# Contents

## Part II   Logic, Inference, and Reasoning

# 1
# Inferentialism: State of Play

## 1.1   What is meaning?

We may say, and we often do say, that what makes the difference between a word and a kind of sound that is not a word is that the former *has meaning*. Yet what does this mean? Thousands of books and articles have been written about the nature of meaning and I have no intention to survey them all here (needless to say, this would not be a humanly accomplishable task). For our present purposes it suffices to note that despite the immense efforts that have been put into these investigations no general agreement about the nature of meaning has yet been reached.[1]

The question regarding the nature of linguistic meaning is approached in multifarious ways. The first crossroad is opened up by the question of whether the phrasing 'has meaning' should be taken at face value, as expressing a relation between the word and some preexisting entity called *meaning*. Many philosophers have taken this for granted and have not seen it as disputable. A word, it is often claimed, stands for – or represents, or expresses – its meaning, and the reason it can do so is that we humans are simply symbol-mongerers: we have the peculiar ability to let one thing stand for another.[2] However, the trouble is that it is very difficult to explain, in a non-mysterious way, how we do it and what the relation so established consists in. Is there some unanalyzable power of our minds that is capable of establishing symbols, and is the symbol bound to what it symbolizes by some mental fiber? It seems to me that it remains utterly mysterious not only what the nature of such mental mechanism would be, but especially how the mind could establish such public links as are essential for public language, and what these would consist in.[3]

1

It also seems to me that attempts at explaining the links directly in a naturalistic, especially causal way have not been very successful.[4] Thus I am convinced that even if we disregard direct attacks on the coherence of such representational conceptions of meaning, due to Quine, Sellars, Davidson, and others,[5] there are reasons to be skeptical about the prospects of fleshing out such a theory in a non-mysterious way.

These quick glosses, of course, are not to be taken as serious criticism, their purpose being only to remind the reader that no such approach has gained general acceptance as an explication of the concept of meaning, and that an effort to look elsewhere for another, more plausible explanation of meaningfulness is understandable. (Inferentialism, as presented in this book, is often thought to be a counterintuitive doctrine, so it warrants keeping in mind the problems plaguing rival conceptions of meaning to see that they face obstacles the circumvention of which might outweigh some amount of prima facie counterintuitiveness.)

But, of course, we need not take the meaning talk at face value; we could take it instead as metaphoric talk about some properties of words. Maybe what is characteristic of words – as contrasted with sounds that are not words – is not, or is not literally, that they stand for something, or express it, or represent it, but rather that they have some peculiar property. (The fact that we tend to talk about having a property as about being related to some reification of the property is not in itself mysterious, for it is something we do as a matter of course: we do not hesitate to speak about things having height, color, etc.[6])

One of such explanations, the popularity of which has been on the increase over recent decades (especially thanks to the impact of the legacy of the later Wittgenstein), is that what characterizes a word is the way it is employed within our language games. According to this view, what we call meaning is, in fact, a reification of use. But the trouble is that all kinds of things around us have uses, and yet it seems that to be meaningful as a linguistic expression is something very different from being used, say, as a hammer. Could the difference consist merely in the complexity of the respective uses?

One alternative way of conceiving the difference is to distinguish between items like hammers, which merely have uses, and items like words, which have *roles*, where a role in the sense entertained here is something that is conferred on an item by *rules*. Here is where the underlying idea can be elucidated by comparing words with chess pieces (a comparison frequently used in this book): just as to make a piece of wood (or, for that matter, whatever substance) into a rook it is enough to subordinate it to the rules of chess, what makes a type of sound into an

expression meaning thus and so are again certain rules – rules constitutive of our language games.

It seems to me that this opens up a non-mysterious way to explain meaning (chess does not seem to be a mystery!), and because such ways are in short supply, it is a view we might want to take seriously. Hence the idea is that what makes linguistic meaningfulness (aka having meaning) categorically different from other kinds of usefulnesses are the *rules* that govern the enterprise of language. According to this view, it is the fact that they are constituted by these rules that makes meanings into something special.[7] Moreover, the fact that meanings presuppose a very specific kind of rules (including, be it only in the background, a framework of most basic rules, rules related to what we call *logic*) makes them into a sui generis, into entities of a kind that has nothing comparable in our world.

