Natural Language and Logical Consequence
An Inferentialist Account

Abstract The relationship between natural language and formal logic is of central epistemological and methodological importance, both for research in formal logic and natural language semantics. Despite this, it sees little explicit attention in contemporary philosophy of logic. In a recent paper, Michael Glanzberg (2015, in: Caret & Hjortland (eds.), *Foundations of Logical Consequence*, 71-120) has explicitly argued that natural language cannot contain a relation of logical consequence. The chief reason he gives is that natural language semantics cannot both model linguistic competence *and* contain a sufficiently large range of models to include such a relation, model-theoretically conceived. I counter this by presenting a use-theoretic approach to linguistic meaning and proof-theoretic conceptions of logical notions. I argue that in this use-theoretic setting, the ascription conditions for logical constants and consequence relations do not exceed the meaning-theoretic background. I present this argument both for so-called 'modest inferentialism', which combines inferentialist (meta)semantics with truth-conditional semantics, as well as the Brandomian variant of 'strong inferentialism'. I further argue that here are good reasons to be optimistic that natural language *does* contain at least one logical constant, which in turn generates a relation of logical consequence.

1 Introduction

Natural language is ubiquitous in the study and teaching of formal logic. This might plausibly suggest a corresponding epistemological and methodological indispensability. Consider, for example, the situation of an introductory class on logic, where I wish to teach my students the logic of conjunction. In order to do so, I need to rely on my student's prior understanding of and. For otherwise they could not make sense of the Tarskian characterisation of $\wedge$

$A \wedge B$ is true if, and only if, A is true *and* B is true.

or be assured that the inferences formalised in the standard introduction and elimination rules for conjunction

$$\frac{A \wedge B}{A}\wedge\mathrm{E} \qquad \frac{A \wedge B}{B}\wedge\mathrm{E} \qquad \frac{A \quad\quad B}{A \wedge B}\wedge\mathrm{I}$$

are valid. Moreover, discussing the validity of logical laws is done by means of arguments coined in *natural* language, not a *formal* one, as the countless examples from the literature demonstrate (e.g. McGee, 1985; Ben-Yami, 2010).

In fact, we can make an even stronger observation. In order to *understand* a piece of formalism in logic *as* pertaining to such things as correct reasoning – and not just as some sort of (uninteresting) algebra – the formalism must be translatable into natural language. In this vein, as Hanfling observes, Russell's claim that his analysis of sentences such as Men exist as $\exists x(man(x))$ could not have been done in ordinary language, but only by inventing a new one, is dubious. For if the statement in the new language is translatable into the old, the claim could have already been made in the latter. If it is not translatable, then it remains mysterious how we could understand the claim to begin with (Hanfling, 2000, 160). This thought naturally generalises, thus, it would appear that understanding formal logic is parasitic on a prior understanding of natural language.

Against this, it is sometimes claimed that the focal point of formal logic research, viz. logical constants and consequence relations, correspond *at best* weakly to any natural language counterparts. For example, it seems that and has a much more complicated role in inference or truth-conditional contribution than its formal cognate $\wedge$ (e.g. Strawson, 1952, 79-82). Hence, perhaps, we should expect little congruence to be forthcoming concerning formal logic and natural language. Accordingly, it could be argued that natural language should not play a decisive role in logic research, and appeals to notions of formal logic in the study of linguistic meaning are on equally unstable ground.[1]

These two perspectives on the relationship between natural language and formal logic stand in obvious tension with each other. Crucially, if it should turn out that there is at best a weak correspondence between the two, this would have radical implications

[1]A common appeal is to logical analyticity. For an example, cf. Aloni et al. (2017).

about our practices in formal logic, be it teaching or research, and formal semantics. For with what *right* can we make appeals to logical notions in semantics or rely on natural language when teaching and researching logic, if natural language is either devoid or only distantly related to formal logic? Due to these ramifications, the nature and implications of this tension deserve serious attention.

Michael Glanzberg (2015) has recently argued that there is a serious mismatch between natural language and formal logic. He argues that natural language *cannot* contain a relation of logical consequence (Glanzberg, 2015, 71f.). In this paper, I will argue that this conclusion is false. I locate the issue with Glanzberg's argument in his premisses concerning an adequate account of logical notions and his general outlook on the nature of linguistic meaning. I conclude that *if* we reject his adoption of mainstream referential truth-conditional semantics and model-theoretic conceptions of logical notions in favour of use-theoretic and proof-theoretic ones, respectively, the mismatch can be averted. Furthermore, this *inferentialist* account should provide the necessary tools to identify logical constants and consequence relations in natural language.

To this end, I will proceed as follows. In section 2, I will provide a detailed discussion of Glanzberg's arguments and assumptions. Next, section 3 introduces inferentialism as a semantical and metasemantical account of linguistic meaning. Given that I will have rejected referential truth-conditional semantics at this point, I will also reject model-theoretic conceptions of logical notions and introduce the proof-theoretic ones, based on logical inferentialism. In the penultimate section 4, I will argue that, in contrast to Glanzberg's case, there is no 'deep' mismatch between natural language and formal logic. Moreover, we can be optimistic that we can both find logical constants and consequence relations in natural language. Section 5 will conclude with an outlook on future issues that stem from this investigation.

Before we proceed, two remarks are in order. First, due to reasons of space, and in contrast to Glanzberg, I will focus only on sentential connectives and ignore quantifiers. This has the added benefit of avoiding the difficult issue of the logicality of second-order logic, arguably a common phenomenon in natural language. Second, there are obvious ways in which natural language and formal logic mismatch, e.g. on the level of syntax. Thus, given its central relevance in the above considerations, I shall only focus on the level of *meaning*.

2 Glanzberg's Arguments

In his *Logical Consequence and Natural Language* (2015), Michael Glanzberg argues against what he calls the *logic in natural language thesis* (henceforth: LNLT):

> A natural language, as a structure with a syntax and a semantics, thereby determines a logical consequence relation.
> (Glanzberg, 2015, 75).

The general idea is to use logical consequence – "the core feature" of a logic – as a measuring tool for the closeness of natural language and formal logic (Glanzberg, 2015, 71). To this end, Glanzberg makes a series of assumptions, to which we turn first.

2.1 The Meaning-Theoretic Assumptions

Glanzberg assumes that truth-conditions capture an important aspect of "what [speakers] know when they know what their sentences mean", hence they form an ingredient in linguistic competence (Glanzberg, 2015, 85). He rules out other approaches to linguistic meaning, such as use theories, on the basis of their lack of empirical success.[2] Lastly, he also assumes *compositionality* and *referentialism*. The former is the idea that the truth-conditions of a sentence are determined by the referents[3] of its constitutive expressions and their mode of (syntactic) combination (cf. Partee, 2004, 153), while the later is the position that those referents *are* the meanings of subsentential expressions (cf. Glanzberg, 2015, 87f.). Coming from these assumptions, it should be no surprise that Glanzberg also endorses model-theoretic conceptions of logical notions, and considers them the most plausible route to logical consequence in natural language (Glanzberg, 2015, 75, 79).

According to these conceptions, a logical consequence relation obtains whenever the premisses preserve the designated values across a space of models (cf. Beall et al., 2019, sect. 3.1). This accounts for the necessity constraint on logical consequence, i.e. the idea that when such a relation obtains, the conclusion follows by necessity from the premisses (cf. Beall et al., 2019, sect. 1). In addition to this, logical consequence should also be *formal*, i.e. it should obtain only in virtue of the form of the relevant sentences, not their specific content (cf. Beall et al., 2019, sect. 2). Glanzberg follows other authors in taking the relevant form to be provided by the *logical constants* (Glanzberg, 2015, 76f.), about which he further relies on the model-theoretic conception, too (Glanzberg, 2015, 77). Specifically, logical constants are those expressions whose extensions do not vary under arbitrary permutations of the domain of discourse (cf. MacFarlane, 2017, sect. 5).

In summary, we see that Glanzberg endorses standard referential truth-conditional semantics as well as the accompanying model-theoretic conceptions of logical notions.

