanings by using words, *n*, as in '*E*' means *n*, where *n* is deployed in a mere proto-act. Thus '*Red*' means *red*, but we cannot substitute *redness* to get *The meaning 'red' is redness*. That is because *red* means something different to *redness*. Which distinct properties do these two terms correspond to? One would be hard pressed to answer. Secondly, *the meaning of 'if' is if* and *the meaning of 'not' is not*. But do we want to countenance properties of negativity and conditionality, and

so on? That is just an abuse of the term *property*. I suggest there are similar infelicities to be found with the idea that meanings are relations.

The admission that meanings are not properties or things of any other general category should not send shock waves through our thinking. Rather we should see matters in these terms. There is absolutely no interest in finding a category of being to put meanings in. If one believes in the metaphysics of properties, facts, propositions, and meanings, then one will be interested. That is because, having done natural science and conceptual analysis, there are questions left over; real questions about reality that we need to have answered. But within GE there is little motivation to go down that path.

§76 Folk Semantic and Pragmatic Principles

I have been arguing throughout this book that there is no theory of meaning, qua formal normative science. But this does not rule out talk of meaning as a perfectly legitimate enterprise. Call *folk semantics* that body of doctrine in which we talk of the meanings of words and sentences, the reference of names, the truth and falsity of sentences, and the correctness of using certain terms. There is nothing wrong with our description of language practice as a rational rule-governed activity. GE just rejects this orientation and its attributed properties as deeply explanatory. *Deep* here means that in using such properties as correctness or reference, or in attempting to give theories of what such properties are, we shall not come to an understanding of how language, thought, and a greater reality interconnect.

Folk semantics is our home grown theory of meaning, which to some extent every minimally reflective speaker possesses, or can be made to assent to. Roughly, here are some of the things we assent to about meaning, which I present in the form of theoretical commitments. Truth and meaning are importantly linked in folk semantics. First truth: GE, as a functional analysis of language agency, does not entail any semantic thesis. It does not entail the T-schema, or, the variant: *Where 'S' is used to say P, 'S' is true if and only if P.*

But what GE does is explain the intuitive plausibility of the principle above. We have already seen how in §20. We also explain the intuitive appeal of the thesis that what a sentence says corresponds to its truth-conditions. Where ‘*S*’ says that *P* and ‘*S*’ is true if and only if *P*, then what ‘*S*’ says corresponds to the conditions under which it is true. If we equate what is said in uttering a sentence with what it means, then we get the view that meaning is truth-conditions.¹²⁰

Intimately connected to these claims about sentence-meaning are commitments to compositionality principles. A bulwark compositionality principle of great plausibility is:

The meaning of ‘S’ in a language L is a function of the meaning of its parts and their mode of composition and context.

We can accept it and explain its intuitive plausibility. The L-processing dispositions in relation to a given sentence *S* are partly a function of the L-processing dispositions we have in relation to parts of *S*. It is this functionality that is expressed in asserting the claim above.

Closely related to folk semantics is what we might call folk pragmatics. Principles and theorems like the following are intuitively compelling:

One ought not to assert what is false;

If U asserts that S then U obligates herself to supply reasons for her belief that S.

The first principle fits in with the idea that assertions aim at the truth. That amounts to nothing more than that where we are disposed to defend a Π -property, we are disposed to defend the corresponding truth-attribution in which that Π -property is defended. We approve of people who direct themselves to not producing claims that are false. Confronted with O, whom we judge to assert falsely that *P*, we are inclined to be uncomfortable, unless we have less than honourable intentions. We are strongly motivated to disabuse O. This is not that we

¹²⁰ Educating a speaker about *conventional implicature* can draw the speaker back from this last move—see Barker 2003. Declarative sentences may have more to their meaning than what they *say*; they may also carry *conventional implicatures*. This is a matter of detail and need not concern us here.

hold the speaker's character for disapproval. Rather we dislike the situation of O's asserting falsely. The second is the kind of folk principle that Brandom (1983) uses to analyze assertion. I have argued that it has no role in an account of language agency, but it does have the status of an intuitive principle in folk pragmatics. In asserting *O asserts that S*, U is disposed to defend approval of O's being disposed to provide reasons for her assertion and to disapprove of inabilities to do so. U has these attitudes because they correspond to her own dispositions as an assertor.

Intuitive Appeal and Justification

GE's relation to folk-semantics is this: it explains the language agency of the semantic and pragmatic terms deployed in theorems of folk semantics, and explains why we are disposed to assert them, that is, defend the Π -properties so specified. This is just as we give an explanation of why we are inclined to assert logical axioms and inference forms.

What we have not provided with GE is a justification of these principles of folk semantics. Note, however, that we can give an explanation of why we are inclined to assert that we ought to assert these principles and that these principles are justified. Take the T-schema: '*S*' is true iff *S*. To assert the biconditional is to defend, roughly speaking, C-laws linking the Π -properties of the two sentences. Π -['*S*' is true] and Π -[*P*] are easily derived from each other in the NC system. Furthermore, we approve of speakers who are disposed to assert sentences with easily processed Π -properties. Such sentences are intuitively obvious. These facts about the T-schema and our motivations imply that we are disposed to assert that a speaker ought to assert the T-schema. It is certainly correct that we ought to assert the T-schema. There is an argument: *The T-schema is intuitively obvious. We ought to assert what is intuitively obvious. So we ought to assert it.*

The theory of speech-act structure does not carry, built in, any *ought*-statements. It is motivationally neutral, we might say. So GE does not entail the correctness of the principles. However, speakers with a certain set of motivations will be disposed to assert certain claims about what ought to be asserted. For them, principles will be justified. We have these desires, so, we will judge that they ought to assert these principles. I am going to make such a statement: they ought to assert these principles.

The Explanatory Fortunes of Folk Semantics

What should our theoretical stance to folk semantics and pragmatics be? The orthodox view is as we have described: there is a discipline, call it *Semantics*, that articulates in a formal, normative manner, the conditions of truth of sentences in compositional terms. This discipline aims at being a refinement of folk semantics. Just as physics is a development of folk physics, semantics is a development of folk semantics. We might view semantic theories in some forms as axiomatizations of folk semantics with added theoretical posits that allow that formalisation to run smoothly.¹²¹ Logic can be thought of as an attempt to systematise and rationalise logical intuitions. Semantics under this conception does something similar.

GE rejects these theoretical orientations to folk semantics. It does not dispute folk semantics as a body of commitment. But folk semantics differs fundamentally from folk physics. The former is a body intuitive principles about (physical) objects, events, and processes. Folk semantics is a body of intuitive principles about non-objects: meanings. The subject matter of folk physics is open to theoretical reduction because it is the operation of events and physical object participants: its domain is one of objects. The subject matter of folk semantics, meanings, is not a domain of objects: meanings are non-objects. If meanings were kinds of events we could reasonably ask about their natures. But they are non-objects. This is not to claim that meanings are primitive, simple entities of an abstract kind—which is a problematic thesis. (Compare §56.)

Our theoretical orientation to folk semantics, then, should not be that of our orientation to a domain of events and objects. There is no naturalistic reduction of meanings or facts of meaning. Exploring theorems of folk semantics may provide illumination about something, but this is not what people are hoping for in a philosophical theory of meaning—one that illuminates meaning's place in nature and the status of normative facts.

¹²¹ For example, the program of research based in the work of Lewis (1970)—extended by Jackson (1998) and those associated with the ANU—might embrace the idea that folk semantics is a network of theoretic terms which implicitly define a set of functional properties. The functional properties are then realised in the nature world by some pattern of natural lower level physical properties. Though for some possible problems with that idea see Menzies and Price (200+).

We can live without a theory of meaning in the expanded sense demanded by Semantics. (We cannot live without our folk theory.) The hunger for a theory of meaning in the expanded sense can be satisfied, with some adjustment of our expectations, by a theory of the language agency that allows there to be an interpretative perspective at all. The theory that delivers this account is not itself couched within any interpretative perspective. Rather, it is couched within the functional-causal orientation of natural science. More particularly, it is a theory from within cognitive science, cleansed of any idioms of content.

