Penultimate draft of a paper forthcoming in *Ergo*. Please check to see if it has been published before citing as forthcoming.

Ergo AN OPEN ACCESS JOURNAL OF PHILOSOPHY

Meaning, the Context Principle, and the Sequent Calculus

REA GOLAN
Ben Gurion University of the Negev, Be'er Sheva, Israel

Natural deduction inference rules seem to reflect the way we actually reason. Hence, many if not most inferentialist theories maintain that meaning is conferred on linguistic expressions by natural deduction rules, rather than the more abstract alternative of sequent rules. In the present paper, I argue, to the contrary, that an inferentialist theory of meaning must take a somewhat metainferential form, whereby the meanings of linguistic expressions—in particular, the logical constants—are conferred by sequent rules, conceived of as licensing inferences between inferences. To that end, I propose that sequent rules are to be understood as *inferential semantic clauses*, i.e., as playing a role analogous to that of the semantic clauses in truth conditional theories of meaning. I establish this proposal on the basis of an argument to the effect that inferentialists must adopt an extended version of Frege's context principle, according to which only within a complete sentence potentially serving either as a premise or as a conclusion of an inference do words have meaning. The resultant inferentialist theory is not disconnected from the way we actually reason. Moreover, on its basis one can make a meaning-theoretic case for rather unintuitive features of many sequent calculi, such as multiple conclusions and various ways of going substructural.

1. Introduction

Inferentialism is a semantic theory that connects the *meanings* of linguistic expressions with their *use* as premises and conclusions in inferences. (Brandom 1983, 1994, 2000; Dummett 1991; Francez 2015; Hlobil and Brandom 2024; Peregrin 2014; Prawitz 1965) Traditionally, inferentialists hold that meaning is conferred on linguistic expressions—in particular, the logical constants—by corresponding natural deduction rules, since those rules seem to reflect the way we actually reason. (Dummett 1991; Francez 2015; Prawitz 1965) Owing to various consider-

Contact: Rea Golan <reagolan@gmail.com>

ations, however, the primacy of natural deduction systems has been gradually superseded by sequent calculi (Dicher and Paoli 2021, forthcoming; Hlobil and Brandom 2024; Schroeder-Heister 2012).

My overall purpose in the present paper will be to ground this shift from natural deduction to sequent calculus in general inferentialist meaning-theoretic considerations. I plan to do so against arguments to the effect that, unlike natural deduction systems, sequent calculi do not reflect our actual inferential practice, and, in fact, many of them involve unintuitive features. Thus, the latter calculi violate what Steinberger (2011: 335) calls the principle of answerability:

> *Principle of Answerability*: Only such deductive systems are permissible as can be seen to be suitably connected to our ordinary deductive inferential practices.[1]

To the contrary, I shall argue that an inferential theory of meaning must take a somewhat metainferential form, whereby the meanings of linguistic expressions—in particular, the logical constants—are conferred by rules governing sequent-to-sequent derivations, conceived of as inferences between inferences. Moreover, I shall argue that the principle of answerability is not as informative a principle; that it does not rule out many—albeit not all—of those sequent calculi that exhibit unintuitive features such as multiple conclusions or various kinds of substructurality.

The key insight on which my account will be based is that inferentialists have to extend Frege's context principle. Whereas Frege (1953) held that only within a complete sentence do words have meaning, I suggest that only within a complete sentence potentially serving either as a premise or as a conclusion of an inference do words have meaning. Drawing on this extended principle, I shall argue that the sequent operational rules should be understood as *inferential semantic clauses*, i.e., they play a role analogous to that of semantic clauses in truth conditional theories of meaning. It is this way of understanding sequent rules that has the above hinted at meaning-theoretic implications.

The rest of the paper is organized as follows. In Section 2, I discuss Frege's context principle and argue that inferentialists have to extend it. In Section 3, drawing on the extended principle, I introduce and expound my proposal to understand sequent rules as inferential semantic clauses. In Section 4, I suggest that, despite appearances to the contrary, the resultant semantic theory is not disconnected from the way we actually reason. In Section 5, I discuss the principle of answerability in more detail. In particular, I argue that both relevant implication rules and multiple-conclusion rules—*prima facie*, two striking examples of rules

1. Based on this principle, Steinberger argues against multiple-conclusion systems, but it is only natural to extend his argument to other unintuitive features of sequent calculi. I discuss concrete examples in Section 5 below.

that are not suitably connected to our actual inferential practice—need not be ruled out by the principle. I also point out that my strategy of defending such rules is not without limitations.

2. The Context Principle

According to Frege's context principle, "only in the context of a sentence do words have any meaning." Frege (1953: 73) The idea underlying this principle is that the proposition—or rather, the sentence that expresses that proposition—is the minimal unit of pragmatic significance, namely, the minimal unit to which a pragmatic force may be attached. As Dummett puts it:

> "Since it is only by means of a sentence that we may perform a linguistic act—that we can say anything—the possession of a sense by a word or complex expression short of a sentence cannot consist in anything else but its being governed by a general rule which partially specifies the sense of sentences containing it." Dummett (1973: 4)

On the face of it, the context principle is in tension—albeit not an unresolvable tension—with yet another tenet of meaning theories, namely, compositionality.[2] The tension stems from the fact that, whereas the context principle is clearly a top-down principle, compositionality is bottom-up, according to which the meaning of a word consists in its systematic contribution to the meanings of the sentences in which it occurs. Thus, compositionality tells us that to know the meaning of a negative sentence of the form $\neg p$, we first have to know the meaning of negation, but the context principle seems to be stating the opposite: that negation has meaning only in the context of a complete sentence, most notably, a sentence of the form $\neg p$, whose principal operator is negation.

To resolve the tension, Dummett makes a distinction between the *order of recognition* and the *order of explanation*. Dummett (1973: 4-7) From a cognitive point of view—that is, when the order of recognition is concerned—grasping the meaning of a complex expression must be conceptually preceded by grasping the meanings of its subexpressions. Yet, when we are concerned with the order of explanation, i.e., the pursuit of an answer to the more theoretical question of what it is for an expression to have a meaning in the first place, our starting point must be the minimal unit that is of pragmatic significance, namely, the minimal unit that can be used in performing a speech act—a sentence; the meaning of a subsentential expression is thus explained derivatively, based on the meanings of sentences in which it occurs.

2. I do not address here the historical question of whether Frege was committed both to the context principle and to compositionality at the same time. See Milne (1986) for a discussion.

While I accept this distinction and recognize that it resolves the tension, I doubt whether a sentence is the minimal unit of pragmatic significance; hence, it is unclear whether the order of explanation could begin with sentences. In particular, I doubt that assertion—the speech act with which Dummett purports to explain the priority of sentences over their parts—is autonomous and can itself be understood independently of other assertions.