*Inferentialism*, the topic of this book, is a specific version of this view, according to which the most important kind of rules that constitute meanings are *inferential* rules. The term was coined by Robert Brandom (1994; 2000) as a label for his theory of language, which draws extensively on the earlier views of Wilfrid Sellars (1949; 1953; 1954). (Brandom has engaged the term especially to contrapose it to the common *representationalism*, i.e., the doctrine that meaningfulness consists in representing, i.e. in 'standing for'.) However, the term is also naturally applicable (and is growing increasingly common) within the philosophy of logic,[8] and indeed it is in the context of logic that we can most clearly see how inferential rules are supposed to give rise to meanings. Let us, therefore, now turn our attention to logic.

## 1.2   Inferentialism and logic

Probably the first expression of what we can, retrospectively, see as inferentialism is a passage from the pioneering work of modern logic, Frege's *Begriffsschrift*:

> The contents of two judgments can differ in two ways: first, it may be the way that the consequences which can be derived from the first judgment combined with certain others can always be derived also from the second judgment combined with the same others; secondly, this may not be the case. The two propositions 'At Plataea, the Greeks defeated the Persians' and 'At Plataea, the Persians were defeated by the Greeks' differ in the first way. Even if one can perceive the slight difference in sense, the agreement still predominates. Now I call

the part of the contents which is the same in both, the conceptual content. (Frege, 1879, p. v)

The idea that the (logically relevant) content of a sentence is determined by what is inferable from it (together with various collateral premises) anticipates an important thread within modern logic, maintaining that the notion of content interesting from the viewpoint of logic derives from the concept of inference. This has led to the conclusion that the meaning or significance of logical constants is a matter of the inferential rules, or the rules of proof, that govern them.

It would seem that inferentialism as a doctrine about the content of logical particles is quite plausible. Take, for instance, the conjunction sign; it seems that to pinpoint its meaning, it is enough to stipulate:

$$\frac{A \wedge B}{A} \qquad \frac{A \wedge B}{B} \qquad \frac{A \quad B}{A \wedge B}$$

(The impression that these three rules do institute the usual meaning of $\wedge$ is reinforced by the fact that they may be read as describing the usual truth table: the first two saying that $A \wedge B$ is true only if A and B are, whereas the last one that it is true if A and B are.) This led Gentzen (1934; 1936) and his followers to study the inferential rules that are constitutive of the functioning (and hence the meaning) of logical constants. For each constant they introduced an *introduction* rule or rules (in our case of $\wedge$ above, the last one) and an *elimination* rule or rules (above, the first two). Gentzen's efforts were integrated into the stream of what is now called *proof theory*, which was initiated by David Hilbert – originally as a project to establish secure foundations for logic[9] – and which has subsequently developed, in effect, into the investigation of the inferential structures of logical systems.[10]

The most popular objection to inferentialism in logic was presented by Prior (1960/1961). Prior argues that if we let inferential patterns constitute (the meaning of) logical constants, then nothing prohibits the constitution of a constant *tonk* in terms of the following pattern:

$$\frac{A}{A \; tonk \; B} \qquad \frac{A \; tonk \; B}{B}$$

As the very presence of such a constant within a language obviously makes the language contradictory, Prior concluded that the idea that inferential patterns could furnish logical constants with real meanings must be an illusion.

Defenders of logical inferentialism (prominently Belnap, 1962) argue that Prior only showed that *not every* inferential pattern is able to confer meaning *worth its name*. This makes the inferentialist face the problem of distinguishing, in inferentialist terms, between those patterns that do, and those that do not, confer meaning (from Prior's text it may seem that to draw the boundary we need some essentially representationalist or model-theoretic equipment, such as truth tables), but this is not fatal for inferentialism. Belnap did propose an inferentialist construal of the boundary: according to him it can be construed as the boundary between those patterns that are conservative over the base language and those that are not (i.e., those that do not, and those that do, institute new links among the sentences of the base language). Prior's *tonk*, when added to a language that is not itself trivial, will obviously not be conservative in this sense for it institutes the inference $A \vdash B$ for every $A$ and $B$.[11]

The Priorian challenge has led many logicians to seek a 'clean' way of introducing logical constants proof-theoretically. Apart from Belnap's response, this has opened the door to considerations concerning the *normalizability* of proofs (Prawitz, 1965) and the so-called requirement of *harmony* between their introduction and elimination rules (Dummett, 1991; Tennant, 1997). These notions amount to the requirement that an introduction rule and an elimination rule 'cancel out' in the sense that if you introduce a constant and then eliminate it, there is no gain.