§77 Interpretation

Meaning attribution is based in a form of simulation. In claiming that a term or sentence means something, a speaker U defends a Π -property based in a disposition to L-process speakers utterances in a certain way, using a repertoire disposition $RD[\Psi, n]$ —in their own sphere of Π -properties. Meaning attribution is largely automatic because L-processing is largely automatic, which is to say, sub-personal. Hence, we can say that we perceive meanings. Nevertheless, there are cases in which attribution becomes a conscious, hermeneutic problem. What justifies a speaker in her attribution of meaning to a speaker O? There are a range of cases that I consider below: (i) meaning attribution in the case of a familiar language; (ii) unknown languages used by humans; (iii) aliens. We shall approach these questions not so much from the point of view of attempting to provide a theory of justification of meaning-attributions, as from the point of view of a theory of justification-talk applied to meaning-attribution.

(i) The Familiar Language

Because of its automatic character, the process being sub-personal and non-inferential, it is not out of place to speak of perceiving speakers' meanings. One cannot but L-process another's language activity, if in a familiar tongue, no more than one can stop seeing something as red. We see colours possessing colour in one case, and sense words possessing

meanings in the other.¹²² It might seem that there is no real analogy between perceiving colour and meaning possession. One might think perception of colour and understanding are not the same since:

- (a) In perceiving the redness of something, something causes U's state, and objects possessing redness;
- (b) Colour perception is based in the natural, real features of objects: their possession of colour properties. But meanings are non-natural and based in arbitrary practice.
- (c) Also those who have not undergone a certain process of inculcation—we call it learning a language—are not able to perceive the meanings possessed by terms.

But none of these represent definitive objections to the analogy. First, regarding (a) there is something causing meaning recognition: language activity. There is no major objection in saying it is sentences possessing meaning that cause meaning recognition. Regarding (b) there is a sense in which language is a conventional matter, but its realisation in creatures is natural. Repertoire dispositions that are one of the prime bases of language agency are natural facts about speakers and the world. Regarding (c) even those who perceive colour have to go through a process of inculcation. All perception—in the sense of *seeing as*—requires some level of cognitive training.

In relation to justification, meaning is in the same boat as colour. Statements of meaning-attribution in the case of a familiar language are akin to statements of colour attribution to objects under normal perceptual circumstances. What justifies an agent in asserting that something is red on the basis of sensory experience of relevant features of an object O? To deal with this question we need to ask first how statements of justification work. In what follows I provide a theory.

Justification

Claims of justification in relation to assertion are claims of epistemic value. To say that O is justified in an assertion is to say that O has good reason for the assertion. Claims of

¹²² Meanings are not properties—see §75—but possessing a meaning is a property. Moreover, we can be direct realists about meaning perception. The things immediately perceived are words-possessing-meaning.

goodness are claims in which we defend states based in approval—see chapter 5. The properties approved of are to do with the activity of assertion in relation to reason. Claims of justification involve defence of Π -properties based in approval of the reasons a speaker has for assertion. We saw that in relation to reason, that assertion of the rightness of inferences involves defence of the C-laws in terms of which we process those inferences. Thus, in asserting *O is justified in asserting/judging/believing P*, U defends a Π -property based in dispositions to defend the C-laws, in terms of which U L-processes the inferential activity that underpins U's assertion.¹²³

We are now ready to return to the question of the justification of colour-judgements. Our goal is to explain the intuitive plausibility—in the sense of §76—of certain kinds of epistemic principles. For example, the following is intuitively plausible:

PC: If an object O seems to be red to U, and conditions appear to be normal in relation to colour, then U is justified in judging that O is red.

From the perspective of a general scepticism, we might put such an epistemic principle in doubt. If a speaker cannot justify her belief that she is not a brain-in-a-vat, she cannot justify her belief that something seemingly before her is red.¹²⁴ But leaving such general scepticism aside the principle holds. **PC**'s intuitive plausibility just follows from what claims of justification express, and the activity of assertion about colour, in relation to the grounding of the Π -properties concerned. Our assertions about colour are just grounded, structurally, in activation of colour-tracking states. If so, we approve of assertions with this character in respect to their grounding. None of this entails that **PC** is correct. All we are providing here is

¹²³ Justification does not imply truth. An assertion can be false but justified. Justification does not imply knowledge. I have not given an analysis of the language agency of knowledge claims. But the rough picture I am attracted to is that in asserting *O knows that P*, U defends a state of approval in relation to treating O's belief or assertion that P as a reason for asserting P. That vague idea and needs development.

¹²⁴ The present treatment of statements of justification obviously does not provide a solution to scepticism. One can always introduce doubt, of a certain theoretical kind, conveyed in sceptical hypotheses that the normal conditions condition is not met.

an explanation of its intuitive plausibility—the same thing we did for principles of folk semantics in §76.¹²⁵

We apply a similar idea to meaning attributions—again put in intentional terms:

PM: If a verbal act *V* seems to U to mean *n*, and conditions appear normal to U in relation to production of meaning, U is justified in assigning the meaning *n* to *V*.

Again, the intuitive plausibility of this principle is based in the activity that underpins our assertions about meaning, based in L-processing. We approve of assertions meeting these conditions, since we ourselves embrace these conditions. Part of being a language agent is L-processing, and, in a sense, we cannot conceive of doing otherwise than L-process. The principle **PM** has a similar air of triviality to **PC**. The *normal conditions* clause boils down to the condition that nothing indicates that we should not make the assertion. This is just an expression of defeasibility. **PC** conveys that we allow appearances to rule—and this is integral to the activity of L-processing. **PC** is a kind of conservative principle: unless something looks bad, go for it.¹²⁶

L-processing is characterised by the poverty of the stimulus. Having only experienced a small, finite number of cases of use by O, U spontaneously L-processes O's terms. But how can U rule out that O's inclination will not radically diverge from her inclinations of use in the future? It seems U cannot. L-processing and S-processing are not (entirely) exercises in predication. It could be that O is a cleverly disguised robot, pre-programmed to produce a set of gestures and sentences which just happen, by pure chance, to accord with U's dispositions to L-process. Sceptical hypotheses are always available. But this does not mean that U is unjustified in making her meaning attributions.

Say O diverges in his behaviour: O uses terms in ways at variance with U's own inclinations. What ought U to say? Should U say *O means N* but is confused, or that he

¹²⁵ Epistemic principles can only be entailed by sets of principles which include at least one epistemic principle. But with respect to the question of explaining the intuitive plausibility of an epistemic principle we can appeal to non-epistemic facts.

¹²⁶ In other words, we can express some of the spirit of the principle thus:

If U is spontaneously disposed, on the basis of L-processing to assert, *V means P*, and U is not disposed to defeat this disposition with some relevant P-property, then U is justified in asserting: *V means P*. Again there is an air of triviality perhaps, but no more to the kinds of inductive inferences we examined in §22.

means something else? This question takes us to our next topic: radical interpretation of human languages.

(ii) *Radical Interpretation of Conspecifics*

The interesting case of interpretation from the point of view of much discussion of meaning is the case of radical interpretation.¹²⁷ In the case of the radical interpreter there is no perception of verbal meaning. The interpreter is not in a position to do that. This does not mean that there is not perception of meaning. That is because, in the case of interpreting humans, there is a vast array of meaningful behaviours in relation to which we perceive meaning. There is a broader symbol-using functional system that extends beyond particular language-symbols: that is, demonstrative behaviours, facial gestures, and comportment to features in environments. In which case, there is automatic simulation in relation to these activities, even if this is not the simulation that we call L-processing, (the kind of simulation that spontaneously occurs in relation to verbal productions of known language). Call this *M-processing*. It is upon the foundation of such spontaneous simulation in relation to the non-verbal behaviours of conspecifics that the possibility of radical interpretation is built. It is because of such repertoire systems that interpretative hypothesis testing is possible. One interacts in primitive conversational forms with subjects of interpretation, and responses are measured using these pre-verbal communication systems, in conjunction with other hypotheses and background assumptions.

Meaning-attributions under conditions of explicit interpretation is guided by finding agreement in use and explaining disagreement. Given that an interpretative subject O seems to use her term *E* in many ways that match U's dispositions to use *red*, U may be disposed to L-process O's term *E* with Γ -[*red*]. But agreement may not be perfect. The disagreement can be explained in two ways: differences of belief and differences of cognitive resources—see §26. Say that although U has good evidence that O's tokens of *E* mean red, O calls grass *E*. Does O not mean red by *E* after all? Does O just have a strange theory of grass? Or does O have a perceptual system that breaks down when viewing grass? What fact determines which

¹²⁷ This is the case that Quine (1960) and Davidson (1984) spend so much time discussing.

hypothesis is the right one? GE gives no simple precise answer to that question: there is none to give. If U's judgements about O's other assertions indicate belief in the magical power of grass to hide its true colour, U may judge that O means red by *E*, but just has strange beliefs about grass. If no strange beliefs about grass are revealed, then U may have to put it down to some cognitive anomaly. These are all testable hypotheses.