To understand the latter claim, one may do well to recall a fundamental inferentialist thesis, namely, that the pragmatic significance of an assertion can only be understood within a "game of giving and asking for reasons."[3] Thus, Brandom (1983: 643-644) claims that a parrot who reliably utters "It's getting warmer" only as the temperature climbs past 80 degrees (Fahrenheit) cannot be conceptually understood as making an assertion, because it can neither justify this claim nor draw relevant inferences from it.[4] That is to say, an assertion can only be understood as something that requires justification by prior assertions and that justifies subsequent assertions. The way Brandom puts it:

> "What is it that we are doing when we assert, claim, or declare something? The general answer is that we are undertaking a certain kind of commitment. [...] The idea is that assertings (performances that are overt undertakings of assertional commitments) are in the fundamental case what reasons are asked for, and what giving a reason always consists in. The kind of commitment that a claim of the assertional sort is an expression of is something that can stand in need of (and so be liable to the demand for) a reason; and it is something that can be offered as a reason. [...] The idea exploited here, then, is that assertions are fundamentally fodder for inferences." Brandom (1994: 167-168)

Now, my above claim is a straightforward implication of Brandom's thesis: if assertions are fundamentally unautonomous, "fodder for inferences," then the order of explanation cannot begin with assertions.[5] Hence, when we address

3. A referee pointed out that this is so only under the assumption that the ratiocinative use of language comes first in the order of explanation. After all, there are other, non-ratiocinative uses of language. This is definitely right, but the assumption is a common one in the general literature on meaning theories. See, e.g., the introduction to Evans and McDowell (1976), and passim.

4. One may argue, at this point, that the very notion of inference is itself in need of further elaboration. For the sake of neutrality I leave this notion unanalyzed in this paper. It should be noted that whether or not the notion of inference is analyzable has no effect on my above claim about the the order of explanation.

5. A similar claim can be found in Sellars, who traces it all the way back to Kant: "Kant was on the right track when he insisted that just as concepts are essentially (and not accidentally) items which can occur in judgments, so judgments (and, therefore, indirectly concepts) are essentially (and not accidentally) items which can occur in reasonings or arguments." Sellars (1953: 314)

the question of what it is for a linguistic expression to have a meaning, the minimal unit with which we should begin is not a single sentence—the content of an assertion—but a sentence standing in inferential relations with some other sentences, i.e., a sentence serving—or rather, potentially serving—either as a premise or as a conclusion in an inference. To put it systematically, whereas Frege's context principle amounts to the statement that "only in the context of a sentence do words have any meaning" Frege (1953: 73), the extended context principle amounts to the statement:

> *Extended Context Principle*: Only in the context of a sentence potentially serving either as a premise or as a conclusion in an inference do words have any meaning.

Notice that this principle appeals to sentences as *potentially* serving as premises and conclusions in inferences. Otherwise, a sentence's meaning could strictly depend on the specific inference in which it occurs, making it impossible to account for the meaning of the same sentence uniformly. For example, "That's red" may be said to mean one thing when it figures as a conclusion in the inference from "That's scarlet" to "That's red," and another when it figures as a premise in the inference from "That's red" to "That's colored." Regarding the view under consideration, though, the meaning of "That's red" is to be understood in terms of its *systematic* and hence *potential* contribution to the goodness of all relevant inferences rather than in terms of its contribution to a specific inference. In effect, the same phenomenon already occurs in the case of truth conditional theories of meaning: just as the meaning of negation in truth conditional theories is associated not with its contribution to the truth value of a specific sentence, but with its systematic and hence potential contribution to the truth value of all relevant sentences, so, too, should the meaning of a sentence (or the parts thereof) be associated, in an inferential theory of meaning, with its systematic and hence potential contribution to the goodness of all relevant inferences, namely, inferences in which the sentence figures either as a premise or as a conclusion.

3. Sequent Rules as Inferential Semantic Clauses

I would like to argue that the extended context principle provides us both with a new understanding of sequent rules and with an argument to the effect that an inferential theory of meaning should account for the meanings of linguistic expressions "metainferentially," namely, based on the operational sequent rules conceived of as licensing inferences between inferences. My argument relies on an analogy with truth conditional theories of meaning and how Frege's context principle underlies the semantic clauses in such theories.

In truth conditional theories of meaning, let us recall, each semantic clause specifies the meaning of a given expression by giving truth conditions to sentences whose principal operator is that expression. Giving truth conditions in this manner already presupposes Frege's context principle: the relevant semantic value—be it truth, falsity, or any other truth value—is not assigned to the expression in question "in isolation," as it were. Thus, the word "or," which stands for the concept of disjunction ($\vee$), is not assigned any semantic value on its own; rather, semantic values are assigned to *sentences*, e.g., sentences in the form of $A \vee B$ whose principal operator is disjunction. Moreover, the semantic values—truth values—are assigned to such sentences compositionally, based on the truth values assigned to the disjuncts A and B (or based on the truth values of the disjuncts at several points of evaluation in the case of an intensional disjunction.) In other words, the semantic clause for disjunction specifies the *contribution* of disjunction to the truth value of a sentence of the form $A \vee B$ compositionally, based on the truth values of A and B.

Analogously, in an inferential theory of meaning, a semantic value should be assigned not to an expression such as $\vee$ "in isolation," but always in an appropriate context. Under the extended context principle, the appropriate context is not a complete sentence of the form $A \vee B$, but an inference in which such a sentence figures either as a premise or as a conclusion. Relatedly, the relevant semantic value in an inferential theory of meaning cannot be a truth value—truth values are not assigned to inferences—but the goodness of an inference, or validity.[6] It follows that an inferential semantic clause must specify the inferential contribution of an expression—say, a disjunction—in terms of the goodness of an inference in which a sentence whose principal operator is disjunction figures either as a premise or as a conclusion. If we want our semantics to be compositional as well, such a contribution must be specified in terms of the ways in which the disjuncts contribute to the goodness of inferences. In sum, then, the inferential semantic clauses for disjunction must specify the inferential contribution of $A \vee B$—its premissory and conclusory roles—based on the premissory and conclusory roles of A and B. In this way, we arrive at semantic clauses in the form of sequent rules such as

$$\frac{\Gamma, A \vdash C \quad \Gamma, B \vdash C}{\Gamma, A \vee B \vdash C} \qquad \frac{\Gamma \vdash A}{\Gamma \vdash A \vee B} \qquad \frac{\Gamma \vdash B}{\Gamma \vdash A \vee B}$$

The left rule specifies the premissory role of $A \vee B$ compositionally, based on the premissory roles of the disjuncts A, B, and each of the two right rules specifies the conclusory role of $A \vee B$ compositionally, based on the conclusory role of one of A, B. As for the other sequent rules (see Table 1 and Table 2), it is immediately

6. I will be talking about the goodness of an inference rather than its validity so as to remain neutral on the question of whether such goodness is purely logical, or whether we adopt a more permissive view that allows for materially yet non-formally good inferences.

noticeable that the single-conclusion sequent rules for conjunction, implication, and the quantifiers can be interpreted in a similar way; in Section 3.2 below, I explain why the negation sequent rules can also be interpreted along the same lines; in Section 5, I explain why multiple-conclusion rules can be so interpreted.[7]

The extended context principle thus provides us with a new understanding of the operational sequent rules, namely, as *inferential semantic clauses*. Notice that we have here two independent sorts of semantic clauses specifying premissory and conclusory roles, respectively. This situation resembles, though it is not entirely analogous to, the semantic clauses in some paraconsistent logics, in which truth conditions and falsity conditions come apart (see, e.g., Priest (2006: 73-81).) It is just that, instead of having two independent semantic values (each with its own semantic clauses), we have here only one semantic value—the goodness of inferences—but two ways in which a sentence may contribute to such goodness, namely, either as a premise, or as a conclusion.