---

[2]As a quick rejoinder, two observations are important. First, use theories are generally less developed than their 'referentialist' cousins (cf. Murzi and Steinberger, 2017, 214). As such, it must be checked whether the lack of empirical success is due to a genuine failure to make correct predictions, or due to a lack of testing. Second, it is important to understand the explanatory ambitions of each theory correctly. The *explanans* that especially *inferentialists* start with – the inferences accepted as correct by the linguistic community – forms the *explanandum* in standard formal semantics. Thus, we would clearly beg the question against the inferentialist if we accuse them of failing to account for correct inferences (cf. sect. 3 below).

[3]In what follows, by referent I shall mean any kind of denotation that is assigned to a linguistic expression, not just in the narrow of sense of proper names and their bearers.

2.2 The Argument from Logical Constants

In order to argue against the LNLT, Glanzberg makes a further point concerning truth-conditions. Consider the following *relativised* truth-conditions for an ordinary sentence, as they can be encountered in historical model-theoretic semantics, and where $\mathcal{M}$ is some (first-order) model:

$$\textsf{Erin is happy} \text{ is true in } \mathcal{M} \qquad \text{iff} \qquad \textsf{Erin}^{\mathcal{M}} \in \textsf{happy}^{\mathcal{M}},$$

i.e. the sentence is true if, and only if, whatever is picked out by Erin falls under whatever subset happy designates. The issue for Glanzberg, following Lepore (1983, 173-178), is that a speaker can understand these truth-conditions without thereby knowing what the sentence is about. For by such conditions alone, it is not determined who Erin is, or what happiness consists in. Thus, such truth-conditions, i.e. as relativised to models, cannot capture linguistic competence (Glanzberg, 2015, 88).[4] In order to avoid this unwelcome consequence, we should simply forgo the relativisation, and instead consider so-called *absolute* truth-conditions, such as:

$$\textsf{Erin is happy} \text{ is true} \qquad \text{iff} \qquad [...] \qquad \text{iff} \qquad \text{Erin is happy},$$

where [...] is a place-holder for the formal interpretation(s) of the sentence. Since it involves the relevant referents and concepts, it cannot depart from modelling linguistic competence. Furthermore, as Glanzberg claims, both contemporary "Montagovians" and "Neo-Davidsonians" work within *absolute semantics*, which only considers absolute truth-conditions, such as the example above. Hence, so Glanzberg concludes, the *status quo* of formal semantics is that it deals in absolute semantics either way (Glanzberg, 2015, 90).

Given this absoluteness of natural language semantics, the falsity of the LNLT follows by direct comparison of the meaning-theoretic assumptions and the corresponding conceptions of logical notions. How *could* natural language endow any logical expressions – intuitively: and, or, not, etc. – with a logical consequence relation, much less identify them as *logical* to begin with, if the semantics are absolute? For in order to be identified as logical constants, the candidate expressions must pass the invariance test and generate entailments[5] that preserve truth over a *sufficiently large* range of models (cf. Beall et al., 2019, sect. 1). This range of models, however, is precisely absent in

[4]The question arises whether formal semantics really ought to be in the business of supplying a theory of understanding (cf. Yalcin (2017) and Balcerak-Jackson (2021) for alternative views). It seems that as far as the research practice of formal semantics is concerned, the modelling of truth-conditions via compositionality is all that matters. In general, however, I would submit that given the conceptual ties between linguistic meaning and linguistic understanding, a *total* divorce of accounting for the latter from accounting for the former would be a strange occurrence indeed. Hence, to the degree that formal semantics is to be seen as supplying a *theory of meaning*, linguistic competence cannot be ignored (cf. Dummett (1996) for excellent points in this regard).

[5]An entailment from P to Q obtains whenever the truth-conditions of P are included in those of Q, meaning that whenever P is true, Q is as well (Glanzberg, 2015, 80).

absolute semantics. Hence, natural language does not provide a (semantic) criterion of logicality (Glanzberg, 2015, 102) and any entailments candidate expressions would give rise to could only ever be *lexical*, i.e. on par with other analytic entailments such as from x is a bachelor to x is unmarried (Glanzberg, 2015, 98f.). Thus, no relation of logical consequence *could* be generated in natural language on the basis of logical constants. This is Glanzberg's *argument from logical constants* (henceforth: ALC).[6]

Thus, given the ascription conditions for logical constancy and consequence, these notions turned out to be 'unrealisable' in natural language, given the meaning-theoretic assumptions. Let us dub this, accordingly, a *conceptual mismatch*.

2.3 Different Readings of Glanzberg

At this point, I wish to remark on my reading of Glanzberg. I call it the *conceptual reading*, since it reads Glanzberg as being committed to a modal claim, based on conceptual considerations. To wit: We *cannot* identify logical constants as logical or model a logical consequence relation in natural language with them, *while* still modelling linguistic competence. The reason for this is the mismatch on the level of ascription conditions between logical notions and that of linguistic meaning. This may seem at odds with what Glanzberg overtly argues against, namely the LNLT, which is no modal claim. Thus, one could read Glanzberg in a different light, i.e. as claiming that modelling logical consequence and modelling linguistic meaning are simply different theoretical enterprises, as revealed by the current practices in these fields, making no further conceptual or modal claims. Call this the *methodological reading*.

However, I would submit that the methodological reading is, for various reasons, both uninteresting and uncharitable, hence I take Glanzberg to be committed to the conceptual one. First, if it is merely a divergence in current practices, we would only need to "adjust our toolkit accordingly" (Sagi, 2020, 187). That is, since all we need to obtain a relation of logical consequence in natural language – as generated by appropriate logical constants – is a sufficiently large range of models, we could simply add this range. Refusing to do so, simply because extant semantics do not have such a range to begin with, would be to beg the question (cf. ibid.). Second, however, this line of reasoning ignores Glanzberg's own claim about the reach of his argument(s), when he writes

[6] Readers of Glanzberg might object to this presentation for leaving out his distinction between restrictive and permissive views on logic (cf. Glanzberg, 2015, 78f.). However, the argument presented here makes use of absolute semantics, hence it is equally damaging to both approaches (cf. Glanzberg, 2015, 103, 106). Additionally, countering an argument against restrictive views should automatically vindicate permissive views as well. As such, we can safely ignore this complication. Also, it must be noted that Glanzberg offers two further arguments – the argument from absolute semantics and from lexical entailments (Glanzberg, 2015, 103). However, my reply only concerns the ALC, hence I shall ignore the others in what follows. That said, my presentation of the ALC implicitly includes a presentation of the argument from absolute semantics, since the former relies on the latter (cf. Glanzberg, 2015, 102; Sagi, 2020, 194).

> [The] semantics of natural language – the study of speakers' semantic competence – *cannot* look at [a non-trivial range of models], and still capture what speakers understand. To capture what the speakers understand, semantics must be absolute, and so blind to what happens to a sentence across any non-trivial range of models.
> (Glanzberg, 2015, 91, emphasis added).

In other words, given that we are restricted by 'what speakers know when they know what their sentences mean', the divergence is not a mere accident, but a fundamental incompatibility, just as the conceptual reading assumes. Third, even if the above quote could be interpreted in order to fit the methodological reading, a further issue looms. It seems to me that the *only* phenomenon warranting the inclusion of enough models into semantic theorising to model logical notions would be *logic itself*. For no other phenomenon requires as many models, given the model-theoretic conceptions. Thus, it appears that only considering extant practices in semantics is precisely the wrong place to look for an answer concerning the LNLT. For we could only justify the inclusion of enough models by *assuming* the truth of the LNLT.[7]

Lastly, and building upon the third point, it is prudent to consider conceptual mismatches before delving into extensional considerations. For if there would indeed be an incompatibility between logical notions and linguistic meaning, then extensional investigations would be rendered superfluous. Hence, instead of being bogged down by such debates, I would, in any case, recommend considering the conceptual level *first*. Thus, I shall assume the conceptual reading in what follows.

3 Inferentialism and Proof-Theoretic Conceptions

In the remaining sections of this paper, I will argue against the conclusion of the ALC. My point of departure is to reject his meaning-theoretic background assumptions and assume a use-theoretic approach. The latter takes meaning to be fully determined – if not constituted – by use. This will motivate an abandonment of model-theoretic conceptions of logical notions in favour of proof-theoretic ones. As I will show in section 4, these use-theoretic assumptions about meaning and the corresponding proof-theoretic conceptions of logical notions do *not* generate a conceptual mismatch. In fact, we have reasons to be optimistic about the existence of logical constants and consequence relations in natural language.