If U has access to a description of the functional structure of O's mind, then, U could greatly improve her decision making. Knowledge of cognitive structures can rule out certain hypotheses. Nevertheless, they will not determine the right interpretative answer. Cognitive structures place necessary conditions on meaning, not sufficient.

Explaining difference always occurs against the background of attributions of sameness. We take interpretative subjects to be like us in their underlying natures as language users. That judgement of sameness, as a judgement from within a purely interpretative stance to a subject, is not a judgement about functional identity of cognitive structure. Rather, it is an expression of the possibility of intelligibility. O is like us in that O can be made sense of. What justifies the judgement is simply that O is a conspecific. On the other hand, if U and O find each other intelligible, their signalling systems and underlying NC systems must have a structural isomorphism up to some point.

(iii) Radical Interpretation of Aliens

The really challenging case of radical interpretation is in relation to aliens; extra-terrestrial creatures with whom we share no natural evolutionary history. In this case we differ not only on language systems but also in relation to communication systems of a non-verbal kind. Hollywood cases aside, in relation to such creatures there is no S-processing or M-processing. If there are such inclinations—weird features of the creature make us spontaneously read it as sad or happy—these are almost certainly misleading. In this case, we have a totally theoretical relation to the problem of interpretation.

I think it is possible that we could only proceed by constructing a functional model of the language-agency systems of such creatures. By studying this system, and its interaction with our subjects, we might begin to develop an explicit theory of their S/M-processing

dispositions, and their pre-linguistic natural sign systems. Once we have this system in operation, we might begin to construct interpretative hypotheses.

§79 Dissolving Meaning Scepticism

GE is perfectly consistent with there being facts of meaning. We have assumed that there are facts of meaning. GE explains what goes on when we assert sentences like *O's use of 'red' means red*. The explicit treatment of *fact* is given in chapter 8. But there is a powerful line of argument which concludes that there are no facts of meaning in this sense. The argument has its source in Wittgenstein's examination of rule-following and private languages, but has taken on a crisp and cogent guise in the work of Kripke (1982). This argument has no force against a theory like GE. I now examine the argument and show how it fails.

The essence of Kripke's version of meaning scepticism is encapsulated in the following argument:¹²⁸

P1: Semantic facts are either non-primitive, that is, derived, or primitive.

P2: They cannot be derived.

P3: They cannot be primitive.

C: If so, there are no semantics facts.

We can grant P1. All the argumentative work rests on P2 and P3. I briefly look at Kripke's arguments for P2 and P3.

For P2: What facts about my brain in its environment determine that I mean addition by 'plus'? The natural thought is the natural facts, such that I have a certain disposition to use 'plus' in a certain way, are what determine that I mean addition rather than some other function. But such a disposition cannot constitute the semantic fact. First, it is doubtful that a finite mind, which has only dealt with a finite number of cases of addition, can somehow fix in advance all applications of the term. Secondly, dispositions can go wrong. What makes a

¹²⁸ The formulation of the argument comes from Satne (2005). That work inspires the exposition that follows.

certain use of ‘plus’ resulting in an incorrect sum, an incorrect use of ‘plus’ rather than just a use of ‘plus’ with a different meaning?

For P3: Appeal to primitive semantics facts is also unacceptable. The primitive facts might be qualia states—feelings of a certain kind. But that is hopeless since how can a feeling determine that I ought to apply a term to cases currently unconceived. We might hold that semantic facts are just basic. But that meaning facts are basic is unpalatable since it conflicts with a naturalistic conception of human beings.

If facts are not naturalistically reducible, and cannot be primitive, we must conclude that there are no semantic facts. Conclusion C in the argument above holds. Kripke argues that we have to make sense of our practice of attributing meaning and rule following to others. Since there can be no truth-conditions for claims of the form, *O means by red by ‘red’*, we need to look to assertability conditions. Speakers assert such claims if and only if they judge that they agree in their practice with the attributee. But this sceptical solution entails a communitarian conception of rule following. In so far as we attribute meanings to a speaker we cannot think of that speaker in isolation from others. This seems to rule out community error. Nevertheless, is it not possible that our community of *red* users might all go wrong? Kripke’s sceptical solution will not allow that, since to use the term correctly is just to be in agreement with the community.

Much has been written on how to provide responses to Kripke’s sceptical argument. I will not review it here. I shall argue that meaning scepticism is an artifact of the objectual view of predication—§1-3.

GE and Meaning Scepticism

Is GE susceptible to this argument? GE’s approach to facts of meaning lies in neither of the camps that Kripke sets out in his two premises. It can avoid both because GE does not inhabit the terrain of orthodox semantics. Here is why in more detail:

Non-Reductionism: GE does not reduce semantic facts to facts of disposition. It certainly rejects the kind of analysis according to which *red* in U’s mouth fixes the set of things which U would react to in a certain way—see §30-31. GE rejects the idea that the

meaning of a predicate is an extension. Thus grasping the meaning of *red* is not being dispositionally related, or related in any other way, to an extension.

Non-Primitivism: Within GE semantic facts are primitive: they are irreducible. There is no truth-conditional analysis of sentences of the form *O means red by 'red'*. We can, however, provide an expressive analysis of the agency and activity of using such sentences. We have done so above in §73-75. In asserting *O means red by 'red'*, we defend a Π -property based in a disposition to L-process tokens of *red* in a certain way. This is not an assertion that U agrees with O, or that their dispositions are the same. U and O may in fact differ in their dispositions. U may interpret a speaker O as meaning red by her term, even though U knows she will not call things *red* which U will call *red*. There may be many reasons for this—U may judge that O is under the influence of some drug, that her visual acuity is flawed.

Facts of meaning are irreducible, but not metaphysically primitive. Facts of meaning supervene on natural facts about speakers—their dispositions to use terms. Here our model of supervenience should be that one provided by emotivists. Given my attitudes are sensitive to objects having certain properties, my moral-statement making will be sensitive to natural facts. That then constrains counterfactual judgements—had the moral facts been different, the natural facts would have to have been. We can say the same for meaning-attributions. To assert that O means something by a term is to express attitudes of taking someone a certain way, interpretative attitudes, or rather, L-processing dispositions of certain kinds.

To assert that U means something by *rouge*, is not to say that there is a physical state or otherwise in U's head—or head embedded in an environment—that fixes an extension, and consequently fixes in advance the correct use of the term *rouge*. No mechanism will do that. But fortunately, the meaning of *rouge* is not an extension—a set of correct objects of application of *rouge*. The rightness of a use of *rouge*, judgements of the form *O is right in applying 'rouge' here*, is not fixed by prior facts of extension. Predications using *rouge* are not analysed in terms of set-membership in relation to the set of red things—see §31. Thus grasping the meaning of *rouge* is not about grasping an extension. The fact that O grasps the meaning of *rouge*, or means red by it, is not a dispositional state, it is not a mental capacity, say a primitive one based on the model of perception, to grasp a property, set, or function. It

is such visions of grasping meaning that generate the paradox. And again, these visions have their sources in representationalism, and its objectual analysis of predication, since it readily furnishes the idea that grasping meaning is grasping conditions for correctness fixed by sets or properties qua universals.

Private Language

It might seem that this approach falls foul of Wittgenstein's attack on private languages. But I think not. In GE, the meaning of a sentence is not an inner state. The meaning is not the public speech-act type either. That a speaker uses a term with understanding is not an inner matter either. GE makes no attempt to analyse understanding in terms of the language system. On the other hand, that a speaker believes something or wants something is an inner matter. GE does not rule out scepticism about other minds. But it is not clear that realist theories of mental states could do so, or ought to do so.¹²⁹ But this is not content scepticism, but simply scepticism generated by the gap between theory and evidence.

¹²⁹ Fodor (1998) discusses why theories of content need not rule out scepticism about other minds.