It is for this reason that the extended context principle makes our theory of meaning metainferential in nature: its semantic clauses must specify the goodness of one inference in terms of the goodness of other inferences, namely, they must be of the metalinguistic form of inferences between deducibility statements (as in, e.g., Hacking (1979)), or, to put it more simply, inferences between inferences.[8] Indeed, it is a virtue of the theory that it provides a conceptual justification for reading sequent rules metainferentially and hence literally.[9] In this regard, the theory contrasts with views that regard such rules as mere principles under which a consequence relation may be closed Cobreros et al. (2013), or interpret them in terms of an abstract consequence relation through translation functions. (Dicher and Paoli (2019, 2021, forthcoming))

To be sure, the idea of inferential semantic clauses requires further unpacking. For a start, I discuss in the following three subsections the issues of (i) how inferential semantic clauses should be formulated, (ii) the roles the structural rules play in this picture and the negation clauses, and (iii) what conditions must

7. There are actually various variations of the sequent rules that appear in Tables 1 and 2 (see Negri and von Plato (2001) for comprehensive discussions), but my suggestion to view them as semantic clauses is independent of which version we adopt.

8. A referee for this journal mentioned that more orthodox proof-theoretic semantics approaches such as Prawitz's also seem to specify the goodness of one inference in terms of the goodness of other inferences. In particular they had in mind the definition of proof-theoretic validity which results from quantifying over atomic bases. Francez (2015: 214-219) There is much to be said here, but notice that in the latter approach it is not inferences that matter, but rather proofs. Hence, it is at the very least questionable whether this approach can be genuinely regarded as a "used-based" theory of meaning.

9. To be precise, we should distinguish between metainferential rules and metainferences proper, namely, the instances that fall under such rules. Clearly, sequent rules fall under the former category, not the latter.

Table 1. Single-conclusion rules

Axioms:

$$\Gamma, A \vdash A$$

Structural Rules:

$$\frac{\Gamma, A, A \vdash B}{\Gamma, A \vdash B}\ \textit{contraction} \qquad \frac{\Gamma, A, B \vdash C}{\Gamma, B, A \vdash C}\ \textit{permutation}$$

$$\frac{\Gamma \vdash A}{\Gamma, B \vdash A}\ \textit{left weakening} \qquad \frac{\Gamma \vdash}{\Gamma \vdash A}\ \textit{right weakening} \qquad \frac{\Gamma \vdash A \qquad \Gamma, A \vdash B}{\Gamma \vdash B}\ \textit{cut}$$

Operational Rules:

$$\frac{\Gamma \vdash A}{\Gamma, \neg A \vdash}\ \neg L \qquad \frac{\Gamma, A \vdash}{\Gamma \vdash \neg A}\ \neg R$$

$$\frac{\Gamma, A \vdash C \qquad \Gamma, B \vdash C}{\Gamma, A \vee B \vdash C}\ \vee L \qquad \frac{\Gamma \vdash A}{\Gamma \vdash A \vee B}\ \vee R1 \qquad \frac{\Gamma \vdash B}{\Gamma \vdash A \vee B}\ \vee R2$$

$$\frac{\Gamma, A, B \vdash C}{\Gamma, A \wedge B \vdash C}\ L\wedge \qquad \frac{\Gamma \vdash A \qquad \Gamma \vdash B}{\Gamma \vdash A \wedge B}\ R\wedge$$

$$\frac{\Gamma \vdash A \qquad \Gamma, B \vdash C}{\Gamma, A \to B \vdash C}\ L\!\to \qquad \frac{\Gamma, A \vdash B}{\Gamma \vdash A \to B}\ R\!\to$$

$$\frac{\Gamma, A(t/x) \vdash C}{\Gamma, \forall x A \vdash C}\ L\forall \qquad \frac{\Gamma \vdash A(y/x)}{\Gamma \vdash \forall x A}\ R\forall$$

$$\frac{\Gamma, A(y/x) \vdash C}{\Gamma, \exists x A \vdash C}\ L\exists \qquad \frac{\Gamma \vdash A(t/x)}{\Gamma \vdash \exists x A}\ R\exists$$

be met in order for an operational sequent rule to qualify as a semantic clause. It should be noted from the outset, though, that within certain restrictions that are discussed below, there is no problem in allowing for various ways of formulating sequent rules. My agenda in this regard may be described as "methodological pluralism." That is, in contrast some authors—notably, Dummett (1991)—I do not subscribe to the view that meaning-theoretic considerations support one specific logic over all other logics, even though such considerations may help us decide between rival logics. Nor do I commit myself to a version of logical pluralism. Rather, I regard various ways of formulating sequent rules or even going substructural as meaning-theoretically feasible, without taking a stance on issues such as which logic is preferable, or whether there is meaning-variance between rival logics.[10]

3.1. How to formulate an inferential semantic clause.

One may object that if we are to seriously consider sequent rules as semantic clauses, then just like truth conditional semantic clauses, sequent rules should take the metalinguistic form of *conditionals* whose antecedents and consequents are deducibility statements, rather than *rules* licensing inferences between those statements. I believe that such an objection amounts to a red herring. Conditionals and inference rules are closely related, regardless of the language (metalanguage) in which they are formulated. [11] Thus, nothing essential seems to depend on whether we formulate our clauses as conditionals or inference rules.

A related objection would be that semantic clauses usually take the symmetric form of a *biconditional* rather than a conditional. Thus, a semantic clause for disjunction (in many logics) is of the form "$A \vee B$ is true *if and only if* either A or B is true." Inference rules, on the other hand are typically asymmetric, as they do not allow us to draw each premise from the conclusion. In response, I would like to point out that there is no theoretical consideration that obliges us to reject asymmetric clauses. Most notably, compositionality only requires that the truth values of complex sentences be determined—rather than interdetermined—by the truth values of their constituents. Indeed, there exist logics in which some (truth-conditional) semantic clauses are essentially formulated as conditionals rather than biconditionals, e.g., da Costa's calculi C_n. Da Costa and Alves (1977) Hence, symmetry does not seem essential to the very notion of a semantic clause. That said, it is also worth mentioning that many—though not all—of the operational sequent rules are in effect symmetric, or *invertible* as they are often called in the literature, as it is possible to derive each premise sequent from the conclusion

10. Thus, my approach is in the spirit of Restall (2005), except that I do not take a stand on whether there is meaning-variance between rival logics.

11. See Ryle (2009: 244-260) for an extensive discussion of this point.

sequent.[12]

A referee for this journal objected that ordinarily, compositionality serves as a supervenience principle: settling how things are with the meanings of the components of a sentence, as well as with the way those components are combined into the sentence, must settle how things are with the meaning of the sentence as a whole. And this does seem to require a biconditional, to make sure that the way things are with the sentence's meaning as a whole is really settled, not just constrained. For example, if our (truth-conditional) semantics for disjunction has only the clause "$A \vee B$ is true if either A or B is true", then we have failed to settle the truth conditions of $A \vee B$. For, in cases where A and B are both untrue, it is compatible with our semantics that $A \vee B$ is true in some such cases and untrue in others. So we have not settled its truth conditions. *Mutandis mutatis*, this objection extends directly to sequent rules construed as inferential semantic clauses.