[7]For this reason, I consider Sagi's other primary counter against Glanzberg – the rejection of the absoluteness of semantics (cf. Sagi, 2020, 185f.) – insufficient. It is true that some range of models might be needed account for the common fact that speakers do not know the precise extensions of all their terms. However, it remains dubious that such ignorance suffices to reintroduce *enough* models for the purposes of modelling logical notions. Since these two counters – rejecting absoluteness and the charge of question-begging – are Sagi's grounds for arguing against Glanzberg's remaining points (cf. 2020, 190, 194), I will not discuss her rejoinder to Glanzberg further.

3.1 A Quick Primer on Use Theories

Since it is not the purpose of the present paper, I shall not defend this switch of theories of meaning in detail. Thus, my conclusion will be conditional: *If* we switch from referentialist truth-conditional semantics to a use-theoretic one, we can avert the mismatch. That said, if my reasoning below is sound, a use-theoretic approach can vindicate the epistemological and methodological points raised in the introduction, whereas mainstream referential truth-conditional semantics cannot. This itself would be an argument in favour for the use-theoretic approach.

Naturally, this is not to say that the literature is devoid of other arguments in favour of use-thereotic approaches. The arguments, however, should be familiar to most philosophers of language. First, referring or denoting something is not the hallmark of meaningfulness, as the myriad of expressives, greetings, commands, etc. demonstrate (cf. Wittgenstein, 1953, §27; Glock, 2018, 65; Kaplan, 1997). The same goes for truth and the diverse range of speech acts beyond assertion. Second, even in the case of those expressions that have a referent, such as proper names, their referent cannot be identified with their meaning, on pain of a category mistake in Ryle's sense (cf. Ryle, 1949, ch. 1). For I can love and marry the bearer of a name, but not its meaning (cf. Wittgenstein, 1953, §40; Glock, 2018, 65f.). Third, meaning is essentially something that can be explained, learned and taught. Yet the way we do this, especially when we do not have a prior language to fall back on, is by mastering the *correct use* of words. Thus, at least knowledge of meaning is a matter of knowing how to correctly use an expression (cf. Skorupski, 2017, 74f.; Glock, 2018, 66).

While this last point primarily deals with metasemantics – specifically what constutites knowledge of meaning – it nevertheless provides further motivation for adopting a use theory more generally. What all meaningful expressions seem to have in common is a *set of rules* that determine their correct use. Combined with this third point, what speakers must master thus are these rules, i.e. (tacit) knowledge-how of following them. However, if knowledge of meaning is knowledge-how of following rules, then it is unclear what motivation remains to deny that meaning *itself* is constituted by such rules (more on this in 3.2 below). Furthermore, from these considerations, especially the essential learnability of meaning, it can be further argued that the rules in question must be entirely determined by the practices of the linguistic community. For it is unclear how else such rules could come into the practice and be mastered by its practitioners.

None of this ought to persuade the staunch advocate of referentialist truth-conditional semantics. These arguments warrant their own, separate discussion, and more in their defense would need to be said.[8] Nevertheless, I hope to have shown that use-theoretic approaches have merit, and to have exemplified their general point of departure. We shall now move on to the details of the approach.

[8] For example, one would need to reject semantic externalism, according to which use does not fully determine meaning (e.g. Putnam, 1975). For an excellent defense against externalism, cf. Hanfling (2000, ch. 12).

3.2 Inferentialism in a Nutshell

So far, 'use' has been left unspecified, yet arguably not all aspects of use are equally important ingredients in meaning. In addition to this, recall that given the overall goal of this paper, the interest lies in a theory of meaning for public, natural languages, specifically those parts that could contain logic. Thus, we will focus on the *declarative* part of a language. In our case, then, the central notion for the elucidation of 'use' will be that of an *inference* in a wide sense, as not only incorporating inferences between sentences, but also from situation to sentences or from sentences to actions, for example (cf. Peregrin, 2009, 160).[9] The general idea behind *inferentialism* is that at least the meaning-*determining* aspect of an expression's use in declarative language is its inferential use (Murzi and Steinberger, 2017, 197ff.). As such, the aforementioned rules of use become rules of *inference*, and linguistic competence is achieved by possessing the ability to correctly use an expression *in inference*.

Let us unpack this cluster of theses in more detail. We have already mentioned the distinction between semantics and metasemantics. The latter is concerned with providing an account of what determines or explains the meaning of an expression (metaphysical), and what constitutes knowledge of said meaning (epistemological). As such, the following metaphysical and epistemological theses can be distinguished (cf. Murzi and Steinberger, 2017, 199):

(MT) An expression's inferential use, i.e. its embedding in the overall inferential practice of the linguistic community, determines its meaning.

(ET) To understand an expression's meaning is to have knowledge(-how) of its use in inference.

I have already motivated both (MT) and (ET) in section 3.1. The argument from the learnability of meaning straightforwardly supports (ET), since for the *declarative* part of language, the mastery in question is arguably the making of correct assertions. More specifically, it is a matter of knowing how to support them and making further inferences from them (cf. below). Thus, we shall assume both from here on out.

Thus, we have the following picture so far. There is the so-called *ordinary inferential practice*, i.e. the public practice of justifying assertions and inferring from them. This practice gives rise to certain rules of use for *sentences* first and foremost, and subsentential expressions second (cf. below), in the form of rules of inference (cf. Peregrin, 2009, sect. 6). These may then further determine any (more) traditional semantic values (cf. below) that one might wish to include in one's account of meaning (cf. fig. 1). The question that is raised at this point is what the meaning of a declarative sentence is to consist in, given these metasemantical background assumptions. As

[9] As such, inferentialism is really a mislabel. What distinguishes it from other theories of meaning is the central explanatory role of rules, *paradigmatically* rules of inference. That said, I will follow the literature and keep referring to it as 'inferentialism'. The reader is implored to keep the misleading nature of the label in mind.

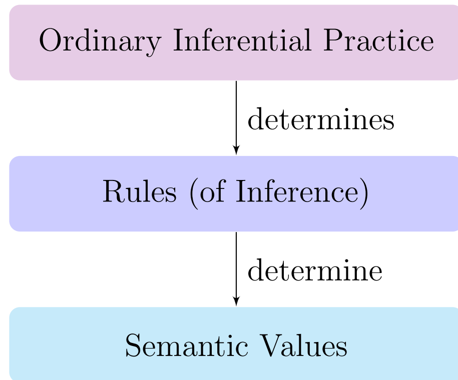

Figure 1: The different levels of meaning determination

Strawson has famously argued, it is a category mistake to equate sentence-meaning with truth-conditions, for neither sentence-types nor tokens are correctly called true or false (cf. also Glock, 2003, 119f.). It is what a sentence says on a particular occasion, i.e. the *proposition* expressed by an assertoric use of a sentence, that can be called true or false (Strawson, 1950, 325ff.). Of course, what a sentence says on a particular occasion is, among other things, such as context, a function of its meaning (cf. Glock, 2003, 154). As such, sentence-meaning can be equated with "[...] *general directions* for [the sentence's] use in making true or false assertions" (Strawson, 1950, 327). Thus, in line with both these observations and the learnability argument, we shall adopt the following *semantical thesis*:

(ST) The meaning of an expression is the set of inference rules that govern its correct use, as sanctioned by the inferential practice at large.

Semantic values are then taken to be those entities that are determined by such rules, be they truth- or assertibility-conditions for the propositions expressed on particular utterances, or referents of referring expressions, etc.