Chapter 11

Realism and Metaphysics

§79 Meaning-Theory and Metaphysics: Partners in Crime

GE eschews representation and truth playing any role in the fundamental account of language function. GE's theory of language activity is anti-representational in the sense that predicates are not assigned worldly correlates to explain their function. Intimately connected to this idea is the denial of objectualism about predication—see §3 and §31. Giving up objectualism about predication enables us to bring in a theory of non-object terms—see §45 and chapters 7 and 8. Referring phrases for properties, relations, facts, and propositions are based in nominalizations. Having given the theory of the language activity using these terms, there is no theory about what the referents of these terms are, either in a physical sense of natural science or a metaphysical sense. To pursue such questions is to deeply misunderstand the nature of how language agency functions. But that means there can be no metaphysics of properties, relations, propositions, or facts. Furthermore, there can be no metaphysics of objects in the sense of persisting things or events. That is because, to believe in the problem of object constitution, as a metaphysical problem, one must implicitly assume that predication functions by representation, and thus that terms like *instantiates* are assigned some worldly correlate as part of their function. But they are assigned no such correlates. This is not to say *instantiates* is meaningless. The language agency of '*Instantiates*' is *meaningful* does not involve assigning worldly correlates to the term *instantiates*. (See §45, §55, §56, §64, for some highlights.)

This is not to claim that all activities that fall under the heading of metaphysics cannot be countenanced by the proponent of GE. That would be an outrageous claim. All we are saying is that there is no metaphysics of properties, relations, facts, propositions, or meanings, and in addition, no metaphysics of object constitution. GE is anti-metaphysical

about the fundamental features of reality. That is to say that there is no metaphysical perspective on these features of reality.

Although GE is anti-metaphysical in this sense, it is not anti-realist. Realism is always a topic specific doctrine: we are realist about the subject matter *R* of some domain of talk. So we can be realists about the physical world, mental states, the past, mathematical reality, and so on—Dummett (1963) and Taylor (2005). Roughly, to be realist about *R* is to hold that there are facts in *R* and truths holding of these facts, where these facts and truths are epistemically independent. Epistemic independence means that these truths and facts about *R* do not depend on what we believe or assert about *R*, or indeed, what, with ideal information, we would believe or assert. (See Taylor (2005)). Broadly speaking, GE is realist about any given domain of reality that we have good reason to believe is a domain of reality. That is because *fact*, *truth*, and indeed, *object*, and *thing*—the conceptual tools we use to describe domains of reality—have no epistemic connotation at all in GE. If so, GE explains our natural realism: our dispositions to assert epistemic-independence claims about the physical world, about mental reality, and so on.

The realism that flows from GE in this sense is a realism without being a form of metaphysical realism. It is not metaphysical because those questions about the nature of properties, relations, and object-constitution involve category error according to GE, despite the fact that these are very tempting errors to make. They are cognitive illusions. If we have to find any word for the view of reality that emerges from GE it is *internal realist*. This does not mean we accept everything that people have associated with that term. Thus we reject Putnam's (1987) characterisation, which involves conceptual relativity—see §70. This conception with its talk of conceptual schemes partly constructing reality is too metaphysical in flavour for anything flowing from GE. Conceptual relativity is an implicitly semantic characterisation of certain features of language activity tied to regions in which faultless disagreement is possible.

In what follows, I develop some of these thoughts. We need to say more about the content of the claim that this form of realism is not metaphysics. That requires saying more about metaphysics, more about realism, and more about the relation between the two.

§80 *Metaphysics*

Both realists and metaphysical realists are meant to accept a basic realism, which we might sum up as the thesis of mind-independent reality. It is just that metaphysical realists believe that mind-independent reality has a structure requiring metaphysical description. But what is metaphysical description? Metaphysics is meant to be the most general articulation of the nature of reality. It defines itself by its subject matter, which is everything. But it approaches everything from the most general perspective possible. It concerns itself with the basic, general features of reality. It is this aspect of metaphysics that leads some philosophers to be slightly free and easy in their use of the term *metaphysical*. For example, assertions of *There are properties* or *There are facts* are labelled *metaphysical*, and one is said to be already involved in metaphysics in making them. That is not true. One needs to be doing more than making such assertions to be doing metaphysics.

The second ingredient of metaphysics is method. Metaphysicians don't do science and they do not merely study concepts, or, in GE's terms the structure of language agency. So what do metaphysicians do? The assumption of metaphysics, methodologically, is this:

Having provided a theory of the structure of language and natural science, there are further questions about the nature of existence of the entities posited both in science and semantics. That is the domain of investigation reserved for *metaphysics*.

In short, metaphysics is about a certain vantage point on investigating what there is and what its nature is. GE is in large agreement with metaphysical orientations about what there is, in that it accepts that there are properties, relations, facts, propositions, events, and property instances, and so on. The difference is in how we conceptualise our explanatory duties in understanding these things. For the metaphysician there is a special task that they must perform in understanding these things that only arises when we have resolved questions of language function and natural science. But for GE no special task arises after such questions are resolved: there are no more explanatory tasks left. How can this be so?

It can only be that the methodology of metaphysics has an assumption about language-function—and here we mean the language that the theorist speaks him or herself. Metaphysics assumes that the language functions in representational terms. The assumption is that there is a discipline called *Semantics* whose task is to set out the pattern of correspondence relations between words and world-parts. Having applied that correspondence analysis to all theoretical enterprises of natural science, we are left with questions of the structures we have deployed in envisioning correspondence relations. The structures are categories, like that of property, relation, fact, proposition, instantiation, and so on. The task of metaphysics is to theorize about these things. And because it implicitly accepts objectualism about predication, it conceptualizes these things as objects with natures. Their natures must be uncovered, and metaphysics is there to do it.

How we view the language in which we pose questions about the world will affect the kind of space we can find for certain forms of investigation of the world. From GE's perspective the task of metaphysics described above can only be a cognitive illusion. That is because GE is totally opposed to any idea that representation can play an explanatory role in the analysis of language function. Why does the cognitive illusion arise? For the same reason that another cognitive illusion arises: the apparent need for a theory of meaning. What lies behind both illusions is the objectual view of predication—see §2-3 and §31—which for whatever reason, is naturally appealing to us.

§81 Realism

Perhaps one will not be convinced that metaphysics can be dropped and realism retained. I shall spell out more the realist claim now, then look at the argument, in §82, that the undermining of metaphysics disrupts realism or involves some illicit inference from the structure of language to the structure of reality.

The basis of realism is a thesis about the epistemic-independence of objects, facts, truth. If objects are mere possibilities of predication, then we are not realist about objects. But in GE objects are not possibilities of predication. If facts are mere possibilities of certified truth, then realism is renounced. But GE does not accept that either. If properties are

similarities that are products of our practices with general terms, then we are not realists about properties. But GE accepts no such view. GE is not a theory about what objects are, what facts are, what propositions are, or what properties are. If no theories of facts, truths, propositions, properties, and so on are given, it is hard to draw any anti-realist conclusions from GE about these things. Typically anti-realist conclusions follow from theories about truth or objects, and so on, which propose that truth or objects are constituted partly by epistemic practice. Thus Dummett (1978) holds that truth can be identified with verification, and McDowell (1996) holds that objects are possibilities of predication. But such views are false. And nothing like them follows from GE. Nothing about what truth is or what objects are is entailed by GE.

Epistemic-Independence

Let us focus on the epistemic-independence of truth and fact. One way of conveying epistemic-independence is through counterfactuals. Thus GE claims that statements of truth involve defence of Π -properties grounded in dispositions to make assertions—20. It might seem then that GE won't be able to respect this thesis; that GE implies a form of idealism. The counterfactuals whose assertability we have to explain have the form below, where S is true and concerns non-human reality where the fact that S is not caused by human belief:

C: Even if everyone had not believed P , then it would still have been true that S .

The only thing that might give the impression that GE entails its falsity, is the false impression that the consequent of **C** is equivalent to: *it would still have been the case that everyone defended Π -[S]*, based on the false belief that *It is true that S* is semantically equivalent to *X defends Π -[S]*. But this is false, and believing it to be true involves making an elementary mistake about the reporting/expressing distinction. The assertion that S is true does not entail that anyone tokens Π -[S]. The claim that something is true is not a claim that a person or persons defend some Π -property. X 's accepting Π -[S] does not explain or constitute S 's being true. If so, **C** is true, which is what we want.