Yet, notice that this objection does not disqualify noninvertible rules as such. Specifically, in the case of disjunction, all that is required is that whenever we can derive a sequent containing $A \vee B$—say, a sequent of the form $\Gamma \vdash A \vee B$—we can do so either from $\Gamma \vdash A$, or from $\Gamma \vdash B$, and so the latter sequents need not both be derivable. Namely, we can stick to the single-conclusion noninvertible rules, to the extent that whenever a sequent containing some complex expression can be derived, it can be derived with at least one—even if not all—of the expression's rules.

Now, in Section 3.3 below, I shall pose the condition that the structural rules must all be conserved by the operational rules. Provided this condition, every sequent that can be derived, can in principle be derived using only the operational rules, with a proof tree whose leaves are all atomic sequents (the uninformed reader should consult with Section 3.3 at this point). In such a case, the only way to introduce a complex expression such as $A \vee B$ on either side of the turnstile goes through one of the expression's operational rules. Thus, meaning is fully settled by the operational rules (conceived of as inferential semantic clauses) even if they are noninvertible, and so the referee's objection does not hold.[13]

3.2. *The structural rules and negation.*

I believe that the structural rules express in this picture certain presuppositions about the notion of inference underlying the inferential semantic clauses. Here again, the analogy between inferential and truth-conditional theories of meaning is illuminating. Just as in a truth-conditional theory of meaning based on classical logic "truth" and "falsity" are presupposed to be mutually exhaustive and exclusive—and the semantic clauses are understood accordingly—so, too, the

12. One can find an extensive discussion of invertibility in Negri and von Plato (2001).

13. For the same reason, there is no need for a sequent calculus with elimination rules such as the one advocated by Dicher and Paoli (forthcoming).

structural rules should be regarded as expressing basic presuppositions we have about the notion of inference underlying the inferential semantic clauses.

Thus, the axioms, which take the form of $\Gamma, A \vdash A$, express the presupposition that our notion of inference is reflexive. These axioms serve as starting points for our derivations.[14] The structural rules of contraction and permutation

$$\frac{\Gamma, A, A \vdash B}{\Gamma, A \vdash B} \qquad \frac{\Gamma, A, B \vdash C}{\Gamma, B, A \vdash C}$$

express, respectively, the presuppositions that premises may be used as many times as we like, and that the order in which we use them does not matter. The cut rule

$$\frac{\Gamma \vdash A \qquad \Gamma, A \vdash B}{\Gamma \vdash B}$$

expresses the presupposition that the underlying notion of inference is transitive. The left weakening rule

$$\frac{\Gamma \vdash A}{\Gamma, B \vdash A}$$

expresses the presupposition that the underlying notion of inference is monotonic.

Finally, Special attention should be given to the right weakening rule

$$\frac{\Gamma \vdash}{\Gamma \vdash A}$$

which sets the convention that an empty right-hand side in a sequent is a placeholder for every sentence A. With this convention in mind, it is easy to understand the negation rules as semantic clauses:

$$\frac{\Gamma \vdash A}{\Gamma, \neg A \vdash} \neg L \qquad \frac{\Gamma, A \vdash}{\Gamma \vdash \neg A} \neg R$$

$\neg L$ specifies the premissory role of $\neg A$ compositionally, based on conclusory role of A: if one may derive A from Γ, then one may derive any sentence from Γ together with $\neg A$; $\neg R$ specifies the conclusory role of $\neg A$ compositionally, based on the premissory role of A: if one may derive any sentence from Γ together with A, then one may derive $\neg A$ from Γ.[15]

14. Faithful to my methodologically pluralist agenda, I do not assume that these axioms are indispensable. However, to replace them, one has to introduce alternative axiomatic starting points for derivations, e.g., atomic sequents like $p \vdash q$ that represent "material inferences" such as the one from "It is raining" to "The streets will be wet" (Sellars (1953), see also Hlobil and Brandom (2024) for an extensive development of this idea). Thus, in contrast to authors like Dicher and Paoli (forthcoming), I insist that in any case, our derivations have to have axiomatic starting points. In Section 3.3, I shall argue that this has to do with compositionality.

15. This reading of the right rule gets really close to Dummett's suggestion that the genuine

3.3. *The necessary conditions for an operational sequent rule to qualify as a semantic clause.*

First and foremost, the semantic clauses must not violate our presuppositions about the underlying notion of inference. Here too the analogy between inferential and truth-conditional theories is illuminating. Consider a binary connective $A\#B$, whose clause states that $A\#B$ is both true and false if either member of A, B is false, and true otherwise. Clearly, $A\#B$ cannot be integrated into a meaning theory based on classical logic, as its clause violates the presupposition that truth and falsity are mutually exclusive.

Perhaps the most famous proof-theoretic analogue of $A\#B$ is Prior's connective *tonk* Prior (1960), which is governed by the following rules:

$$\frac{\Gamma, B \vdash \Delta}{\Gamma, A tonk B \vdash \Delta} \qquad \frac{\Gamma \vdash A, \Delta}{\Gamma \vdash A tonk B, \Delta}$$

When introduced into a system in which the cut rule is present, *tonk* trivializes the deducibility relation, making $A \vdash B$ derivable for every A and B.

Such triviality is particularly bad from the perspective of the inferential theory outlined above. Consider again the analogy between inferential and truth-conditional theories of meaning. In the latter theories, sentences cannot all be assigned the same truth value, as that would render the semantic clauses redundant and the theory uninformative . Indeed, from a semantic perspective, there would be no point in assigning truth values to sentences—indeed, no point in considering truth values as having any semantic significance— if each and every sentence is assigned the same value (Priest (2005: 68-69); see also Evans and McDowell (1976: xix-xxiii)). Likewise, if sequent rules are regarded as inferential semantic clauses—indeed, as having any semantic significance—then it cannot be the case that every inference is good.

Consequently, the sequent rules for tonk are illegitimate in the context of a transitive deducibility relation precisely because, in trivializing that relation, they undermine the theory of meaning in which they purport to take part. Since trivializing the consequence relation is not a viable option, introducing the *tonk* rules in the presence of cut should be regarded as violating one of our presuppositions about the underlying notion of inference, namely, the one expressed by the cut rule.