Thus, if you use the introduction rule for conjunction and then use the elimination rule, you are no better off than in the beginning, for what you have proved is nothing more than what you already had:

$$\frac{\dfrac{A \quad B}{A \wedge B}}{A}$$

The reason *tonk* comes to be disqualified by these considerations is that its elimination rule does not 'fit' its introduction rule in the required way: there is not the needed 'harmony' between them; and proofs containing them would violate normalizability. If you introduce it and eliminate it, there may be a nontrivial gain:

$$\frac{\dfrac{A}{A \; tonk \; B}}{B}$$

Prawitz, who has elaborated on the Gentzenian theory of natural deduction, was led, by his consideration of how to make rules constitutive of logical constants as 'well-behaved' as possible, to consider the relationship between proof theory and semantics. He and his followers then developed their ideas, introducing the overarching heading of *proof-theoretic semantics*.[12]

It is clear that the inferentialist construal of the meanings of logical constants presents their semantics more as a matter of a certain know-how than of a knowledge of something represented by them. This may help not only explain how logical constants (and hence logic) may have emerged,[13] but also to align logic with the Wittgensteinian trend of seeing language more as a practical activity than as an abstract system of signs. This was stressed especially by Dummett (1993).[14]

## 1.3   Brandom's inferentialism

Unlike Dummett, Brandom (1994; 2000) does not concentrate on logical constants; his inferentialism extends to the whole of language. As a pragmatist, Brandom concentrates on our linguistic *practices,* on our *language games* and on their place within our human coping with the world and with each other, but, unlike many postmodern followers of Wittgenstein, he is convinced that one of the games is 'principal', namely, the *game of giving and asking for reasons*. It is this game, according to him, that is the hallmark of what we are – thinking, concept-possessing, rational beings abiding to the force of better reason.

To make inferentialism into a doctrine applicable to the whole of language we must make sense of the view that inferences are crucial for *all* kinds of words, including empirical ones. The weakest way to do this would be to claim that an expression cannot be meaningful without playing *some* part in *some* inferences, i.e., that each meaningful expression must be part of some sentences that are inferable from other sentences and/or from which some other sentences are inferable. This is a position that Brandom (2007) calls *weak inferentialism*. This position is clearly not necessarily incompatible with representationalism: believing that to mean something is to represent something is not incompatible with believing that sentences are inferable from other sentences. (Brandom himself conjectures that *everybody* would be a weak inferentialist, but I think that some representationalists would claim that an expression may be meaningful without being part of any sentence, or at least any sentence having inferential links to other sentences.[15])

A stronger version of inferentialism, which Brandom (ibid.) terms *strong inferentialism*, claims that this kind of 'inferential articulation' (i.e., being part of sentences that enter into inferential relationships) is not only a necessary, but also a sufficient, condition of meaningfulness – though construing the concept of inferential rule rather broadly, so that it encompasses 'inferences', as it were, from situations to claims and from claims to actions. (Hence it accepts such 'inferential rules' as *It is correct to claim 'This is a dog' when pointing at a dog*.) This is Brandom's own version, and it is a version to be discussed in this book – though not necessarily in Brandom's own terms, nor sharing his emphases. (Besides these two versions, Brandom also considers *hyperinferentialism*, the claim that inferential articulation is a necessary and sufficient condition of meaningfulness on the narrow construal of inferential rules, and he rejects it as clearly untenable for a language containing empirical vocabulary.)

Why language must be inferentially articulated is because of its crucial role of being the vehicle of the game of giving and asking for reasons. To be able to *give* reasons we must be able to make claims that can serve as reasons for other claims, hence our language must provide for sentences that *entail* other sentences. To be able to *ask for* reasons we must be able to indicate that a claim is in need of being justified, i.e., we must be able to make claims that count as a *challenge* to other claims. (We may, of course, ask for reasons for a claim without explicitly challenging it, but the most primitive way of asking for reasons seems to be a doubt expressed by a challenge.) Hence our language must provide for sentences that are *incompatible* with other sentences; our language must be structured by these entailment and incompatibility relations.