From claims of truth or fact, nothing follows about speakers and their discourse activities or mental states. Thus we deny all of the following:

S is true \rightarrow U defends some Π -[S];

S is true \rightarrow U believes that S ;

It is a fact that S \rightarrow that S is known;

It is a fact that S \rightarrow that S might be known.

We know this because GE does not analysis truth or fact at all. No thesis in folk semantics entails any such principles, and GE respects folk semantics.

Truth-Bearers

Another issue that should be attended to that overhangs from earlier discussion. This is that the truths are truth-bearers. We would want a realism about truth to preserve the idea that even if there had been no thinkers, then many things about the world would still have been true. Here we quantify over things that can be true, that is, propositions. In GE, propositions are non-objects—§55. Propositions are not assertions or assertion-types, which is to say, speech acts—*pace* Barker (2004). Nor are they any kind of object, say abstract entities like Russellian states of affairs or sets of worlds. They lack any metaphysical constitution or any other constitution for that matter. If this is so, there is no basis for the inference below:

It is true that P \rightarrow there is an intentional being who could entertain the thought that P .

This is not correct. A claim about truth is not a claim about any intentional being. Use of the construction, *that P*, does not carry any commitment to there being intentional beings. (Of course, to use it, one has to be an intentional being, but that is another matter.)

Truth-Making

GE accepts truth-making maximalism: the thesis that all truths have truth-makers—see chapter 8. These truth-makers are facts. GE rejects the idea that facts are objects. But this is not the thesis that facts are pale shadows of propositions; that they are truths and so forth. Rather, truths are made true by facts qua non-entities. Is this anti-realist? I think not. We are

not proposing that facts are constituted by, or depend upon, mental states; which is to say, no statement about facts links them to mental or epistemic states of language users.

Bivalence

A facet of realism is assertion of bivalence. But GE does not assert bivalence for all sentences. We saw that bivalence fails for (a) non-objective modes of discourse, from the point of view of objective modes—see §69—and (b) sentences with uniqueness failure for referring terms—see §48. Generally, sentences about the physical world are bivalent.

Whether the sentences concerned are verification-transcendent or not, makes no difference. Dummett's challenge to the realist was that they should provide an account of our grasp of truth-conditions of verification transcendent sentences. As we cannot provide an analysis of this grasp, our assertion of bivalence for verification-transcendent sentences is unjustified. Our disposition to do so must be based in a cognitive illusion that we grasp their content.

GE's reply to Dummett's challenge is that there is no account of grasping meaning, and certainly no account of grasping truth-conditions. Anti-realism comes from reducing meaning to epistemic capacities. Because our epistemic capacities are limited, and because meaning is intimately connected to truth, our access to truth becomes radically limited by the Dummettian move. The world shrinks to that which is epistemically accessible. GE, however, denies that there is any reduction of meaning to epistemic capacities.

The World

The world, the totality of being, is not an object. The world includes objects and events and processes, but also facts, properties, and relations. All these things exist, but not all of them are objects. Not everything that falls under the category of being is an object. Being is not the *superobject*. Being is light in the sense of §64. The claim that reality is not an object is, partly, seen as rejection of the idea of a kind of radical subject-object duality in relation to it. We cannot make sense of the picture in which we place, on one side, our conceptual system, and on the other, reality, where the conceptual system represents reality, and the structure of the conceptual system is read-off that of reality. That picture makes no sense.

§82 *From the Structure of Language to the Structure of Reality?*

I have made some kind of a case that GE is realist. But there are some who might see here a very large fly in the ointment. This is the charge that GE draws inferences about the structure of reality from analysis of language, an entailment of the form:

LR: Language works in way X \rightarrow Reality has structure Y.

How can we allow such inference and retain realism? Take properties: How can we infer what properties are on the basis of true nature of our talk about them? It's as if our choice of language determines what the world is like: a world of properties qua objects as opposed to non-objects. But that is surely not realism.

GE does not accept the inference LR in any form that is objectionable or that undermines realism. Matters turn on how we analyse *structure of reality*. If by this term we mean metaphysical structure, then, since there is no metaphysical structure to infer, the inference cannot hold. There is no structure of this kind to infer. Metaphysicians cannot claim that any pre-theoretical commitment is contravened here, since, the notion of metaphysical structure is not one that is part of ordinary thought.

If by *structure of reality* we mean its fundamental physical structure, then LR again fails. In describing the structure of a language, we describe some part of reality, and so are committed to reality having certain structures in some local region. But there is no inference from this to all the physical structures of reality having a certain form, since the language we describe and the cognitive systems underpinning it may not interact with all of reality. There may be vast tracks of reality, its fundamental physical structure, that a given language never makes contact with.

Language Choice

Suppose that GE is right. It is generally the right kind of approach to analyse the structure of language agency for human natural languages. Given this analysis, property-talk is talk of non-objects. But all we can infer is that properties are non-objects if we speak this language. But who says we have to speak this language? Could there not be a language, call it

M-language, in which sentences are correlated with structured entities, made up of objects, particulars, and universals, involving non-mereological forms of composition? For this language truth *is* correspondence.

Our question now should be, what makes this a language and who speaks it? To the extent that we accept GE, our answer must be that the M-language is a fantasy. Given GE, this just is not a language, and it would be very difficult to know what would constitute language agency for it. The philosophers who say that they can describe this language are deluding themselves. They have not described a language. A few stipulations on paper about truth-conditions will not do.

Maybe GE is false, and we should accept some substantial theory of meaning, in terms of truth-conditions with Russellian propositions. If so, it could be that there really are substantial forms. But this is not to accept GE and assert that maybe there are substantial forms. GE condones the assertion that reality lacks a metaphysical nature—it is empty of metaphysical form. That is not a substantial claim about reality. It is just the claim that a certain methodology and purported mode of enquiry into reality is an illusion. Being anti-metaphysical is not being committed to the idea that reality lacks something.

Language is only a little part of the world

Here's another objection. Language activity is only a little part of the world. Metaphysics is concerned with the world in general. So the conclusions we come to about language can have no bearing on these broader issues of the nature of reality. This objection mistakes the dispute between the proponent of GE and the metaphysician. It assumes that the dispute is over what there is. It isn't. It is about our epistemic, explanatory orientation to what there is. Furthermore, it assumes that a discipline, like metaphysics, can proceed without defining its goals and methodologies, and identify itself purely by what its subject matter is. But no discipline can do that. But if metaphysics must have a self-image, a form of awareness of what its concepts and procedures are, it cannot avoid recognition that it is undertaken using a language, and thus, must have some awareness of how the language functions. But then the issue of language function can become pressing. The thesis of GE is this: once we have the account of language function, we see that the methodology that metaphysics

imagines for itself—and sketched basically above—does not exist. In short, metaphysics is that discipline that proceeds and can only proceed, by a kind of blindness to the true nature of the language in which it is necessarily carried out. It is the unselfconscious discipline par excellence.

§82 Self-Application and Last Words

It is a common idea that language analysis, say in the form of semantics, necessarily brings metaphysical assumptions and commitments with it. Thus, if a theory of language function talks of properties or propositions, then, according to this line of thought, properties and propositions are part of the ontology of the theory. GE is a theory that provides its analysis in terms of properties, relations, events, facts, and so forth. Thus, from a conservative perspective, these commitments are ontological commitments of the theory. Metaphysics precedes theory of language.

But this is entirely misguided. The mere fact that GE quantifies over properties, relations, facts, etc., does not mean these are part of its ontology. Properties, relations, and facts are non-objects. GE applies to itself. GE being a theory whose domain is language activity includes that of all theoretical undertakings, including itself. The assertions that state theses of GE are activities that are describable using GE. The language employed in GE is open to GE's own language-agency analysis. There is no requirement that there be some hierarchy, in which the language used to state GE is excluded from GE's own explanatory domain. In the course of providing a language-agency analysis of some locution, one might assert *O has a disposition to F*. That statement is itself open to language-agency analysis, in which the concept of disposition may appear. Is this circular? No! To use a predicate like *has a disposition to F*, the speaker does not have to have the concept of a disposition. So there is no epistemic circularity. Nor are we providing a reductive truth-conditional analysis of disposition talk. So there is no conceptual circularity.