That the rules for a given operator must not violate our presuppositions about the underlying notion of inference should thus be a condition for the possibility of considering those rules inferential semantic clauses. This is how the inferential

introduction rule of negation (in a natural deduction setting) should take the following form: to derive a sentence of the form $\neg A$ we have to show that any atomic sentence follows from A. Dummett (1991: 294-295) theory outlined above deals with *tonk* and its cronies. To guard against such a violation, we should require that the operational rules *conserve* the structural rules. In other words, we should require that the addition of operational rules to the original language—the one that has no operators—not result in extending the deducibility relation defined by the structural rules over that original language.[16]

This point merits further discussion. On the face of it, there is no problem with forcing the extension to be conservative: all we have to do is grant the structural rules the status of basic, nonderivative rules of the systems. After all, the addition of any operational rule to the orginal language as governed by the structural rules will trivially never violate those structural rules. Such a suggestion has already been made by Peacocke (1976), in rebuttal to Hacking's defense of conservativeness. Hacking responds thus:

> "But what is wrong with just adding the operational rules and the structural rules stated for arbitrary formulas built up out of the logical constants? Nothing, except that one is not then defining logical constants in connection with some previous language fragment. Rather one is creating, as a totality, a new system of logistic." Hacking (1979: 298)

Put in our terms, Hacking's rejoinder to Peacocke is that there is nothing *technically* wrong with forcing the structural rules; yet, were we to do so, these rules would no longer express presuppositions about the notion of inference underlying the inferential semantic clauses. Rather, in such a case the structural rules would be on a par with the operational rules in determining the meanings of linguistic expressions.[17] Needless to say, in such a case the operational rules can hardly be regarded as inferential semantic clauses.

In addition to conservativeness, the inferential semantic clauses must not violate the tenets of the theory outlined above. In particular, an inferential semantic clause for a given expression must specify either the expression's premissory role

16. Notice that the above formulation of conservativeness pertains primarily to the structural rules, rather than the deducibility relation. Thus, my formulation of conservativeness follows (Hacking (1979); see also Kremer (1988)), rather than Belnap (1962) This formulation clearly entails that conservativeness should be tested with respect to the atomic language, rather than a language that already includes some operators. See Francez (2015: 76-80) for further discussion of conservativenes.

17. Dicher (2020) puts forward such a thesis, which he calls the "co-determination thesis." His idea, in brief, is that not only structural rules affect the output of the operational rules, but also that operational rule imposes some minimal requirements on the structure of derivations. Thus, he maintains that some connectives, e.g., multiplicative disjunction, determine a multiple-conclusion setting, and so multiple-conclusion rules—though remote from our actual practice of reasoning—may be justified on meaning-theoretic grounds. I shall not discuss this theory in depth here. Suffice it to point out that my theory does better with respect to connecting sequent rules with our actual practice of reasoning. See my discussion in Section 5 below, particularly, footnote 31.

or its conclusory role, and such roles must be specified compositionally, i.e., based on the subexpressions' premissory or conclusory roles.

To meet the first tenet, the rules for a given expression must be *explicit* as well as *symmetrical*.[18] Explicitness requires that there be only one essential occurrence of the expression whose meaning is given in the conclusion sequent of each rule; hence, explicitness guarantees that each rule specifies either a premissory role or a conclusory role. Symmetry requires that there be at least one left rule and one right rule; hence, both roles are indeed specified.

As for the second tenet, notice that the requirement of conservativeness, when applied to the axioms, yields the property known as the *deducibility of identicals*. Namely, every axiom of the form $\Gamma, A \vdash A$ is derivable from the atomic instances of the axioms. Therefore, each sequent may be derived by the operational rules from atomic instances of the axioms.[19] That is to say, the goodness of a complex inference is ultimately determined compositionally, by the goodness of non-complex inferences.[20]

4. Metainferential Rules and the Way We Actually Reason

As I said, one virtue of the semantic theory outlined above is that it sheds light on the nature of sequent rules, which take a metainferential form that is somewhat unintuitive. There is an apparent problem, however, with any use-based theory of meaning that appeals to metainferential rules. After all, such rules are in principle remote from our actual inferential practice: in reasoning we paradigmatically infer a conclusion from certain premises; in any case, we do not infer an inference from other inferences. Hence, our actual inferential practice does not seem to be governed by metainferential rules. Consequently, there arises

18. I employ here the terms coined by Wansing (1994).

19. Recall footnote 14: if we do not endorse those axioms, we must endorse other axiomatic starting points for our derivations, in the form of atomic sequents such as $p \vdash q$. Still, conservativeness implies in that case that each sequent may be derived from such atomic sequents. Hence, compositionality still holds. See Hlobil and Brandom (2024: 103-157) for more detail, including a conservativeness proof.

20. The deducibility of identicals is a *global* condition that holds for the entire system. As for a more *local* and concrete condition on the structure of rules, there is Dummett's *complexity condition* Dummett (1991: 258), according to which the formulas in the conclusion sequent must be of higher complexity than those in the premise sequent(s). (Dummett formulates the condition for natural deduction rules, but the adjustment to sequent rules is straightforward.) Truth be told, I am inclined to think that compositionality requires a condition stronger than Dummett's. For, rules such as $\frac{\Gamma \vdash \neg A}{\Gamma \vdash \neg(A \wedge B)}$ meet Dummett's condition, even though $\neg A$ is not a subexpression of $\neg(A \wedge B)$. For simplicity, and since nothing essential in what follows depends on this, I focus only on to the global condition of the deducibilty of identicals, leaving the discussion of local conditions for another occasion.

the challenge of reconnecting the theory with our actual practice of reasoning, whose meaningfulness—namely, how come the sentences with which we reason have propositional content—the theory aims to explain. Tackling this problem must precede any specific answerability issue, and I plan to do this in two steps, to which Sections 4.1 and 4.3 are respectively dedicated. In Section 4.2, I take a pause and prepare the groundwork for the second step, by describing two ways in which sequent rules can be read.

4.1. The Davidsonian Distinction

At first approximation, since metainferential rules are regarded here as semantic clauses, it makes sense to think of them in a way similar to how Davidson (1977) thinks of some seemingly fishy elements that figure in in the semantic clauses of his Tarskian truth conditional theory of meaning, namely, reference and satisfaction (as applied to open sentences.) To recall, Davidson comes across the following problem. On the one hand, reference and satisfaction fall short of any primary semantic significance in the order of explanation: one cannot perform a speech act merely by uttering a referential expression or an open sentence; moreover, it is questionable whether there are such things as open sentences in a natural language such as English. On the other hand, reference and satisfaction are needed to given compositional truth conditions for certain sentences in Davidson's Tarskian-truth conditional theory.

To resolve this tension, Davidson draws a distinction between *explanation within the theory* and *explanation of the theory*. It is only *within the theory of meaning* that reference and satisfaction play a role. But when it comes to connecting the theory with our linguistic practice—thereby giving *explanation of the theory* as a whole—it is the notion of truth that does the explaining. Davidson points out the close analogy with physics:

> "[W]e explain macroscopic phenomena by postulating an unobserved fine structure. But the theory is tested at the macroscopic level[...] I suggest that words, meanings of words, reference and satisfaction are posits we need to implement a theory of truth. They serve this purpose without needing independent confirmation or empirical basis." Davidson (1977: 254-255)

In a similar manner—again, only at first approximation—I suggest that metainferential rules be regarded as theoretical posits of our inferential theory of meaning. As such, these posits need not have correlates in the actual practice of reasoning, whose meaningfulness the theory purports to explain. In other words, for purely theoretical reasons—i.e., the extended context principle—we have to accommodate inferential semantic clauses in the form of sequent rules, even though such rules do not govern our actual practice of reasoning. The resultant theory of meaning thus explains as a whole the *meaningfulness* of reasoning—i.e., how come the sentences with which we reason have propositional content—but it does so on the basis of posits that are remote from our actual inferential practice.