In fact, for Brandom the level of inference and incompatibility is merely a deconstructible superstructure, underlain by certain normative statuses that communicating people acquire and maintain via using language. These statuses comprise various kinds of *commitments* and *entitlements*. Thus, for example, when I make an assertion, I *commit* myself to giving reasons for it when it is challenged (that is what makes it an assertion rather than just babble), and I *entitle* everybody else to reassert my assertion deferring any possible challenges to me. I may commit myself to something without being entitled to it, i.e., without being able to give any reasons for it, and I can be committed to all kinds of things, but there are certain things the commitment to which blocks my entitlement to certain other things.

Brandom's idea is that living in a human society amounts to steering within a rich network of normative social relationships and enjoying

many kinds of normative statuses that reach into many dimensions. Linguistic communication institutes an important stratum of such statuses (commitments and entitlements) and to understand language means being able to keep track of the statuses of one's fellow speakers – to keep score of them, as Brandom puts it. And the social distribution is essential because it provides for the multiplicity of perspectives the intersections of which make the objectivity of linguistic content possible.

This interplay of commitments and entitlements is also the underlying source of the relation of incompatibility: commitment to one claim excluding the entitlement to others. Additionally, there is the relation of inheriting commitments and entitlements (by committing myself to *This is a dog* I commit myself also to *This is an animal,* and being entitled to *It is raining* I am entitled also to *The streets are wet*), and also the relation of inheritance of incompatibilities (*A* is in this relation to *B* iff whatever is incompatible with *B* is incompatible with *A*). This provides for the inference relation (more precisely, it provides, according to Brandom, for its several layers).

Brandom's inferentialism is a species of pragmatism and of the use-theory of meaning: he sees our expressions as tools that we employ to do various useful things (though they should not be seen as *self-standing* tools like a hammer, but rather as tools, like, say, a toothwheel, that achieve useful results only in cooperation with other tools). Brandom gives pride of place to the practical over the theoretical, seeing language as a tool of social interaction rather than an abstract system. Thus any explication of concepts such as *language* or *meaning* must be rooted in an account of what one *does* when one communicates, hence semantics, as he puts it, 'must answer to pragmatics' (1994, p. 83).

What distinguishes Brandom from most other pragmatists and exponents of various use-theories is the essentially normative twist he gives to the pragmatist attitude to language. Thus we can say that what his inferentialism is about are not inferences (as mental actions or episodes of speakers or thinkers), but rather *inferential rules*. This is extremely important to keep in mind, for it is this that distinguishes Brandom's inferentialism from other prima facie similar approaches to meaning, from theories that try to derive meaning from the episodes of inferring rather than from rules.

## 1.4   'Normative' inferentialism vs. 'causal' inferentialism

This brings us to an issue that must be clarified right at the outset. There is a doctrine that, although superficially similar to the Brandomian

inferentialism, should not be confused with it (as, unfortunately, often happens). This doctrine was discussed in the early nineties by Peacocke (1992), Boghossian (1993), and others and it has become popular under the term 'inferential role semantics'.

What this doctrine shares with the Brandomian inferentialism is the conviction that meaning is an inferential role, *viz.* the role conferred on an expression by our inferential practices. However, the crucial difference lies in the aspect of the practices taken to be relevant for the determination of the role. Whereas this theory concentrates on inferences individual human subjects really carry out, or have dispositions to carry out, Brandomian inferentialism concentrates, as we have seen, on inferential *rules*. Let us discuss this difference in greater detail.

Consider the exposition of the theory given by Boghossian (ibid., p. 73):

> Let's suppose that we think in a language of thought and that there are causal facts of the following form: the appearance in $O$'s belief box of a sentence $S_1$ has a tendency to cause the appearance therein of a sentence $S_2$ but not $S_3$. Ignoring many complications, we may describe this sort of fact as consisting in $O$'s disposition to *infer* from $S_1$ to $S_2$, but not to $S_3$. Let's call the totality of the inferences to which a sentence is capable of contributing, its *total inferential role*. A subsentential constituent's total inferential role can then be defined accordingly, as consisting in the contribution it makes to the total inferential role of the sentences in which it appears.