Realism in a Nutshell

Taylor (2005) have purported to show that Realism is dead. He argues against metaphysical realism using Putnam's model theoretic argument, and attacks forms of internal realism, arguing that they cannot help themselves falling into various forms of anti-realism, and mind-dependence. Taylor's argument, one might think, has a reasonable level of cogency if one restricts oneself to orthodox concepts of how meaning-theory has a role in explaining the relation of thought to reality. From the perspective of GE, there is no cogency. Realism, I tentatively conjecture, is alive and well, and untainted by metaphysics.

References

- Alston, W., 2000. *Illocutionary Acts and Sentence Meaning*. Ithaca: Cornell.
- Altham, J. 1986. 'The Legacy of Emotivism', in G. Macdonald and C. Wright (eds) *Fact, Science and Morality: Essays in A. J. Ayer's Language, Truth and Logic*. Oxford: Blackwell: 275-288.
- Anscombe, G.E.M., 1957. *Intention*. Oxford: Basil Blackwell.
- Appiah, A., 1986. *For Truth in Semantics*. Oxford: Blackwell.
- Armstrong, D. 2004. *Truth and Truthmaking*. Cambridge: Cambridge University Press.
- Ayer, A. J., 1946. *Language, Truth and Logic*, (2nd. ed.) London: Gollancz.
- Austin, J. L., 1950. 'Truth', *Proceedings of the Aristotelian Society*, supp vol 24: 111-29.
- Austin, J.L., 1955. *How to Do Things with Words: The William James Lectures Delivered at Harvard University*. Oxford: Oxford University Press (1962).
- Bach, K., 1987. *Thought and Reference*. Oxford: Oxford University Press.
- Bach, K. and Harnish, R. 1979. *Linguistic Communication and Speech Acts*. Cambridge, Mass: MIT Press.
- Barker, S. J. 2000. 'Is Value Content a Component of Conventional Implicature?' *Analysis*, 60: 185–211.
- Barker, S. J., 2003. 'Truth and Conventional Implicature'. *Mind*, 112: 1–33.
- Barker, S. J., 2004. *Renewing Meaning*. Oxford: Oxford University Press.
- Barker, S. J., 2007. 'Truth-Bearers and the Unsaid', in K. Turner (ed.) *Making Semantics Pragmatic*. Cambridge: Elsevier.
- Barker, S. J., 200+. *An Expressivist Theory of Truth*. Book Ms.
- Barker, S. J., 200+a. 'Leaving Things to take their Chance: Cause and Disposition Grounded in Chance'. Ms.
- Barwise, J. and Perry, J. 1983. *Situations and Attitudes*. Cambridge Mass.: MIT Press.
- Beall, J. C. and Restall G., 2000. 'Logical Pluralism'. *Australian Journal of Philosophy*, 78: 475–93.

- Bennett, J., 1975. *Linguistic Behaviour*. Cambridge: Cambridge University Press.
- Blackburn, S., 1984. *Spreading the Word*. Oxford: Clarendon Press.
- Blackburn, S., 1993. *Essays in quasi-realism*. New York : Oxford University Press.
- Blackburn, S., 1998. *Ruling Passions*. Oxford: Clarendon Press.
- Boolos, G., 1984. 'To be is to be a value of a variable (or to be values of some variable)'. *The Journal of Philosophy*, 81: 430–49.
- Brandom, R., 1983. 'Asserting', *Nous* 17: 637-650.
- Brandom, R., 1994. *Marking It Explicit*. Cambridge, MA: Harvard University Press.
- Brandom, R., 2005. *Between Saying and Doing: Towards an Analytic Pragmatism*. *Locke Lectures*.
- Burge, T., 1973. 'Reference and Proper Names', *Journal of Philosophy* 70: 425-39.
- Cappellen, H. and E., Lepore, 2005. *Insensitive Semantics*. Oxford: Blackwell.
- Chomsky, N., 1969. 'Quine's Empirical Assumptions', in D. Davidson and J. Hintikka (eds) *Words and Objections*: Dordrecht: Reidel.
- Churchland, P., 1986. *Neurophilosophy*. Cambridge, MA: MIT Press.
- Clark, A., 1993. *Associative Engines: Connectionism, Concepts, and Representational Change*. Cambridge, MA: MIT Press
- Cox, D., 1997. 'Putnam, Equivalence, Realism', *Southern Journal of Philosophy* 35: 155-70.
- Cresswell, M., 1985. *Structured Meanings: The Semantics of Propositional Attitudes*. Cambridge, Mass.: MIT Press.
- Currie, G., and Ravenscoft, I., 2002. *Recreative Minds: Imagination in Philosophy and Psychology*. Oxford: Oxford University Press.
- Davidson, D. 1967. 'Truth and Meaning', *Synthese* 17: 304–323, reprinted in Davidson 1984.
- Davidson, D., 1968. 'On Saying That', *Synthese* 17: 130-46, reprinted in Davidson 1984.
- Davidson, D., 1969. 'True to the Facts', *Journal of Philosophy* 66: 748–64, reprinted in Davidson 1984.
- Davidson, D., 1974. 'On the Very Idea of a Conceptual Scheme', *Proceedings and Addresses of the American Philosophical Association*: 5-20. Reprinted in his *Inquiries into Truth and Interpretation*. Oxford: Clarendon Press, 1984.

- Davidson, D., 1979. 'Moods and Performances', in A. Margalit (ed.) *Meaning and Use*. Dordrecht: Reidel, reprinted in Davidson (1984).
- Davidson, D., 1984. *Inquiries into Truth and Interpretation*. Oxford: Oxford University Press.
- Davidson, D., 1984a. 'Communication and Convention'. *Synthese*, 59: 13–17, in Davidson 1984.
- Davis, W. 2005. *Meaning, Expression, and Thought*. Cambridge: Cambridge University Press.
- Dennett, D., 1989. *The Intentional Stance*. Cambridge, Mass.: MIT Press.
- DeRose, K., 1991. 'Epistemic Possibilities', *Philosophical Review*: 581-605.
- Devitt, M., 1996. *Coming to Our Senses*. Cambridge: Cambridge University Press.
- Devitt, M. and Sterelny, K., 1987. *Language and Reality: An Introduction to the Philosophy of Language*. Cambridge, Mass.: MIT Press.
- Divers, J. and Miller, A., 1994. 'Platitudes and Attitudes: A Minimalist Conception of Belief', *Analyst Preprint #2*.
- Dodd, J., 2000. *An Identity Theory of Truth*. Basingstoke: Macmillan.
- Donnellan, K., 1966. 'Reference and Definite Descriptions'. *Philosophical Review* 77: 281–304.
- Donnellan, K. 1974. 'Speaking of Nothing'. *Philosophical Review* 83: 3–31.
- Donaho, S., 1998. 'Are Declarative Sentences Representational?', *Mind* 107: 33–58.
- Dorr, K., 2002. 'Non-Cognitivism and Wishful Thinking'. *Nous* 36: 97-107.
- Dreier, J., 1996. 'Expressivist Embeddings and Minimalist Truth', *Philosophical Studies* 83: 29-51.
- Dretske, F., 1988. *Explaining Behaviour: Reasons in a World of Causes*. Cambridge, Mass.: MIT Press.
- Dummett, M. 1959. 'Truth', *Proceedings of the Aristotelian Society* 59: 141–162, reprinted in Dummett 1978.
- Dummett, M., 1963. 'Realism'. Read at Oxford University Philosophical Society in March 1963 and first printed in Dummett 1978, 145–165.