Yet, Davidson's distinction between explanation within the theory and explanation of the theory applies only imperfectly to the inferential theory outlined above, for at least two reasons. First, reference and satisfaction are, in the final account, of mere derivative semantic significance for Davidson. By contrast, I have construed sequent rules as semantic clauses; hence, they play a pivotal role in the semantic theory I have outlined. Second, inferentialism is, after all, a use-based theory of meaning. It would therefore be odd for such a theory to be construed such that its principles cannot directly govern the actual practice on which it purports to draw. Thus, even though a variation of Davidson's distinction seems to apply to our case, the relation between metainferential rules and the rules governing our actual practice of reasoning must be further refined.

4.2. Two Readings of Sequent Rules

As a first step toward refining the Davidsonian distinction, we should bear in mind that there is an alternative reading of sequent rules. Rather than as licensing inferences between inferences, they can be understood as natural deduction rules. As Dummett puts it:

> "A sequent is simply a device for carrying along the premises on the strength of which the succedent is asserted at a particular stage of the deduction, or the hypotheses under which it is asserted, with the sentence asserted. Sequents can thus be used to represent a deduction from contingent premises to a contingent conclusion; and we may call the succedent of the last sequent in a deduction so represented the final conclusion of the whole deduction, and the sentences comprising its antecedent the initial premises of the deduction. But the whole point of allowing inferences that discharge hypotheses is that, in such a case, we cannot describe the inference as a transition from the assertion of certain statements as premises to the assertion of some other statement as conclusion: the conclusion is asserted on the strength of its being possible to derive certain statements from certain hypotheses. At least one of the premises of the inference is therefore not a statement but a deduction, most easily representable by a sequent; for convenience, we therefore represent every line as a sequent, whether a hypothesis is discharged or not." Dummett (1991: 186)

In essence, Dummett points out that (single-conclusion) sequent rules can be read as natural deduction rules, equipped with a device for keeping track of undischarged assumptions at each stage of the derivation.[21] I shall refer to this reading—according to which sequent rules are not that far off the actual practice of reasoning—as the "natural deduction reading." The natural deduction reading differs from the reading of sequent rules as licensing inferences between inferences, which I shall call the "metainferential reading."

To stress the difference between these two readings, let us consider a few examples. First, consider the right conjunction rule:

$$\frac{\Gamma \vdash A \quad \Gamma \vdash B}{\Gamma \vdash A \wedge B}$$

In the metainferential reading, the rule is read literally, as licensing an inference between inferences, namely, an inference from the two *premise inferences* $\Gamma \vdash A$ and $\Gamma \vdash B$ to the *conclusion inference* $\Gamma \vdash A \wedge B$. In the natural deduction reading, by contrast, the rule licenses an inference from the two *assumptions* A and B to the *conclusion* $A \wedge B$. The sentences in the set Γ take no part in this derivation; the set Γ is thus only mentioned in the latter reading, where it serves as a mere reminder that the assumptions A and B are not given to us "for free," but rather, were previously derived from some undischarged assumptions in Γ.[22]

Next, consider the right implication rule:

$$\frac{\Gamma, A \vdash B}{\Gamma \vdash A \to B}$$

In the metainferential reading, this rule is read literally as licensing an inference between inferences. In the natural deduction reading, by contrast, we have a rule that allows us to infer the conclusion $A \to B$, provided that we may infer B from the assumption A. The assumption A is discharged in the derivation of $A \to B$, and so it no longer appears on the left-hand side of the turnstile, which is the

21. A referee pointed out that there might be an exegetical issue here. Dummett may be interpreted as concerned first and foremost with sequents in the strict sense, not with sequent calculi, and their rules. By contrast, as is clear in my above discussion, I take it that Dummett is concerned with sequent calculi and their rules. In any case, for what follows my interpretation need only be feasible.

22. The above rule is the *context-sharing* rule for conjunction. Notice that the *context-mixing* right rule

$$\frac{\Sigma \vdash A \quad \Gamma \vdash B}{\Sigma, \Gamma \vdash A \wedge B}$$

is read in a similar manner. In the metainferential reading, the context-mixing rule licenses an inference between inferences, whereas in the natural deduction reading, this is the same rule as the context-sharing rule, licensing the derivation of $A \wedge B$ from A and B. Thus in the natural deduction reading, the context-mixing rule is the same as the context-sharing rule, except that the former allows for cases in which each premise is derived from a different set of undischarged assumptions.

place for the undischarged assumptions from which $A \to B$ may be derived. As before, the sentences in the set Γ take no part in this derivation.

Finally, consider the left disjunction rule:

$$\frac{\Gamma, A \vdash C \qquad \Gamma, B \vdash C}{\Gamma, A \vee B \vdash C}$$

Once again, the metainferential reading of this rule is literal and straightforward: the rule simply licenses an inference between inferences. On the other hand, the natural deduction reading of the left disjunction rule—indeed, of any left rule—is less intuitive. For, the right-hand side of each sequent is the same, namely, the sentence C, and so the "inferential action" that the rule licenses has predominantly to do with managing one's assumptions at each step. It is for this reason that Dummett calls the sentences on the left-hand side of the bottom sequent the "initial premises" of the rule. One way of understanding his point is that such a rule should be partly read bottom up, as it were, as it allows us to reach the final conclusion from the initial premises via the top sequents. Thus, in the natural deduction reading, the initial premise of this rule —the "major" premise, as it is called in the natural deduction literature—is the disjunction $A \vee B$. The other two premises are in the form of the inferences from assumption A to C and from assumption B to C. The final conclusion is C, and when it is derived, assumptions A and B are discharged, and so they do not appear on the left-hand side of the bottom sequent. That is to say, in the natural deduction reading, we have an inference whose conclusion is C, and whose premises are: (i) $A \vee B$, (ii) the inference of C from assumption A, and (iii) an inference of C from assumption B. The set Γ, which mentions additional undischarged assumptions, remains the same throughout the derivation that the rule licenses, and so the assumptions in Γ are not read as taking any part in that derivation.

More generally, here is a procedure for reading a given sequent rule in a natural deduction manner.[23] First, we observe what sentences appear on the left-hand side of all sequents; those are regarded as undischarged assumptions that take no part in the derivation licensed by the rule. At this point, there are two options: either there are active sentences on the left-hand side of the sequents, or there aren't. In the latter case, we are presented with a right rule that makes no use of discharging assumptions, like the right conjunction rule; such a rule is read from the top down, where all the inferential activity takes place on the right-hand side of the sequents. In the former case, there are two options: either there is an active formula on the left-hand side of the bottom sequent, or there isn't. If

23. The procedures described above is based on the familiar translations between natural deduction and sequent calculi devised by Prawitz (1965: 90-91), and further elaborated, e.g., by von Plato (2003). I do operate here at a more abstract level, though, as the distinction between left and right rules is explained, rather than used, during the procedure.

there isn't, we are presented with a right rule that makes use of assumptions (to be discharged) like the right conditional rule; again, such a rule is read from the top down. Finally, if there is an active formula on the left-hand side of the bottom sequent, we are presented with a left rule that is read partly bottom up: the rule licenses the derivation of the final conclusion (the formula on the right-hand side of the bottom sequent) from the initial premise (the active formula on the left-hand side of that sequent) via the top sequents, which are regarded as inferences from assumptions that are discharged.