The role, then, is determined by what a subject does, or is disposed to do. In this sense, this theory appears to be a subspecies of 'conceptual role semantics',[16] and thereby a subspecies of the *functionalism* well known in the philosophy of mind.[17] As the functioning that plays the crucial role here is the causal functioning of the human brain (at least insofar as we see mind as supervening on the brain; otherwise it would be a pseudo-causal functioning of the mind), we can call this variety of inferentialism *causal* inferentialism. Hence there is a basic difference between this variety of inferentialism and the normative variety promoted in this book.[18]

The difference is more far-reaching than it might prima facie seem, and to appreciate it we must clarify the nature of the *rules* that play such a crucial role in the characterization of inferentialism. In Chomskian linguistics (and elsewhere too), rules are considered as something that can be directly implemented within the human brain; hence they are

again certain causal mechanisms. But this – and this is the key point – is *not* the notion of rule essential for inferentialism. Rules as understood here are not causal determinants of human conduct, but rather something that it is *not* causally necessary, for any given subject, to follow; it is merely *proper* for the subject to follow them.

However, what does it mean that something is *proper* for a subject? Does this not lead us to some esoteric stratum of reality populated by *proprieties*?[19] Not really; for a propriety is nothing other than a resultant of certain attitudes of many people. It follows that to be able to accommodate proprieties, we need to consider the subject in the context of a society, with the interlocking stances of its members creating a filigree web of social relationships. A human as a social being not only reacts to her natural environment, but also reacts to her peers' reactions. In the course of time she develops what I tend to call *ought-to-be*-thinking (appropriating the terminology of Wilfrid Sellars), which means that she perceives some ways of behaving and acting as agreeable and others as reprehensible.[20] And what I call a propriety, or an (implicit) rule, grows out of such attitudes resonating throughout the surrounding society.

It follows that rules are far from etheric entities beyond the causal order; they are a social, and, especially, what we usually call *institutional*, matter. (As Wittgenstein and the post-Wittgensteinian discussion has taught us, rules in the relevant sense of the word cannot exist other than in the public, social space – for it is only this space that provides for *following the rule* not collapsing into *thinking one is following the rule*.[21]) Thus they are not a matter of merely resonating attitudes, but rather they tend to invoke a superstructure of customized and institutionalized reactions to *improper behavior* ('punishments') as also to *proper* ones ('rewards') that are often wielded in a cooperative manner. And such institutions, though they are a matter of the causal order, are *not* a matter of the causal structures of an individual brain. The existence of a rule is thus a matter of the interlocking patterns of attitudes, actions, and reactions of many people.

Saying that an inferential role of an expression that amounts to its meaning is instituted by such social rules, rather than individual dispositions, has profound consequences. First, there is straightforwardly room for *error*: the way somebody uses an expression may be *wrong*; her individual disposition may not chime with the social rule. And, second, social rules may govern only what is socially accessible; they may govern how we act, not directly what we think. As a result, what is governed by such rules will be the usage of words, expressions, and

especially sentences, not our handling of any mental contents such as beliefs. (Though insofar as beliefs can be thought about as internalized assertions, the subjective mental reality may be thought about as influenced – if not formed – by the intersubjective normative one.)

## 1.5   Is inferentialism circular?

There is an objection often thought fatal to inferentialism, and so we will address it immediately. This is the objection that inferentialism is (viciously) circular: making an inference, so the model version of the objection goes, we must move from some propositions to a proposition, hence from sentence meanings to a sentence meaning; how, then, can inferences *constitute* meanings? To illustrate the crucial difference between the causal and the normative versions of inferentialism, let me consider the difference in the impact this objection on the two versions, in particular the fact that the normative version, unlike the causal one, is largely immune to it.

Consider this objection in greater detail: drawing inferences we typically move, so the story goes, from some beliefs to a new belief, i.e., from propositions to a proposition. These propositions should be definite: it should be clear exactly which propositions they are. I may, for example, move from the propositions that *if it rains, the streets are wet* and that *it rains* to the proposition that *the streets are wet*, and obviously I must be in their possession before I can make this inference. Hence the inference would seem to presuppose propositions, rather than help them into being.