- Dummett, M. 1976. 'What is a Theory of Meaning?(II)', in G. Evans and J. McDowell, (eds.), *Truth and Meaning: Essays in Semantics*. Oxford: Clarendon Press, 67–137.
- Dummett, M., 1978. *Truth and Other Enigmas*. London: Duckworth.
- Dummett, M., 1981. *Frege : Philosophy of Language* (2nd edn). London: Duckworth.
- Dummett, M., 1991. *The Logical Basis of Metaphysics*, London: Duckworth.
- Egan, A, Hawthorne, J., Weatherson, B., 2004. 'Epistemic Modals in Context'. *Contextualism in Philosophy*, Oxford University Press: 131-170.
- Evans, G., 1975. 'Identity and Predication', *Journal of Philosophy* 72: 343-363.
- Evans, G., 1977. 'Pronouns Quantifiers and Relative Clauses (I)', *Canadian Journal of Philosophy* 7: 777-797.
- Field, H., 1972. 'Tarski's Theory of Truth', *Journal of Philosophy* 69: 347–375.
- Fine, K., 1985. *Reasoning with Arbitrary Objects*. Oxford: Oxford University Press.
- Fodor, J., 1975. *The Language of Thought*. New York: Crowell.
- Fodor, G., 1987. *Psychosemantics*. Cambridge, Mass.: MIT Press.
- Fodor, G., 1998. *Concepts: Where Cognitive Science Went Wrong*. New York: Oxford University Press.
- Fodor J. and Lepore, E., 2001. 'Brandom's Burdens: Compositionality and Inferentialism', *Philosophy and Phenomenological Research*. LXIII: 465-481.
- Frege, G., 1918. 'Thoughts', in his *Logical Investigations*. Oxford: Blackwell, (1977).
- Frege, G., 1892. 'Über Sinn und Bedeutung' in *Zeitschrift für Philosophie und philosophische Kritik* 100: 25-50. Translation: 'On Sense and Reference' in Geach and Black (1980).
- Frege, G., 1892a. 'Über Begriff und Gegenstand', in *Vierteljahresschrift für wissenschaftliche Philosophie* 16: 192-205. Translation: 'Concept and Object' in Geach and Black (1980).
- Frege, G., 1950. *The Foundations of Arithmetic*. J. L. Austin (trans.) Oxford: Blackwell.
- Gasking, D., 1955. 'Causation and Recipes', *Mind* 54: 479-487
- Gauker, C, 2002. *Words without Meaning*. Cambridge, Mass.: MIT Press.
- Geach, P., 1960. 'Ascriptivism'. *Philosophical Review* 69: 221-5.

- Geach P. and Black, M., (eds) 1952. *Translations from the Philosophical Writings of Gottlieb Frege*. Oxford: Blackwell.
- Gentzen, G., 1934. 'Untersuchungen ueber des logische Schliessen'. *Mathematische Zeitschrift* 39, 176-210.
- Gibbard, A., 1990. *Wise Choices, Apt Feelings*. Camb., Mass.: Harvard University Press.
- Grice, P., 1957. 'Meaning'. *Philosophical Review* 66: 377-388.
- Grice, P., 1969. 'Utterer's Meaning and Intentions', *Philosophical Review* 78: 147-77.
- Grice, P., 1971. 'Utterer's Meaning, Sentence-Meaning, and Word-Meaning' in J. Searle (ed.) *The Philosophy of Language*. Oxford: Oxford University Press.
- Grover, D., 1992. *A Prosentential Theory of Truth*. Princeton, N.J : Princeton University Press.
- Gupta, A., 1993. 'A Critique of Deflationism', *Philosophical Topics* 21: 57-81.
- Hale, B., 1993. 'Realism, Anti-Realism, Irrealism and Quasi-Realism' in P. French, T. Uehling Jr., and H. Wettstein (eds.), *Midwest Studies in Philosophy* 12: 201-225.
- Hale, B., 2002. 'Can arboreal knotwork help Blackburn out of Frege's abyss?' *Philosophy and Phenomenological Research* LXV: 144-49
- Hare, R., 1952. *The Language of Morals*. Oxford: Oxford University Press.
- Horgan, T., and Timmons, M., 2000. 'Nondescriptivist Cognitivism: Framework for a New Metaethic', *Philosophical Papers* 29: 121-53.
- Hornsby, J, 1997. 'Truth: The Identity Theory', *Proceedings of the Aristotelian Society* 97: 1-24.
- Horwich, P., 1990. *Truth*. London: Blackwell.
- Horwich, P., 1998. *Meaning*. Oxford: Oxford University Press.
- Humberstone, L., 1992. 'Direction of Fit'. *Mind* 101: 59-83.
- Jackson, F., 1987. *Conditionals*. Oxford: Basil Blackwell.
- Jackson, F., 2000. *From Metaphysics to Ethics: A Defence of Conceptual Analysis*. Oxford: Oxford University Press.
- Jackson, F. and Pettit, P., 1998. 'A Problem for Expressivism', *Analysis* 58: 239-251.
- Jackson, F. and Pettit, P., 2003. 'Locke, Expressivism, Conditionals', *Analysis* 63: 86-92.

- Jackson, F., Graham, O. and Smith, M. 1994. 'Minimalism and Truth Aptness', *Mind* 103: 287-302.
- Kaplan, D., 1977. 'Demonstratives'. in *Themes from Kaplan*, J. Almog, J. Perry, and H. Wettstein (eds.), New York: Oxford University Press, 1989: 481-564.
- Katz, J., 1990. *The Metaphysics of Meaning*. Cambridge, Mass: MIT Press.
- Katz, J., 1994. 'Bearerless Names', *Philosophical Review* 103: 1-39.
- Kölbel, M., 2002. *Truth without objectivity*. London: Routledge.
- Kripke, S., 1980. *Naming and Necessity*. Cambridge, Mass.: Harvard University Press.
- Kripke, S., 1982. *Wittgenstein on Rules and Private Languages*. Oxford: Blackwell.
- Lewis, D., 1970. 'General Semantics', *Synthese* 22: 18-67. Reprinted in D. Davidson and G. Harman (eds) *Semantics of Natural Language* (1972; Dordrecht: Reidel).
- Lewis, D., 1970a. 'How to Define Theoretical Terms', *Journal of Philosophy* 67: 427-446.
- Lewis, D. K., 1986a. *Philosophical Papers*, 2. Oxford: Oxford University Press.
- Lewis, D. K., 1986b. *On The Plurality of Worlds*. Oxford: Basil Blackwell.
- Lewis, D. K., 1986. *On The Plurality of Worlds*. Oxford: Basil Blackwell.
- Lewis, D. K., 1997. 'Finkish Dispositions', *Philosophical Quarterly* 47: 143-158.
- Lewis, D. K., 2001. 'Forget the Correspondence Theory of Truth', *Analysis* 61: 275-80
- Lewis, D., 2003. 'Things qua Truthmakers', plus postscript in Rodriguez-Pereyra, G (ed.) *Real Metaphysics*. London: Routledge: 25-42.
- Lycan, W. G., 1984. *Logical Form and Natural Language*. Cambridge Mass: MIT Press.
- MacFarlane, J., 2003. 'Future Contingents and Relative Truth', *The Philosophical Quarterly* 53: 321-36.
- MacFarlane, J., 2005. 'Making Sense of Relative Truth', *Proceedings of the Aristotelian Society* 105: 321-39.
- MacFarlane, J., 200+. 'Epistemic Modals are Assessment-Sensitive,' forthcoming in an Oxford University Press collection on epistemic modals, ed. B. Weatherson and E. Egan.
- McDowell, J., 1977. 'On the Sense and Reference of a Proper Name', *Mind* 86: 159-85.
- McDowell, J., 1978. 'Physicalism and primitive denotation', *Erkenntnis* 13: 131-52.
- McDowell, J., 1996. *Mind and world*. Cambridge, Mass : Harvard University Press.

- Martin, C., 1984. 'Dispositions and Conditionals', *Philosophical Quarterly* 44: 1-8.
- Menzies, P., and Price, H., 200+. 'Is Semantics in the Plan?', forthcoming in Braddon-Mitchell, D., and Nola, R., (eds) *Naturalism and Analysis*. Cambridge, Mass: MIT Press
- Molnar, G., 2000. 'Truthmakers and Negative Truths', *Australian Journal of Philosophy* 78: 72-86.
- Moltmann, F., 2002. 'Properties and Kinds of Tropes: New Linguistic Facts and Old Philosophical Insights', *Mind* 113: 1-41.
- Montague, R., 1974. *Formal Philosophy; selected papers by Richard Montague*. R Thomason (ed.) New Haven: Yale University Press.
- Mumford, S., 2006. 'Truthmaking and Negative Truth', *Aristotelean Society*.
- Neale, S., 1990. *Descriptions*. Cambridge, Mass.: MIT Press.
- Neale, S., 1995. 'The Philosophical Significance of Gödel's Slingshot Argument'. *Mind* 104: 761–825.
- Oppy, G., 1997. 'The Philosophical Insignificance of Gödel's Slingshot Argument'. *Mind*, 106: 121–141.
- Pagin, P. 2004. 'Is Assertion Social?', *Journal of Pragmatics* 36: 833-59.
- Pagin, P. 2005. 'Assertion', *Stanford Encyclopaedia of Philosophy* (Online).
- Parsons, T., 1980. *Nonexistent Objects*, New Haven: Yale University Press.
- Pearl, J., 2000. *Causality: Models of Reasoning and Inference*. Cambridge: Cambridge University Press
- Price, H., 1988. *Facts & The Function of Truth*. Oxford: Basil Blackwell.
- Price, H., 1990. 'Why "not"?'. *Mind* 99: 221–238.
- Price, H., 1991. 'Agency and Probabilistic Causality', *British Journal for the Philosophy of Science* 42: 157-76.
- Price, H., 1992. 'Agency and Causal Symmetry', *Mind* 101: 501-520.
- Price, H., 1994. 'Semantic minimalism and the Frege point', in Tsohatzidis, S. (ed.), *Foundations of Speech Act Theory: Philosophical and Linguistic Perspectives*, London: Routledge: 132-55.
- Priest, G., 1995. *Beyond the Limits of Thought*. Cambridge: Cambridge University Press.