Concerning the semantic values for sentences, we now face two options. On the one hand, we can understand the rules that constitute a sentence's meaning as giving us general directions for making *true* or *false* assertions, as Strawson recommends. Thus, the resulting semantic values will be truth-conditions. Call this this *orthodox route* (cf. Murzi and Steinberger, 2017, 200). On the other hand, we can follow most inferentialists in eschewing truth as an explanatory notion, and instead opt for different ones. Call this the *unorthodox route* (Murzi and Steinberger, 2017, 200).[10] For example, as Brandom argues, two aspects are relevant to a sentence's meaning: the conditions under which an

[10]To this option, I shall also count the 'deflationary' addition of truth-conditions and referents to semantic values in the way of Brandom, who argues for a prosentential theory of truth and a pro-anaphoric theory of reference (Brandom, 1994, ch. 5). Truth-conditions and referents would then only play an 'expressive' role, however, not a genuinely explanatory one (Brandom, 1994, 283).

assertion of it would be *warranted* – sometimes also just called *assertibility conditions* – and what appropriate inferences can be drawn from it. This is due to the fact that in order to be credited with full understanding of a sentence, both aspects are crucial (Brandom, 2000, 64ff.). Call this the *two-aspect model of meaning*. However, this model can equally well be adopted by someone who assigns explanatory relevance to the notion of truth. After all, the notion of truth seems to be conceptually tied to the notion of assertion (cf. Miller, 2002, sect. VI; MacFarlane, 2009, sect. 3). Thus, there is reason to argue that the rules referred to in (ST) must supply truth-conditions by being general guides to making true or false assertions. Against this, it has been argued that *verification-transcendent* truth-conditions cannot be accommodated given use-theoretic metasemantic assumptions,[11,12] and that concerning truth's conceptual connection to asserting, putting forward a proposition as true during an assertion amounts to nothing more than licensing specifically inferential work (Brandom, 1983, 638f.).

While I have strong sympathies for both, in what follows I will primarily consider the unorthodox route, leaving a detailed discussion of the orthodox approach to section 4.3. The chief reason is that it is not so clear how exactly one ought to work out the details for the orthodox variant (cf. ibid.). However, I will argue that my reasoning applies to the orthodox approach as well, provided one can embed such semantics into the metasemantics given by (MT) and (ET). In general, then, we take metasemantics to be concerned with the top-layer of figure 1, whereas semantics is (primarily) concerned with the middle one. Lastly, semantic values such as referents, truth- or assertibility-conditions are treated as *resulting* from meaning, i.e. the middle layer of rules.

From the preceding paragraphs, a few inferentialist trends should have become apparent. First, the inferentialist reverses the usual order of explanation with respect to standard formal semantics. A sentence's use in inference determines corresponding rules, which in turn constitute sentence-meaning, where the propriety of the relevant inferences are taken to be antecedently given by the ordinary inferential practice (cf. Brandom, 1983, 640). As such, the propriety of inference rules explain why a sentence has the meaning it has, not the other way around. Second, the primary semantic unit for the inferentialist is sentence-meaning, with the meaning of subsentential expressions being accounted for in terms of their contribution to the overall sentence-meaning. In this way, compositionality is preserved (cf. Murzi and Steinberger, 2017, 199ff.; Peregrin, 2009).[13]

Two further distinctions in the landscape of inferentialism need comment. First, the

---

[11] Cf. Miller (2002) for excellent discussion of this so-called *manifestation argument*.

[12] It is important to remark that for everyday empirical propositions, their assertibility-conditions at least include, if not equal, their usual truth-conditions. For the warrant of an assertion of Student x has finished their assignment is clearly given if I can show the finished work, i.e. if I can verify the statement. The unorthodox route is simply more cautious about the role of truth-conditions beyond such examples, preferring to use assertibility-conditions as a more general device.

[13] Moreover, given the repeated appeals to the notion of *correct use*, inferentialism is a generally *normative* theory of meaning. However, it must be pointed out that none of the arguments below turn on this, which is why I will ignore the topic in what follows. Cf. Glock (2018) for a defense of the normativity of meaning.

foregoing discussion can be understood as pertaining to *global inferentialism*, i.e. the position that the inferentialist account covers the whole of (declarative) language, rather than some of its parts.[14] Alternatively, the preceding remarks can be understood as pertaining to so-called *local inferentialism*, which is an inferentialist account of a specific area of discourse, such as the moral domain or specialist languages such as formal logic (cf. Murzi and Steinberger, 2017, 201).

Second, *global* inferentialism as it is championed by Brandom, for example, is a *holistic* theory: The meaning of an expression is determined by the entire inferential network it is embedded in (cf. Murzi and Steinberger, 2017, 201). However, this seems to go too far. To use an example due to Murzi and Steinberger (2017, 202), it would seem odd to think that the notion of measurable cardinal would be relevant in determining the concept of a cat. Furthermore, since the inferentialist mode of explanation via inferential links will always involve more than one concept, something like *semantic atomism* – the idea that every concept is fully determined in isolation of others – is ruled out by default. Hence, we shall adopt *molecularism* for our background 'global inferentialism', which holds that certain concepts are determined *en bloc*, but do not require further concepts beyond this 'cluster' (cf. Murzi and Steinberger, 2017, 201).[15] One such cluster will be that of formal logic, to which we turn next.

3.3 Logical Inferentialism

As is to be expected, *logical inferentialism* contends, at its core, that the meaning of logical constants is determined by their use in *deduction* (Murzi and Steinberger, 2017, 204). Given the two-aspect model of meaning, we need grounds for warranted assertions and appropriate inferences from them for those compound sentences that feature the connective in principal position.[16] This finds a natural expression in Gentzen-style systems.[17] Thus, for example, the meaning of conjunction is given by its rules of inference as they appear in those systems:

$$\frac{A \quad B}{A \wedge B}\wedge\mathrm{I} \qquad\qquad \frac{A \wedge B}{A}\wedge\mathrm{E} \qquad\qquad \frac{A \wedge B}{B}\wedge\mathrm{E}$$

[14]The critical literature on (global) inferentialism is far too large to be detailed. Famous criticism stems from Williamson (2007, ch. 3 and 4), and interesting one from MacFarlane (2009). For rebuttals to the former, especially from the use-theoretic camp and including defenses of conceptual truths, cf. Schroeder (2009), Murzi and Steinberger (2017, sect. 4) and Büttner (2021).

[15]I have deliberately coined this discussion in purely explanatory terms, and not in epistemological ones, as it has been in the cited paper. The reason is that (i) the reasoning applies in the explanatory domain equally well and (ii) it is this domain we will be primarily concerned with in the sequel.

[16]A such, the Gentzian idea of introduction rules as *defining* the connective and the corresponding elimination rules being simply consequences of such definitions (cf. Gentzen, 1934, 80) will be rejected in the sequel.

[17]Of course, other kinds of proof systems are available as well, such as sequent-based ones. As Steinberger (2009, 202f.) argues, however, using a sequent calculus would presuppose an understanding of disjunction. For this reason, I will stick to natural deduction systems, but nothing in my argument hinges on this.

Given (ST) above, we shall see such rules as *constitutive* of a connective's meaning. In addition to this, these rules are determined by the overall inferential practice in which it figures, as per (MT). Furthermore, understanding a connective is to be seen as the capacity to use it in accordance with such rules. As such, all three of our meaning-theoretic assumptions find an expression in logical inferentialism.

Two reasons are readily available as to why this account is to be preferred over its competitors, especially truth-table based ones. First, as Boghossian argues (2011, 493), ascribing competence with respect to logically relevant expressions such as and is based on the speaker according with certain rules of use rather than others, much in line with the learnability argument above (cf. also Steinberger and Murzi 2017: 205). Second, as Rumfitt (2000, sect. II) argues, the truth-table based approach cannot distinguish between truth-functionally equivalent expressions – such as $A \vee B$ and $\neg(\neg A \wedge \neg B)$ – despite their differences in inference. However, whether an inference proceeds by one step – as in A, therefore $A \vee B$ – or by many – A, therefore $\neg(\neg A \wedge \neg B)$ is arguably of interest to the logician (Rumfitt, 2000, 785). As such, the inferential approach simply enjoys greater expressive flexibility for the purposes of the logician.

But which rules are supposed to be constitutive of meaning and which are not? A natural thought in this regard is to focus on only those rules which are free of other connectives and, in a certain sense, *basic* or 'gap-free'. To get a grip on this notion, consider (trivial) analytic entailments in natural language. Typically, their justification can only proceed with reference to use. For example, if someone questions my inference from x is right of y to y is left of x, there is nothing I can say apart from that this is what these expressions mean, that this is how they are used. Thus, those rules of inference that cannot be justified with reference to further rules must be basic or 'gap-free',[18] i.e. their conclusion is not mediated by any further rule.[19] Thus, we shall require meaning-constitutive rules to be gap-free in what follows, as individuated on the basis of whether further justification can be provided for them.

So far then, the meaning of a logical constant is to be given by their gap-free introduction and elimination rules, as they can be found in natural deduction systems. However, this is not enough. For example, the following rule is arguably satisfying all

[18] The modality appealed to here turns not on the (potential) *practical need* for further justification, but the *justificatory possibility* of breaking down rules into intermediate ones.