The same holds for concepts insofar as they are seen as constituents of propositions. The proposition that *if it rains, the streets are wet* incorporates implication (rather than, say, conjunction). Hence I must be in possession of the concept of implication already *before* I put together this proposition, and hence before I carry out any such inference. Hence again, claiming that the concept of implication is *forged* by inferences of this kind seems to lead us to a vicious circle: we need implication to be able to substantiate the inferences.[22]

A way of circumventing this objection that might immediately come to mind is to insist that inferences are essentially linguistic, i.e., that they are carried out primarily with sentences, and only secondarily with propositions that the sentences express. But prima facie this does not help, for it would seem that for such a linguistic move to deserve the title *inference* (rather than being just a haphazard passage from one string of letters to another), the sentences must be *meaningful – viz.*

express propositions. So the circumvention would seem to fail because we need propositions *before* we can do any inferences, and again it would seem that inferences thus cannot be constitutive of meanings, especially propositions.

In a recent paper, Boghossian (2014, p. 17) speaks of:

> something that should have been obvious, but that is often lost sight of, including by me…: and that is that reasoning is an operation on thought contents and not on symbols (that have content). That immediately implies that the usual ways of presenting programs of 'inferential role semantics' are confused – a logical constant's role in inference must be explained by its content; its content cannot be explained by its role in inference. Of course, it is always open to an 'inferential role' theorist to give up on the claim that concept possession arises out of the inferential manipulation of symbols, and to insist, rather, that both inference and concepts arise simultaneously out of some pre-cognitive operations on symbols. But it is not easy to see how to flesh out such a view in a plausible way.

I think that what inferentialism provides – or at least struggles to provide – is precisely the fleshing out of this view. Our version of inferentialism presupposes the existence of rules that in turn, as discussed in the previous section, presupposes the social nature of the enterprise of drawing inferences. Inferences are not subjective mental moves, but rather moves in a certain public, intersubjective game, and the rules of the game are constituted together with the constitution of the game itself.

Consider the following 'objection' aimed at chess: chess is played with chess pieces and not with mere bits of wood, hence the piece's role in chess must be explained by its value and its value cannot be explained by its role in chess. Or, put differently, chess moves are not made with bits of wood, but rather with chess pieces, hence we must have the pieces prior to the moves and independent to them. The obvious reply is that it is the rules of chess that confer the values on the bits of wood, i.e., make them into the chess pieces. Hence as soon as we have the distinctions between rules and moves, we may let the former constitute the pieces and the latter then 'operate' on the pieces. In other words, 'the piece's role in chess' is ambiguous, in between the role conferred by the *rules* of chess and the role we confer on it by the ways we use it in games. Once this ambiguity is sorted out, which, in the case of chess, is trivial, the 'objection' looks ridiculous.

And the point is that normative inferentialism can parry the objection of circularity in an analogous way. It can accept that 'a logical constant's role in inference must be explained by its content', whereas at the same time rejecting that 'its content cannot be explained by its role in inference'. We must only sort out the ambiguity of *inference*: the role of a logical constant (or, for that matter, another linguistic item) in inference$_1$ is explained by its content, where inference$_1$ amounts to the inferential moves we actually do with the constant, whereas the content is explained by the role of the constant in inference$_2$, where inference$_2$ amounts to what is *correct* to infer, *viz.* to inferential *rules*.

Of course there is a difference between language and chess consisting in the fact that the rules of chess can be stipulated (in language), whereas those of language cannot have come into being in this way. But this objection does not entail that such rules are nonexistent, and it will be one of the tasks of this book (especially in Chapter 5) to indicate how they could have come into being and established themselves in the form such that this parallel between language and chess turns out to be viable.

Thus, normative inferentialism maintains that for rules, as certain social institutions, there is a story to be told about how they emerged as means of fixations of certain social mechanisms (a story we will tell in detail in Chapter 6), and how they bestowed certain meanings on items the use of which they regulate. No such story appears to be available for a causal inferentialist; the only way a mind can acquire the required dispositions to operate with symbols so that it generates a language (or a 'logic') appears to be some kind of trial-and-error, and due to the holistic nature of linguistic and logical rules, there is no direct feedback that would make this path passable, i.e., that would make it possible to acquire the rules one by one. In contrast to this, the 'social version of the trial-and-error' that leads to the establishment of the rules of language is viable because the 'cultural promulgation' of the social rules makes them survive the demise of any individual mind and hence can wait for the slow feedback given by the external world to the whole system of rules.