- Putnam, H., 1975. 'The Meaning of "Meaning"'. In *Philosophical Papers, Vol. 2: Mind, Language and Reality*. Cambridge: Cambridge University Press.
- Putnam, H., 1981, *Reason, Truth, and History*. Cambridge University Press.
- Putnam, H., 1987, *The Many Faces of Realism*. La Salle, Ill., Open Court.
- Quine, W. V., 1948. 'On What There Is', *The Review of Metaphysics* 2: 21-28, reprinted in Quine (1953).
- Quine, W. V., 1951. 'Two Dogmas of Empiricism', *The Philosophical Review* 60: 20-43.
- Quine, W. V., 1953. *From a Logical Point of View*. New York: Harper & Row.
- Quine, W. V., 1960. *Word and Object*. Cambridge, Mass.: MIT Press.
- Quine, W. V., and Ullian, J., 1970. *The Web of Belief*. New York: Random House.
- Ramachandran, M., 1993. 'A Strawsonian objection to Russell's Theory of Descriptions', *Analysis* 53: 209-12.
- Rayo, A., 2002. 'Words and Objects'. *Nous* 36: 436-464.
- Read, S., 2000, 'Truthmakers and the Disjunction Thesis', *Mind* 109: 67-79.
- Recanati, F., 1986. 'On Defining communicative intentions', *Mind and Language* 1: 213-241.
- Recanati, F., 1993. *Direct Reference: From Language to Thought*. Oxford: Blackwell.
- Restall, G. 1996. 'Truthmakers, Entailment, and Necessity', *Australian Journal of Philosophy* 74: 331-40.
- Richard, M., 1990. *Propositional Attitudes*. Cambridge: Cambridge University Press.
- Ridge, M., 2006. 'Ecumenical Expressivism: The Best of Both Worlds?' *Oxford Studies in Metaethics, Volume Two*. Oxford: Oxford University Press: 302-336.
- Rodriguez-Pereyra, G. 2005. 'Why Truthmakers' in *Truthmakers* J. Dodd and H. Beebe (eds.) Oxford: Oxford University Press: 17-31.
- Rosenkranz, S., 2001. 'Farewell to Objectivity: A Critique of Brandom', *The philosophical Quarterly* 51: 232-237.
- Routely, R., 1982. *Exploring Meinong's Jungle and Beyond*. Atascadero, California: Ridgeview Publishing Company.
- Russell, B., 1905. 'On Denoting'. *Mind*, 14: 479-493.

- Russell, B., 1908. 'Mathematical Logic as Based on the Theory of Types', *American Journal of Mathematics* 30: 222-262.
- Russell, B., 1918. 'The Philosophy of Logical Atomism', in *Logic and Knowledge*. London: Allen and Unwin: 177-281.
- Russell, B., 1940. *An Enquiry into Meaning and Truth*. London: Allen & Unwin.
- Sainsbury, M., 2005. *Reference without Referents*. Oxford: Oxford University Press.
- Salmon, N., 1986. *Frege's Puzzle*. Cambridge Mass.: MIT Press.
- Satne, G., 2005. *El Argumento Esceptico: de Wittgenstein a Kripke*. Buenos Aires: Gama Ediciones.
- Schiffer, S., 1972. *Meaning*. Oxford: Clarendon Press.
- Schiffer, S., 2003. *The Things we Mean*. Oxford: Oxford University Press.
- Searle, J., 1969. *Speech Acts: an essay in the philosophy of language*. Cambridge: University Press.
- Searle, J., 1983. *Intentionality: An Essay in the Philosophy of Mind*. Cambridge: Cambridge University Press.
- Simons, P., 2005. 'Negatives, Numbers and Necessity: Some Worries about Armstrong's Version of Truthmaking', *Australasian Journal of Philosophy* 83: 253–261.
- Sellars, W., 1956. *Empiricism and the Philosophy of Mind*. Cambridge, Mass.: Harvard University Press.
- Skorupski, J, 1999. 'Irrealist Cognitivism', *Ratio* XII: 436-459.
- Smith, M., 1994. *The Moral Problem*. Oxford: Blackwell.
- Soames, S., 1987. 'Substitutivity' in J. J. Thomason (ed.), *On Being and Saying: Essays for Richard Cartwright*. Cambridge Mass.: MIT Press: 99–132.
- Sobel, D. and D. Copp, 2001. 'Against direction of fit accounts of belief and desire', *Analysis* 61: 44-53.
- Stalnaker, R., 1983. *Inquiry*. Cambridge, Mass.: MIT Press.
- Stamenov, M. and Gallese, V., (eds). 2002. *Mirror Neurons and the Evolution of Brain and Language (Advances in Consciousness Research 42)*. Amsterdam: John Benjamins Publishing Co.

- Stich, S., 1983. *From Folk Psychology to Cognitive Science*. Cambridge, Mass.: MIT Press.
- Stoljar, D., 1993. 'Emotivism and Truth Conditions', *Philosophical Studies* 70: 81-101.
- Strawson, P., 1950. 'On Referring'. *Mind*, 59: 320–344.
- Strawson, P., 1950a. 'Truth'. *Proceedings of the Aristotelian Society*, 24: 129–156.
- Strawson, P., 1964. 'Intention and Convention in Speech Acts'. *Philosophical Review* 73. 439-460.
- Strawson, P., 1971. *Logico-Linguistic Papers*. London: Methuen.
- Tarski, A., 1937. 'The Concept of Truth in Formalized Languages', in *Logic, Semantics and Metamathematics*. New York: Oxford University Press (1956): 152-278.
- Tennant, N., 1997. *The Taming of the True*. Oxford: Oxford University Press.
- Travis, C., 1989. *The Uses of Sense*. Oxford: Oxford University Press.
- Tye, M, 1997. *Ten Problems of Consciousness: A Representation Theory of the Phenomenal Mind*. Cambridge, Mass.: MIT Press
- Van Fraassen, 1980. *The Scientific Image*. Oxford: Clarendon Press.
- Van Inwagen, P. 1981. 'Why I do not understand substitutional quantification'. *Philosophical Studies*, 39: 281–285.
- van Roojen, M., 1996. 'Expressivism and Irrationality'. *Philosophical Review* 105: 311-35.
- Von Wright, G., 1971. *Explanation and Understanding*. Ithaca, New York: Cornell University Press.
- Wiggins, D., 2001. *Sameness and Substance Renewed*. Cambridge: Cambridge University Press.
- Wittgenstein, L. 1922. *Tractatus Logico-Philosophicus*, translated by C. Ogden. London: Routledge and Kegan Paul.
- Wittgenstein, L., 1975. *On Certainty*. Oxford: Blackwell.
- Woodbridge, J., 2005. 'Truth as a Pretense', in Mark Kalderon (ed.), *Fictionalism in Metaphysics*. Oxford: Oxford University Press.
- Woodward, J., 2003. *Making Things Happen: a Theory of Causal Explanation*. Oxford: Oxford University Press.

Wright, C., 1992. *Truth and Objectivity*. Cambridge, Mass. London : Harvard University Press.

Yablo, S., 1998. 'Does Ontology rest upon a Mistake?', *Proceedings of the Aristotelian Society*, supp. Vol. LXXII: 229-261.

Yablo, S., 2002. 'De Facto Dependence', *Journal of Philosophy* 99: 130-148.

Zalta, E., 1998. *Intensional Logic and the Metaphysics of Intentionality*. Cambridge Mass.: MIT Press.