[19] This might give the impression that there are no 'gappy inference rules', but they are easy enough to come by in formal settings (cf. Rumfitt's example above). For an analytic example from natural language, consider the following:

$$\frac{x \text{ is red}}{x \text{ is spatially extended}}$$

This inference rule is *not* gap-free, for it *can* be justified with reference to further inference steps – namely that red can only apply to physical objects of a certain minimum size, which *in turn* implies having a spatial extension. Thus, there are indeed gappy inference rules to be found (analogously to what has been called 'synoptic grammatical propositions', cf. Baker and Hacker (2009, 20)). I am indebted to Gil Sagi for giving me the impetus to clarify this point.

these criteria as well:

$$\frac{x \text{ is a human male} \qquad x \text{ is of marriageable age} \qquad x \text{ is unmarried}}{x \text{ is a bachelor}}\ \textsf{Bachelor-I}$$

However, bachelor is not a logical constant. Thus, in order to get to an account of logical constants proper, the general idea behind logical inferentialism needs to be extended to rule out other analytically valid inferences. Traditionally, logical constants are supposed to be those expressions which are 'topic-neutral', i.e. logical reasoning should be applicable in all domains of discourse, hence the meaning of constants cannot contain any 'substantive' content (cf. MacFarlane, 2017, sect. 4). For the inferentialist, such topic-neutrality is naturally accounted for by requiring the corresponding rules of inference to be *purely inferential*. For our purposes, given the focus on sentential connectives, we adopt the following strict criterion.[20] A rule of inference is purely inferential if, and only if, it contains (i) at least one occurrence of the purported logical constant[21] in either the premisses or the conclusion, and (ii) contains only structural (e.g. brackets or commas) or schematic (e.g. A, B, etc.) signs. These schematic signs ought to be placeholders for declarative sentences, hence we require them to substitutable for such *salva congruitate*. Thus, this criterion is clearly not satisfied by Bachelor-I, since no declarative sentence can be substituted for x *salva congruitate*. However, we are also ruling out some sentential operators, such as propositional attitude constructions. For x believes that p allows the (defeasible) inference to x behaves as if p, which is clearly not purely inferential. Thus, the criterion *prima facie* neither over- nor undergenerates, and delineates the topic-neutral sentential operators.

In order to refer to them in the following discussion, and by way of summary, our three criteria for logical constants are repeated below[22]

(LC1) The expression's meaning-constituting rules of inference can be given in the style of a natural deduction system's introduction and elimination rules.

(LC2) These rules are gap-free.

(LC3) These rules are purely inferential.

(LC1)'s first part follows from our general inferentialist outlook on meaning, viz. (ST). The crucial part of its formulation, however, is the can. For in order to avoid the trivial

[20] As such, this choice cannot be held against us, for it makes the following task harder, rather than easier.

[21] 'At least one', since otherwise we would rule out double negation elimination as purely inferential, which would be wrong.

[22] The most famous challenge to logical inferentialism has been posed by Prior (1960). A common strategy against this objection is the notion of harmony (cf. Steinberger (2011) for an overview), which is sometimes considered as a necessary condition for logical constancy. I will remain agnostic on the issue for several reasons. First, harmony is not so easily motivated (cf. Rumfitt, 2017). Second, tonk is clearly not a pre-existing expression in natural language, so for our purposes, it is not an issue either way. Third, the concrete example to be discussed in section 4.2 will be one whose rules of inference are harmonious. Hence, nothing hinges on this agnosticism for the argument in this paper.

syntactic mismatches mentioned in the introduction, we must allow for some kind of translation from informal notation to a more formal one.

Before we can proceed to logical consequence proper, a methodological remark is in order. As the attentive reader may have noticed, the motivation for specifically *logical* inferentialism and notions such as being gap-free all stem from direct considerations about ordinary inferential practice and ascription conditions for linguistic competence. As such, it might be argued, we already incur some commitment with regards to the LNLT, by virtue of infusing the proof-theoretic conceptions with conditions more readily realisable in natural language. Against this, I offer the following counter. Logical inferentialism and the resulting proof-theoretic conception of logical constants should, first and foremost, be seen as an account of meaning in *formal* logic. For there is no reason to suppose that the general criteria for linguistic competence do not extend to specialist languages such as formal logic.[23] Thus, the above account is to be seen as an account of logical expressions in natural language only after we have vindicated the LNLT.

3.4 Proof-Theoretic Logical Consequence

Having set up the meaning-theoretic background and the resulting proof-theoretic conception of logical constants, the last missing piece is logical consequence. Once more, its motivation stems from a critique of the prevalent model-theoretic conception of logical consequence relations.

First, as Shoesmith and Smiley (1978, 105) demonstrate, it is possible for arguments to contain logical fallacies, yet to end up having a conclusion which, model-theoretically speaking, follows from the premisses. Additionally, this issue cannot be rectified by requiring each argumentative step to be valid, for this alone would not rule out the possibility of circular arguments. However, in either case it would be absurd to say that the conclusion *follows* from the premisses, and such arguments would be rightfully rejected as incorrect. Thus, both the validity of individual reasoning steps and their proper order are of importance to correct logical reasoning – arguably a topic that should be guided by an appropriate notion of logical consequence.

Second, knowing that a set of premisses is true and that a certain conclusion A follows logically from them is sufficient to justify the claim to knowledge of A. However, under the model-theoretic picture, knowledge of the validity of an inference would *include* knowing the designated value of the conclusion. As such, it seems that the model-theoretic conception of logical consequence is epistemically inert (cf. Etchemendy, 2008, 267).[24]

Thus, we need a conception of logical consequence that pays heed to the individual steps of an argument, their respective order, and which does not render deduction an

[23]For an impressive account of so-called *formal analyticity* in formal languages, which even circumvents Williamson's argument from expert revisability, cf. Soysal (2019, esp. sect. 2).

[24]Furthermore, as Prawitz (2005, 676) notes, the necessity constraint might not be satisfied either.

epistemically inert enterprise. The *proof-theoretic* conception of logical consequence meets these *desiderata*. At its core, it is committed to the following claim:

(LCons) A is a logical consequence of Γ iff there is a *logically valid argument* (or *proof*) from premisses in Γ to the conclusion A.
(cf. Beall et al., 2019, sect. 3.2)

where, for our purposes, an argument is taken to be the kind of proof tree found in natural deduction systems. As the reader will have noticed, this definition contains logical on both sides, hence we need an account of logical validity for arguments that does not rely on a prior conception of logical consequence. Dag Prawitz (2005, 2006) provides an impressive and systematic account of logical validity in proof-theoretic terms.

As Prawitz points out, justifications have to come to an end at some point, which in case of inferences take the following form:

> Similarly, what can we answer someone who questions the drawing of the conclusion $A \to B$, given a proof of B from A, except that this is how $A \to B$ is used, it is part of what $A \to B$ means?
> (Prawitz, 2005, 682).

In other words, certain kinds of inferences ought to be valid simply in virtue of what the relevant expressions mean. In the inferentialist's case, this meaning *is*, of course, constituted by those same rules. As such, an argument is supposed to be valid if such a rule is applied in its last step, and the preceding subarguments are valid, too. This yields a recursive definition of validity, where the ultimate grounds of justification are meaning-constituting inference rules (cf. Prawitz, 2006, 511-516). If the validity of such an argument is furthermore irrespective of the content of the atomic propositions in the argument, the argument itself is *logically* valid (cf. Prawitz, 2006, esp. 515f.).

This notion of logical validity is custom-fit for our characterisation of logical constants. For the purely inferential nature of their inference rules guarantee logical validity in Prawitz' sense. Thus, we can give the following definition of it:

(Val) An argument from premisses Γ to the conclusion A is logically valid just in case it is a series of applications of either introduction or elimination rules of logical constants.