Consider another variation on the circularity objection, presented by Fodor and Lepore (2007, p. 682):

> [I]f, as we suppose, Brandom understands his Gentzen-style analysis of content as providing a *possession condition* for 'and' (more generally, for the concept of conjunction), then the treatment would seem to be circular on the face of it. So, for example, we're told that 'to define

the inferential role of an expression "&" ... one specifies that anyone who is committed to P and committed to Q, is thereby to count also as committed as to P&Q, and that anyone who is committed to P&Q is thereby committed both to P and to Q' (Brandom, 2000, p. 62). But since expressions for conjunction (*viz.* '&' and 'and') appear on both sides of each equation, it couldn't be that Brandom's definition of 'and' is what is known by someone who has the word (/concept) and in virtue of which he understands the word (/grasps the concept). Nor, for the same reason, could it be what is *learned* when someone learns the word (/concept).

Of course, to *articulate* the inferential rules governing a logical constant we need a language with its logical vocabulary. But this only says that inferential rules cannot always be explicit, and that there is a sense in which rules have to be implicit to human behavior before they can come to be expressed. This is, of course, a nontrivial assumption and Fodor and Lepore question it; again, it is one of the principal tasks of this book to defend it.

I conclude that the allegation of circularity that is sometimes taken as a knock-down refutation of inferentialism rests on a conflation of the causal and normative versions of inferentialism. If we stick to the normative version, it loses its bite. (Clearly this loads a great deal of the burden of explanation onto the concept of rule, which is itself not transparent, but to unpack it is one of the main tasks of the first part of this book.)

## 1.6  Plan of the rest of the book

In what follows we will be talking about *normative* inferentialism, the kind of inferentialism introduced by Brandom. However, what I will be discussing may not be exactly Brandom's version of inferentialism, nor will it be presented within Brandom's preferred framework. I will explore the foundations of inferentialism in my own way (which I believe is in essence compatible with Brandom's).

Let me return to the trivial example of an inferential role: the role of $\wedge$ that is established by the inferential pattern:

$$\frac{A{\wedge}B}{A} \qquad \frac{A{\wedge}B}{B} \qquad \frac{A \quad B}{A{\wedge}B}$$

There does not seem to be much controversy possible over this simple case: as this pattern can be read as straightforwardly equivalent to the

standard truth table for the connective, nothing seems to stand in the way of seeing it as delimiting the meaning of '∧'.

However, serious difficulties emerge as soon as we move on from this case. We can distinguish two directions along which we can move. One obvious direction is to try to extend inferentialism to expressions other than logical constants, to expressions that can be found in natural languages, especially *empirical* expressions. The most general problem then is to establish how this can be done at all: how empirical expressions that seem to be first and foremost means of *representing* the world can be treated inferentially. Another direction along which to move would keep us within the realm of logic, but would strive to scrutinize how the various kinds of logical constants can be accounted for inferentially. (Already standard disjunction, as we will see, presents a problem for the inferentialist.) Here the basic problems are much more technical. These two directions are dealt with by the first and second parts of this book, respectively.

Thus, in the first part we address the general problems of inferentialism with respect to the whole of natural language, including a discussion of the very sources of normativity that underlie the inferential rules governing it. I try to generalize the inferential construal of logical constants to the rest of the vocabulary, thereby reaching an inferential explication of the concepts of meaning and language. Discussing how the concept of inferential rule can be generalized so as to encompass the empirical dimension of language leads to the conclusion that the whole of language (in contrast to its purely logical part) must be understood as a system of *embodied* rules, i.e., of rules that constitutively incorporate the world. I broach the problem that at least some of the rules of our language are bound to remain merely implicit in our linguistic practices. I point out that such rules are carried by the *normative attitudes* of people, leading to the conclusion that normative attitudes result from the fact that we do not only state that something is the case, but also endorse that something *should be* the case. There follows an analysis of how the rules of our languages interlock to provide for propositions and concepts, and finally I discuss this fact from the evolutionary perspective.