Now, the fact that in the natural deduction reading the right rules are read from the top down whereas the left rules are partly read from the bottom up, makes such a reading cumbersome and anything but straightforward when it comes to considering whole proof trees, rather than single rule applications. This is already implied in Dummett's comment that in the bottom sequent of such a proof tree, the antecedent consists of the initial premises of the whole derivation, whereas the succedent is its final conclusion. The apparent problem with this reading is that it is unclear how to reach the final conclusion from the initial premises. To do so, one roughly has to apply first the left rules, climbing all the way up the tree, and then the right rules, climbing all the way back down.[24] But this is only a rough description of the process, because structural rules may also be involved, and because the order in which the rules are applied may matter.[25]

I shall therefore say that the natural deduction reading is *nonuniform*, whereby I mean that it is at the very least unclear how to read in a natural deduction manner a whole derivation consisting of several, consecutive rule applications. Yet, such nonuniformity does not threaten the cogency of the natural deduction reading. Consider the following analogy: the easiness of putting together two pieces of a puzzle (upon recognizing that the two are a fit) does not make puzzle-solving an uncomplicated task. Likewise, the mere fact that each sequent rule can be given a natural deduction reading need not imply that whole derivations can also be easily or uniquely read in a similar manner. After all, reading—to say nothing of constructing—natural deduction proofs is not always a trivial task, even to experienced logicians, as it can be tricky to keep track of where each assumption is discharged. It is nothing if not natural that such trickiness is reflected in the natural deduction reading of sequent derivations.

Summing up, whereas sequent rules license inferences between inferences in the metainferential reading, they at most license inferring a conclusion on the basis of an inference with assumptions that are discharged in the natural deduction reading. Each reading has its own pluses and minuses. The natural deduction

24. Notice that—again, only roughly—such a reading presents the proof tree in what is called a "normal form" in the natural deduction literature. Prawitz (1965)

25. For these reasons, in many cases there is no single natural deduction reading of a proof tree.

reading is intuitive and thus preferable to the metainferential reading with respect to single applications of rules. The metainferential reading is preferable to the natural deduction reading with respect to whole derivations. All in all, it would be ideal to find a way to combine these readings, enjoying the best of all possible worlds. This is the task to which I turn next.

4.3. Knowing-How Description of Rules vs. Theoretical Description of Rules

Despite appearances to the contrary, I would like to argue that the two readings are *compatible* with one another. Just like Rubin's vase, we have here two complementary ways of looking at the same picture, rather than two different pictures. For, the two readings can be regarded as two descriptions *of the same rules* rather than different, incompatible ways of understanding the mathematical structure that is the sequent calculus. To understand this point, we should make a distinction between a *knowing-how* description of rules governing some practice, and a *theoretical* description of the same rules *within* a given theory, which is a refinement of Davidson's distinction between explanation of the theory and explanation within the theory.

Consider, for instance, the rules governing the activity of riding a bike.[26] On the one hand, these rules can be described in a knowing-how fashion, which is what parents do when they teach their child how to ride a bike. Such a description must be systematic, but it need not be purely linguistic, as it may involve gestures, demonstrations, etc. On the other hand, the rules governing riding a bike may also be described theoretically, e.g., *within* the theory of physics. The theoretical description is purely linguistic—and partly mathematical—and involves theoretical concepts such as *momentum*. It is to the theoretical description that we appeal if we are interested in understanding what makes riding a bike physically possible rather than knowing how to ride a bike. On the other hand, one does not have to know the relevant physics behind biking to know how to ride a bike. In general, one does not have to know the theoretical description of the rules that govern a practice or activity within a given theory—not even the concepts this theoretical description involves—to know how to follow the rules by internalizing the relevant knowing-how description.[27]

26. I assume here, for the sake of the argument, that riding a bike is a rule-governed activity.

27. The above distinction is a faily natural outcome of Anscombe's observation that one and the same action admits of more than one correct description, and so the same item of behavior can be understood as intentional under some descriptions but not under others. Anscombe (2000) Likewise, I claim that the rules governing one's actions admit of more than one correct description. Notice futher, that riding a bicycle can be understood as intentional under the knowing-how description of its rules, but it is at the very least unclear whether it may be understood as intentional under the theoretical description of the rules within the theory of physics.

Now, I suggest that the two readings of sequent rules may each be associated with a different description of the rules governing the practice that is reasoning. On the one hand, there is the knowing-how description of these rules, namely, what is required to know how to make an inference, or how to draw a conclusion from certain premises. This description is evidently associated with the natural deduction reading, according to which sequent rules govern the actual practice that is reasoning. Yet, the natural deduction reading is unsuitable for meaning-theoretical purposes because it violates the extended context principle. It turns out that just as one cannot explain, based on the knowing-how description of the rules for riding a bike, what makes riding a bike physically possible, likewise, one cannot explain, based on the knowing-how description of sequent rules, why reasoning is meaningful, namely, why the sentences with which we reason have propositional content. To answer this question, we must appeal to the description of the sequent rules within the inferential theory of meaning, i.e., the one given in the metainferential reading, which construes sequent rules as inferential semantic clauses.

Recall the challenge with which I was grappling all along. It stems from two considerations that pull us in opposite directions. On the one hand, the extended context principle makes us endorse semantic clauses in the metainferential form of sequent rules; on the other hand, the use-based nature of the inferential theory lets us acknowledge as meaning-conferring only those rules that can be viewed as governing our actual practice of reasoning. This challenge can be overcome if we notice that whether or not a given rule can be viewed as governing our practice is dependent, among other things, on our description of the rule and the practice. Think for example about the rules of chess. It is because some such rule allows one, in certain circumstances, to "checkmate one's opponent" that we may say that the rule governs our activity of playing chess. The physical description of the same rule as allowing one to "move some piece of wood on the board in a certain way" is just as correct, but cannot serve as a reason to count the rule among chess rules. That is to say, a rule may be said to govern our practice to the extent that there is a description of it as governing our practice, even if there are other descriptions as well.[28] Likewise, whether or not a rule can

28. Relatedly, but more generally, a linguistic activity can be described in one of two ways: either behavioristically, in terms of the sequences of sounds and the regularities governing their utterances, or as a means of communication that consists in exchanging meaningful expressions. McDowell (1998: 87-131) draws on this Anscombian distinction to criticize Dummett's demand that a theory of meaning for a given language must be *full-blooded*, i.e., one whose learning involves acquiring the concepts of that language, as opposed to a theory that only provides an interpretation of the language to someone who has already learned its concepts. Dummett (1993: 1-93) One problem with full-blooded theories, according to McDowell, is that under the behavioral description, linguistic activity cannot be regarded as rational:

"Not every true characterization of what someone is doing displays it so as to be be viewed as governing our inferential practice is dependent on there being a description of it as governing our inferential practice. But we may well admit of other descriptions of the rule, each with its own purpose. Thus, the knowing-how description of sequent rules is required for regarding such rules as governing our inferential practice, whereas the metainferential description of the rules within the theory of meaning serves to explain why the sentences with which we reason have propositional content.

I believe that the above considerations have have wide-ranging implications. To bring the paper to a close, though, I focus in Section 5 below on one of these implications, namely, how various unintuitive features of sequent calculi may be accounted for in our theory of meaning.