Thus, logical consequence obtains whenever there is a proof from the premisses to the conclusion, while a proof is a sequence of logically valid inference steps, as supplied by the meaning of the logical constants. Since nothing about the characterisation of logical constants depended on any notion of validity, much less logical consequence, the potential circularity in (LCons) has been avoided. Lastly, this characterisation of logical consequence is not epistemically inert. For knowledge of validity is grounded in linguistic competence. Specifically, it stems from the appreciation of certain analytic inferences, whose validity is sanctioned via the rule-governed linguistic practice of the community – whether natural language speakers or specialists creating formal languages. Therefore, no knowledge of the truth-value of the conclusion is presupposed, even though in virtue of the analyticity involved, such inference rules could not fail to be truth-preserving.[25]

4 Averting the Mismatch

Having all our pieces in place, I will now argue that the inferentialist does not face a conceptual mismatch between their meaning-theoretic assumptions and the resulting conceptions of logical notions. The chief reason is that the meaning-theoretic assumptions (MT), (ET) and (ST) fit in well with the ascription conditions for logical constancy – (LC1-3) – and consequence – (LCons) and (Val), in the sense that the former not only do *not* rule out the latter, but provide a congenial embedding for them. Additionally, as we can glean from the expression and, there are at least some *contexts* in which natural language *does* appear to contain a relation of logical consequence.

4.1 Conceptual Congruence for Inferentialists

Coming from our discussion so far, we have seen that the proof-theoretic conception of logical consequence requires a prior demarcation of logical expressions. However, once we do have logical constants, we immediately yield a relation of logical consequence. For according to (Val), every single-step inference according with an I- or E-rule of a logical constant is already an instance of such a relation. Furthermore, there is no reason to suppose that chaining such inferences would be any general issue for speakers of a natural language. Thus, we yield the following principle:

> (I) A conceptual match between proof-theoretic conceptions of logical constants and inferentialist assumptions about meaning automatically yields a conceptual match for consequence relations, too.

Thus, we only need to check for the compatibility between our meaning-theoretic assumptions and the proof-theoretic conceptions of logical *constants*. Concerning this endeavour, another general observation is important. It holds in full generality that whenever some entity falls under some concept F, it follows that the ascription conditions for F are satisfiable. This gives us another principle:

[25] At least in the case of everyday empirical sentences (cf. fn. 12). More generally, recall that under the Brandomian picture, truth enters in a deflationary way: as a tool to characterise propositional content as determined by antecedently valid inferences (cf. fn. 10). Hence, one could potentially argue for general truth-preservation in the unorthodox setting beyond verifiable sentences in this way.

(II) If, for any of the conditions for logicality, we can find rules in natural language satisfying them, then the relevant condition is compatible with the meaning-theoretic assumptions, i.e. satisfiable in natural language.

This principle yields a handy 'companion-in-guilt' strategy for assessing the conceptual compatibility between the meaning-theoretic assumptions and our proof-theoretic conception.

With these tools in hand, consider (LC1) first. It is only concerned with the possibility of presenting certain rules of inference in a specific form, hence operates only on the level of semantics, not metasemantics. Moreover, the idea of I- and E-rules was already observed to flow from the general two-aspect model of meaning. As such, it comes as no surprise that there is an abundance of natural language examples for which such presentation can be found as well, such as Bachelor-I. In general, any analytically definable concept F will have 'introduction' and 'elimination' rules of the form

$$\frac{p_1(x) \qquad p_2(x) \qquad \dots \qquad p_n(x)}{F(x)}\ F\text{-I} \qquad\qquad \frac{F(x)}{p_i(x)}\ F\text{-E}_i$$

where p_i is some salient statement about a given entity x. Thus, by (II), we conclude that there is no mismatch with respect to (LC1).

Next, consider (LC2). It imposes a specific requirement on rules themselves, specifically of epistemological nature. However, there is nothing about (ET) that would render the idea of being gap-free problematic, nor is there anything about the idea of meaning-determination or -constitution that stands in tension with this property. Quite the contrary: In order not to end up with infinite regresses, it would seem that inferentialism *requires* such rules. More importantly still, analytically definable predicates are once more companions in guilt. For Bachelor-I is part of part of the meaning of bachelor, and it cannot be further justified. Thus, by virtue of (II), (LC2) faces no mismatch either.

Consider last (LC3). Arguably, we cannot find companions in guilt in this case, since any such candidates could only be logical themselves. Thus, we need to motivate the conceptual congruence by other means. On the one hand, observe that nothing about (MT), (ET) or (ST) rules out the existence of purely inferential rules *in principle*. There is nothing inherently incoherent about the thought that ordinary inferential practice could give rise to such rules, that they could constitute meaning, or that their nature as purely inferential would form insurmountable epistemic burdens on speakers. For concerning the latter point, the fact that natural language speakers can also learn and become proficient in formal logic ought to forestall such worries.

On the other hand, concerning ordinary inferential practice, such rules would find a natural embedding. For part of this practice is to object to others' arguments by largely two means: either criticism of the warrant for their premisses, or criticism of the validity of their inferential patterns. It becomes important in such cases that (i) all premisses are properly distinguished and identified and that (ii) all inferential links are made explicit. Here, as Brandom has argued (e.g. Brandom, 2000, ch. 1, sect. VII), logical vocabulary would serve a crucial and basic function. For example, conditional constructions identify premisses and conclusions, while conjunction serves the purpose of grouping premisses. However, logical vocabulary is not only congenial for this 'expressive' task, but would also serve a crucial epistemic function. For in order for this practice to be stable, we need a 'rock-bottom': certain inferential links, such as the 'meta-links' of the expressive vocabulary, cannot themselves be assailable. Otherwise, no justification in this practice could ever come to an end. Therefore, purely inferential rules, given their universal applicability, would serve such an expressive *and* 'foundational' role.

In summary, (LC3) is in good standing. None of our meaning-theoretic assumptions ruled it out in principle, and moreover, such rules would serve crucial functions in ordinary inferential practice. Together with our discussion of the other two criteria, we are now in a position to conclude that *overall*, the proof-theoretic conception of logical constants does not result in any mismatch with inferentialist background assumptions about meaning. Thus, *neither* does the proof-theoretic conception of logical consequence, thanks to (I). As such, we can now reject Glanzberg's argument from logical constants. In contrast to his set-up, our ascription conditions for logical constants require only rules of inference, plus further requirements *on* those rules. Thus, none of this goes beyond what can be found according to our meaning-theoretic assumptions, hence there *could*, in principle, be a relation of logical consequence in natural language. Therefore, the LNLT has been vindicated on a conceptual level.

4.2 Extensional Considerations

With conceptual compatibility secured, I want to further argue in favour of the LNLT. If we had reason to believe that natural language contained logical constants, our worries about conceptual mismatches would, according to (II), be alleviated. That said, the result of this section will be less conclusive, giving us merely reason to be *optimistic* about the existence of logical constants in natural language.

Let us turn to and and postulate the following meaning-constitutive rules as a first, naive attempt:

$$\frac{\text{A} \quad \text{B}}{\text{A and B}}\ \textsf{and-I} \qquad\qquad \frac{\text{A and B}}{\text{A}}\ \textsf{and-E} \qquad\qquad \frac{\text{A and B}}{\text{B}}\ \textsf{and-E}$$

Clearly these rules come in the appropriate presentation in the sense that they satisfy (LC1) and (LC3). Furthermore, these rules appear to be gap-free. For if someone were to challenge my reasoning from A and B to A, I could not provide further justification, except by using an answer like "Well, when I have A and B, I therefore also got A, right?".

Unfortunately, there is a host of issues concerning this proposal. The ordinary inferential behaviour of and is not as simple as the above rules suggest, and so we find many counter-examples to the proposal. Consider the following:

(1) A and B came to the party. (*distributive coordination*)
(2) Sarah[1] is clever and she$_1$ is brave. (*anaphora*)
(3) She took the pill and she recovered. (*non-commutativity*)
(4) A and B are siblings [of each other]. (*reciprocity*)
(5) A, B, C and D surrounded the fort. (*plural quantification*)
(6) She remembers John[1] and that he$_1$ promised to join. (*coordination of 'unlikes'*)

(1) is witness to the simple fact that and allows for syntactic coordination of expressions of the same category (cf. Partee and Rooth, 2002), while obeying a version of and-E and and-I, without being of the prerequisite syntactic form. (2) does not allow a reversal of the conjuncts, because of the anaphoric use of the pronoun (cf. Dekker, 2011, 925). Sentence (3) has a similar issue: the order of conjuncts results in a corresponding temporal reading, hence they do not commute (cf. Strawson, 1952, 80), in contrast to our rules for and above, which do yield commutativity for and-sentences. (4) raises the issue that while a version of and-E might hold, A/B is a sibling does not allow the inference to (4), as (4) claims that they are siblings *of each other*.[26] (5) shows that not all uses of and allow for inferences according to and-E, since no single person can surround a fort alone, giving rise to the notion of plural quantification (cf. Linnebo, 2022). Lastly, and can even coordinate expressions from different – 'unlike' – syntactic categories (cf. Bayer, 1996).