The second part of the book concentrates on the inferentialist approach to the meaning of logical constants and to logic in general. We start from the disambiguation of the term *inference* and from the discussion of the relationship between inference and consequence. (It is often claimed that the necessary discrepancy between inference and consequence, as documented by the results of Tarski and Gödel, shows the irreducibility of the truly semantic notions to the 'syntactic' ones, but we argue that the relationship inference vs. consequence can be

construed as the relationship between two layers of inference, namely one based on the usual strict concept of rule and the other based on a looser concept.) In view of the Priorian argument that not every inferential pattern is capable of constituting a reasonable logical constant, I consider the problem of characterization of 'benign' (or 'semantogenic') patterns, as contrasted to the *tonk*ish, 'malign' ones. I conclude, in accordance with Belnap's reaction to Prior's problem, that the inferential patterns constitutive of logical constants should be conservative. I also discuss the kinds of logical constants that can be introduced in terms of inferential patterns straightforwardly, and introduce a hierarchy of inferential (and consequently semantic) systems yielded by relaxation of the concept of inferential rule.

I also offer a story (based on the idea of Brandom) explaining *why* it is that the patterns constitutive of logical constants should be conservative. My claim is that it is because the role of logical vocabulary is basically *expressive* – that its *raison d'être* is to make explicit the inferential relationships between sentences implicit to our non-logical concepts. Exploiting this idea, I then discuss the notion of 'native' logical operators (generic operators needed for making the inferential relationship explicit) and I use it to shed new light on the differences among logical systems. I draw some consequences of this construal of logic for the very nature of logic: I claim that human 'possession of logic' should not be understood as a matter of knowledge of logical laws, but rather as a matter of possessing a certain kind of language, governed by a certain intricate set of interlocking rules. Finally I turn my attention to the interconnection between logic and reasoning, and concluding that the laws of logic are not rules of reasoning in the sense of tactical rules, I claim that they are rather rules that constitute the 'material' that is a necessary vehicle for reasoning.

Individual chapters of the book have absorbed some of the materials (mostly substantially reworked) that I have earlier published in articles. Aside from material from articles that had the character of prepublications and were printed in volumes that were not widely accessible, this also concerns some genuinely published papers. Thus, in the first part of the book Chapter 3 contains bits of the paper 'Inferentialism and Compositionality of Meaning' (*International Review of Pragmatics* 1, 2009, pp. 154–181), while Chapters 4 and 5 include some scattered fragments from 'The use-theory of meaning and the rules of our language games' (K. Turner, ed.: *Making Semantics Pragmatic*, Emerald, Bingley, 2011, pp. 183–204); Chapter 4 incorporates some material from 'Inferentialism and the Normativity of Meaning' (*Philosophia* 40, 2012, pp. 75–97);

Chapter 5 includes parts of 'Semantics without Meaning?' (R. Schantz, ed.: *Prospects of Meaning,* de Gruyter, Berlin, 2012, pp. 479–502); while Chapter 6 overlaps with 'The Enigma of Rules' (*International Journal of Philosophical Studies* 18, 2010, pp. 377–394). In the second part of the book, Chapter 8 draws on the material published (in greater detail) in 'Inferentializing Semantics' (*Journal of Philosophical Logic* 39, 2010, pp. 255–274); Chapter 9 partly overlaps with 'What is *the* logic of inference?' (*Studia Logica* 88, 2008, pp. 263–294), while Chapter 11 contains a small fragment of the paper 'Logic and Natural Selection' (*Logica Universalis* 4, 2010, pp. 207–223).

## 1.7   Summary of Chapter 1

In this chapter we have introduced the general concept of inferentialism as it has come into circulation both in logic and in philosophy of language. We have also attempted to clear away the most widespread misunderstandings, particularly stressing that inferentialism is not what has occasionally been called *inferential role semantics*. What was termed *inferentialism* by Brandom, and what we address in this book, is the doctrine that identifies meanings with roles vis-à-vis inferential *rules*, whereas the kind of inferentialism envisaged by Boghossian, Peacocke, and others is interested in roles with respect to inferences actually or potentially carried out by speakers. We have stressed that a proper understanding of this preempts the most frequent kind of objections to inferentialism, namely the allegations of circularity – objections that have no obvious force against the normative version of inferentialism we present.

# Index