Instead of dealing with all these examples in detail, I wish to make one crucial observation. It is important to notice that even if it turns out that these rules are invalid in certain contexts,[27] this does not prevent them from performing their expressive or foundational function. For there are at least two notions of universal applicability at play here: applicability across topics and across contexts. Thus, while and-E and and-I may not be licensed in all contexts, *when* they are, they do so across all topics, just as their expressive and foundational role requires. In these contexts, then, sentences obtained by applications of and-I or and-E will follow logically from their given premisses, and we yield a logical consequence relation in natural language by (I).

Thus, there are reasons to be optimistic about finding linguistic expressions functioning as logical constants in natural language, at least for certain types of contexts. They do not require any *substantial extra*-linguistic machinery to be deployed in order to be recognised *as* logical, or to generate a relation of logical consequence. For their logical character is apparent in important contexts of use, such as when assessing each

[26] I am indebted to Dolf Rami for the example.

[27] Although, I think at least some examples can be dealt with rather easily by a combination of two things. First, we are free to modify our proposed rules for and as long as we satisfy (LC1-3) – the LNLT only relies on there being *some* logic in natural language. Second, observe that logical expressivism entails that the existence of materially correct inference rules precedes those of logic, for the latter are to make the former explicit (cf. e.g. Brandom, 2000, 22). Furthermore, the inferentialist has a notion of propositional content as detached from syntax: two sentences express the same proposition just in case their grounds of assertion and appropriate consequences are the same. Thus, it might be feasible to regard and's rules as operating on antecedently given propositional content, thereby allowing us to at least circumvent the issues connected to syntactic form.

others' arguments, and we need nothing further to generate a consequence relation, respectively.

This runs counter to what Glanzberg claims (cf. e.g. 2015, 72, and sect. 3.5), namely that we must first have extra-linguistic guidance to what is to be identified as logical. Then, we must abstract away from the meanings of non-logical expressions by introducing a space of models, before finally idealising away all the grammatical complications present in natural language. In the use-theoretic setting, however, the identification process does not need extra-linguistic guidance, since all we need are rules with a specific profile – to wit: gap-free and purely inferential ones. Neither do we need to abstract away from any meanings: The topic-neutral aspect of logical constants is cashed out in terms of having purely inferential rules. This does not force us to go beyond the building materials of our semantics, namely rules of inference. Thus, *no substantial extra*-linguistic guidance is needed to get from natural language to logic. The only step we might have to make as well is to idealise away grammatical complications. This, however, would not endanger my defense of the LNLT, for I explicitly decided to discount syntactical mismatches as irrelevant to the matter at heart, viz. a mismatch on the level of meaning. For example, the sentence in (1) means the same as A came to the party and B came to the party, hence the mismatch is only syntactical in nature.

4.3 The Argument from the Orthodox Point of View

Before concluding this investigation, I would like to return to a point made in section 3.2. Recall that I consciously chose an unorthodox semantics as a basis, leaving the possibility of combining inferentialist (meta)semantics with standard referential truth-conditions open. For those readers that sympathise with such an approach, or for those that work in mainstream formal semantics, it ought to be of interest to see how my argument can modified to accommodate such approaches.

As far as my overall argument is concerned, I submit it works equally well for orthodox approaches. First, observe that there is no inherent incoherence of combining logical inferentialism with referential truth-conditional semantics, provided the possibility for truth-conditions to be determined fully by (inferential) use is accommodated. Nevertheless, since both the validity and the nature of inference rules retains explanatory priority, any other criteria – such as necessary truth-preservation and permutation invariance – become 'symptoms', rather than 'causes', for logicality.

Clearly, the criteria for logical constancy would have to be extended in order to guarantee that the logical constants make a truth-conditional contribution to their compound sentences, via the truth-conditions of the subsentences. Hence, one could adopt the following additional criterion:

(LC4) The introduction and elimination rules must determine a consistent set of truth-conditions for the relevant compound sentences.

This is still a use-theoretic account of logical constancy in virtue of truth-conditions being fully determined by inference rules, themselves born out of use. Of course, the way this determination proceeds is by two "semantic principles": (i) these inferences ought to preserve truth and (ii) no sentence can be both true and false at the same time (Rumfitt, 2000, 806). In contrast to the case of logical constants, (Val) and (LCons) need no amendment: (LC4) simply ensures that any chaining of I- and E-rules preserve truth, and given their purely inferential character as well as meaning-determinative status, formally and necessarily so.

All that is left then is to check whether (LC4) mismatches with these slightly altered meaning-theoretic assumptions. If we have a general story as to how inference rules determine truth-conditions – as the incorporation of orthodox semantics would require – there is no reason to suppose this would be impossible. To the contrary: Analytically definable predicates and even the standard rules for $\wedge$ determine such conditions via both their I- and E-rules. Thus, there would be no mismatch either, thanks to (II), and by (I), the argument concludes. Not even our considerations from section 4.2 are overturned: If, for example, and-I and and-E can be equated with those for $\wedge$, we would yield the corresponding Tarskian clause in the case of and as well.

Unfortunately, this line of reasoning has a few problems. While (LC4) itself may mesh well with logical inferentialism, the issues concern the meaning-theoretic background. Anyone who wishes to pursue such an approach would need to tackle the following problems:

1. Which types of rules – I or E – are relevant for fixing truth-conditions? And how does the answer cohere with the two-aspect model of meaning?

2. Are the semantic principles evoked above inferentialistically acceptable?

3. Why should rules of inference preserve truth to begin with?

Concerning 1., it might be thought that this is the function of the I-rules, with E-rules merely being consequences of the truth-conditions thus determined. The problem is that certain sentences do not behave like this: it is in no way part of the ascription conditions for murder that the act was wrong, but a central consequence (cf. Hanfling, 2000, 146). Yet if both types of rules fix truth-conditions, it would be part of the truth-conditions for This is murder that the act was wrong – perhaps an undesirable consequence of the account. As far as 2. is concerned, the issue lies with truth, falsity and inconsistency being evoked as seemingly independent semantic notions, even if they share conceptual ties with the notion of inference. Perhaps those ties are what might give an answer to 3., but I am less optimistic about them sufficing to preserve the spirit of inferentialism with respect to 2.

Alas, this approach needs further investigation, which must be left for future occasions. However, I hope to have shown that orthodox semanticists have something to gain from my approach as well.

5 Conclusion and Outlook

In this paper, I have scrutinised the claim that natural language contains a relation of logical consequence. We first delved into Glanzberg's arguments supporting the falsity of it. In a nutshell, Glanzberg claimed that there is a conceptual mismatch between the restrictions imposed by the background theory of meaning and the ascription conditions for logical consequence and constants. After introducing the idea of inferentialism and the accompanying proof-theoretic conceptions of logical notions, we saw that none of the ascription conditions for logical constants exceed the background assumptions in the inferentialist case. As logical constants are all that are needed to generate a logical consequence relation, we concluded that under such meaning-theoretic assumptions and conceptions of logical notions, no conceptual mismatch obtains. Additionally, it gives us reasons to be optimistic about the existence of logical constants, and therefore of consequence relations, in natural language.

This result not only counters the observations made by Glanzberg, but makes further contributions. First, as Hjortland pointed out, Glanzberg's findings suggest that 'semantic intuitions' cannot be grounds for knowledge of logic (Hjortland, 2019, 206f.). The reason is that if natural language is devoid of logical consequence relations and constants, the reliability of judgements grounded in linguistic competence is challenged. However, the above result points in a different direction. It sheds a vindicating light on the widespread practice of relying on natural language in formal logic, and of appeals to logical notions in semantic theorising.

Second, to the extent that the reader already agrees with the introductory remarks concerning the closeness of logic and natural language, the result of this paper motivates inferentialism over its main competitor. For if we have good reason to suppose that natural language and logic *must be* at least close, the failure of mainstream semantics and the success of inferentialism to account for this ought to count in favour of the latter.

Third, if my strategies concerning and should prove fruitful, further connectives deserve excavation. If we were to recover purely inferential, gap-free I- and E-rules for (sentential) negation, we could recover a form of propositional logic from natural language *alone*. This would further strengthen the epistemological consequences mentioned above: Knowledge of logic may, at its core, turn out to centrally involve linguistic competence with a *natural* language after all.

Naturally, the assumptions made along the way are far from uncontroversial. This holds even more so for the repeated appeals to semantic analyticity. However, my argument vindicates the practices in logic teaching and research as crucially relying on natural language. It would further vindicate appeals to logical notions in mainstream formal semantics – assuming the issues raised in section 4.3 to be surmountable. Meanwhile, we ought to be able to rest assured that, assuming the inferentialist premisses to be defensible, there is no 'deep' mismatch between natural language and formal logic.

References

Aloni, M., F. Roelofsen, and N. Teiler (2017). What's wrong with believing whether? *Proceedings of SALT 27*, 248–265.

Baker, G. P. and P. M. S. Hacker (Eds.) (2009). *Wittgenstein: Rules, Grammar and Necessity*. Oxford: Blackwell. Volume II of the Analytical Commentary on the Philosophical Investigations, Essays and Exegesis of §§185–242.

Balcerak-Jackson, B. (2021). Model-theoretic semantics as model-based science. *Synthese 199*, 3061–3081.

Bayer, S. (1996). The coordination fo unlike categories. *Language 72*(3), 579–616.

Beall, J., G. Restall, and G. Sagi (2019). Logical consequence. In *The Stanford Encyclopedia of Philosophy (Spring 2019 Edition)*. Edward N. Zalta (Ed.). URL = https://plato.stanford.edu/entries/logical-consequence/.

Ben-Yami, H. (2010). A wittgensteinian solution to the sorites. *Philosophical Investigations 33*(3), 229–244.

Boghossian, P. (2011). Williamson on the a priori and the analytic. *Philosophy and Phenomenological Research 82*(2), 488–497.

Brandom, R. B. (1983). Asserting. *Noûs 17*(4), 637–650.

Brandom, R. B. (1994). *Making It Explicit*. Harvard: Harvard University Press.

Brandom, R. B. (2000). *Articulating Reasons*. New York/Cambridge (Mass.): Harvard University Press.

Büttner, M. K. (2021). Truth in virtue of meaning reconsidered. *Philosophical Papers 50*(1-2), 109–139.

Dekker, P. (2011). Dynamic semantics. In C. Maienborn, K. von Heusinger, and P. Portner (Eds.), *Semantics: An International Handbook of Natural Language Meaning*, Volume 1, pp. 923–945. De Gruyter Mouton.

Dummett, M. (1996). *The Seas of Language*. Oxford: Oxford University Press.

Etchemendy, J. (2008). Reflections on consequence. In D. Patterson (Ed.), *New Essays on Tarski and Philosophy*, pp. 263–299. Oxford University Press.

Gentzen, G. (1934). Untersuchungen über das logische schliessen. *Mathematische Zeitschrift 39*, 405–431.

Glanzberg, M. (2015). Logical consequence and natural language. In C. R. Caret and O. T. Hjortland (Eds.), *Foundations of Logical Consequence*, pp. 71–120. Oxford University Press.

Glock, H.-J. (2003). *Quine and Davidson on Language, Thought and Reality.* Cambridge: Cambridge University Press.

Glock, H.-J. (2018). Why rules ought to matter. In O. Beran, V. Kolman, and L. Koreň (Eds.), *From Rules to Meaning: New Essays on Inferentialism*, pp. 63–80. Routledge.

Hanfling, O. (2000). *Philosophy and Ordinary Language.* New York/London: Routledge.

Hjortland, O. T. (2019). What counts as evidence for a logical theory? *Australasian Journal of Philosophy 16*(7), 250–282.

Kaplan, D. (1997). The meaning of 'ouch' and 'oops': Explorations in the theory of meaning as use. unpublished manuscript.

Lepore, E. (1983). What model-theoretic semantics cannot do. *Synthese 54*, 167–187.

Linnebo, Ø. (2022). Plural quantification. In *The Stanford Encyclopedia of Philosophy (Spring 2022 Edition).* Edward N. Zalta (Ed.). URL = https://plato.stanford.edu/entries/plural-quant/.

MacFarlane, J. (2009). Pragmatism and inferentialism. In B. Weiss and J. Wanderer (Eds.), *Reading Brandom. On Making It Explicit*, pp. 81–95. Routledge.

MacFarlane, J. (2017). Logical constants. In *The Stanford Encyclopedia of Philosophy (Winter 2017 Edition).* Edward N. Zalta (Ed.). URL = https://plato.stanford.edu/entries/logical-constants/.

McGee, V. (1985). A counter-example to modus ponens. *The Journal of Philosophy 82*(9), 462–471.

Miller, A. (2002). What is the manifestation argument? *Pacific Philosophy Quarterly 83*(4), 352–383.

Murzi, J. and F. Steinberger (2017). Inferentialism. In B. Hale, C. Wright, and A. Miller (Eds.), *A Companion to the Philosophy of Language* (2 ed.), Volume 1, pp. 197–224. Wiley-Blackwell.

Partee, B. H. (2004). Compositionality. In *Compositionality in Formal Semantics. Selected Papers*. Blackwell Publishing Ltd.

Partee, B. H. and M. Rooth (2002). Generalized conjunction and type ambiguity. In P. Portner and B. H. Partee (Eds.), *Formal Semantics: The Essential Readings.* Blackwell Publishing.

Peregrin, J. (2009). Inferentialism and the compositionality of meaning. *International Review of Pragmatics 1*, 154–181.

Prawitz, D. (2005). Logical consequence: A constructivist perspective. In S. Shapiro (Ed.), *The Oxford Handbook of Philosophy of Mathematics and Logic*, pp. 671–695. Oxford University Press.

Prawitz, D. (2006). Meaning approached via proofs. *Synthese 148*, 507–524.

Prior, A. (1960). The runabout inference ticket. *Analysis 21*(2), 38–39.

Putnam, H. (1975). *Mind, Language and Reality.* Cambridge: Cambridge University Press.

Rumfitt, I. (2000). 'yes' and 'no'. *Mind 109*(436), 781–823.

Rumfitt, I. (2017). Against harmony. In B. Hale, C. Wright, and A. Miller (Eds.), *A Companion to the Philosophy of Language* (2 ed.), Volume 1, pp. 225–249. Wiley-Blackwell.

Ryle, G. (1949). *The Concept of Mind.* Chicago: The University of Chicago Press.

Sagi, G. (2020). Considerations on logical consequence and natural language. *Dialectica 74*(2), 179–200.

Schroeder, S. (2009). Analytic truths and grammatical propositions. In H.-J. Glock and J. Hyman (Eds.), *Wittgenstein and Analytic Philosophy*, pp. 83–108. Oxford University Press.

Shoesmith, D. J. and T. J. Smiley (1978). *Multiple Conclusion Logic.* Cambridge: Cambridge University Press.

Skorupski, J. (2017). Meaning, use and verification. In B. Hale, C. Wright, and A. Miller (Eds.), *A Companion to the Philosophy of Language* (2 ed.), Volume 1, pp. 73–106. Wiley-Blackwell.

Soysal, Z. (2019). *Formal Analyticity.* Ph. D. thesis, Harvard University.

Steinberger, F. (2009). *Harmony and Logical Inferentialism.* Ph. D. thesis, University of Cambridge.

Steinberger, F. (2011). What harmony could and could not be. *Australasian Journal of Philosophy 89*(4), 617–639.

Strawson, P. F. (1950). On referring. *Mind 59*(235), 320–344.

Strawson, P. F. (1952). *Introduction to Logical Theory.* London: Methuen & Co. LTD.

Williamson, T. (2007). *The Philosophy of Philosophy.* Oxford: Blackwell.

Wittgenstein, L. (1953). *Philosophische Untersuchungen.* Frankfurt am Main: Suhrkamp.

Yalcin, S. (2017). Semantics as model-based science. In D. Ball and B. Rabern (Eds.), *The Science of Meaning: Essays on the Metatheory of Natural Language Semantics*, pp. 334–360. Oxford University Press.