5. What to Say about Answerability?

In light of the above considerations, answerability figures as a somewhat uninformative principle. To understand this, recall that in the natural deduction reading, the contexts do not take part in the inferential activity licensed by the rules; rather, they are regarded as a mere device for keeping track of undischarged assumptions. Thus, insofar as the unintuitive features involved in a given rule are restricted to the structure of the contexts, we may well be able to read the rule in a natural deduction manner, as answerable to our actual practice of reasoning. To make this point clearer, I discuss below how we can account in our theory of meaning for the unintuitive features involved in two cases: relevant implication and multiple-conclusion rules.

Example 1. *Relevant Implication*. In an early attempt to formulate a sequent calculus for the implication fragment of various relevance logics, Anderson and Belnap (1975: 51-57) come across the following problem. Given the usual implication right rule

> intelligible as a bit of rational agency. [...] [O]ur capacity for rational understanding gets its grip on linguistic behaviour precisely under descriptions of the sort Dummett disallows. [...] [Dummett] seems rather to be telling us to describe linguistic behaviour in a way that would obliterate [...] the rational intelligibility that we know it has." McDowell (1998: 112-113)

Notice that the inferential theory I have outlined above cannot, in principle, be full-blooded, as it acknowledges a dichotomy between a knowing-how description of rules and a description of these rules within the theory of meaning. This implication is somewhat surprising, because it runs counter Dummett's insistence on both inferentialism and full-bloodedness. Given the scope of the present paper, I shall not discuss this point in greater detail, and it will have to await another occasion.

$$\frac{\Gamma, A \vdash B}{\Gamma \vdash A \to B} \to R$$

as well as the structural rules of weakening and permutation, one can derive the relevantly-unacceptable principle $A \to (B \to A)$ in the following way:

$$\dfrac{\dfrac{\dfrac{\dfrac{A \vdash A}{B, A \vdash A}\ \textit{weakening}}{A, B \vdash A}\ \textit{permutation}}{A \vdash B \to A}\ \to R}{\vdash A \to (B \to A)}\ \to R$$

To block this derivation, Anderson and Belnap endorse a suggestion by Kripke (1959) to read the comma on the left-hand side of the turnstile in terms of nested implication, rather than extensional conjunction. Namely, they suggest that a sequent of the form $A_1, A_2, ..., A_n \vdash B$ expresses the statement $A_1 \to (A_2 \to (... \to (A_n \to B))...)$ rather than $(A_1 \wedge A_2 \wedge ... \wedge A_n) \to B$. In this reading, the permutation rule does not hold, and so the above derivation is blocked.[29]

Needless to say, reading the comma on the left-hand side of the turnstile in terms of nested implication is a striking example of an unintuitive feature of a sequent rule. In the natural deduction reading, however, this feature does not show up. For, in the natural deduction reading, the right rule is identical to the classical right rule: it licenses the derivation of $A \to B$, provided that we can infer B from A. The sentences in the context Γ take no part in this derivation, no matter how this context is structured. After all, in the natural deduction reading, this context is regarded as a mere device with which to keep track of undischarged assumptions; the extra complication of putting these assumptions together in terms of nested implication should accordingly be regarded merely as a means of restricting the availability of each assumption at different stages of the derivation, following a major principle of relevance logics, namely, that each such assumption must be used in the derivation in a specific way.

Therefore, there is no problem with reading the relevant implication right rule as a natural deduction rule that governs our actual inferential practice, even though the rule involves a highly substructural, and indeed unintuitive, feature. In the metainferential reading, by contrast, the rule figures as remote from our practice: not only does it license inferences between inferences, but the premises in such inferences are put together in terms of nested implication! Yet, the metainferential reading is associated with the description of the rule within the theory of meaning. And just as the description of the rules for riding a bicycle

29. Anderson and Belnap's solution also requires that $\to$ be regarded as a modal operator. Dunn (1974) later solved the problem by having two (structural) modes of combining assumptions: in terms of extensional conjunction and intensional conjunction. Given my purposes in the present paper, I shall not discuss such technicalities in greater detail.

within the theory of physics may be unintuitive or even cognitively unavailable to the rider, so too may the meaning-theoretic description of the rules govening one's reasoning be unintuitive and even cognitively unavailable to the reasoner.

Example 2. *Multiple Conclusions.* Introduced originally by Gentzen (1964), multiple-conclusion rules are well known: they make use of sequents in which a bunch of conclusions appear together on the right-hand side of the turnstile (see Table 2). Usually, a multiple-conclusion sequent of the form $A_1, \ldots, A_n \vdash B_1, \ldots, B_m$ is regarded as equivalent to the statement $(A_1 \wedge ... \wedge A_n) \rightarrow (B_1 \vee \ldots \vee B_m)$. Following the mainstream view on this matter (Shoesmith and Smiley (1978); Steinberger (2011); Tennant (1997: 5)), I take it that there is no room for such disjunctive multiple conclusions in the actual practice of reasoning.[30]

Now, based on his principle of answerability, Steinberger rules out other ways of understanding multiple conclusions:

> "[T]he principle of answerability precludes a purely formalistic reading of the sentences occurring in the conclusion as an alternative to the disjunctive reading. [...] Unless there is a way of showing how such a [...] story ties in with our ordinary practice, the multiple-conclusion setting would be condemned to the status of mere artifice, which, though perhaps of mathematical interest, would be wholly devoid of meaning-theoretic significance and thus useless to the inferentialist." Steinberger (2011: 340-341)

It is at this point that I take issue with Steinberger. In light of the theory outlined above, I claim, we can accommodate multiple-conclusion rules in our theory, to the extent that these rules can suitably be given a natural deduction reading in which they are described as governing inferences with only one conclusion. The other conclusions—those whose derivation is not licensed by the rules—can and should be read in a rather formalistic manner.[31]

30. See Restall (2005: 12) for a rebuttal of this view and Steinberger (2011: 341-346) for a rejoinder. It is also worth mentioning that there is an alternative to such arguments, namely, the kind of bilateralism associated, mainly, with Restall (2005, 2020) and Ripley (2013). For them, a multiple-conclusion sequent of the form $\Gamma \vdash \Delta$ is understood as the statement that it is incoherent to assert all the sentences in Γ and deny all the sentences in Δ. Likewise, Kneale views multiple-conclusion rules as specifying "developments," or "limits" of the premises, namely, "the field within which truth must lie if the premisses are to be accepted" Kneale (1956: 245) (see also Shoesmith and Smiley (1978: 133-147) for a discussion.) Yet, these views are too revisionary. We should do our best to find a way of reconnecting multiple-conclusion rules with our actual practice of reasoning.

31. In this regard, my theory is similar to Dicher (2020). To put it briefly, Dicher argues that multiple-conclusions should be thought of as epiphenomena of the meanings of the connectives. Thus, he argues, there is no need to worry about the place of multiple conclusions in our practice of reasoning. Rather, one should look for features of the practice that are suitably connected

Table 2. Multiple-conclusion rules

Axioms:

$$\Gamma, A \vdash A, \Delta$$

Structural Rules:

$$\frac{\Gamma, A, A \vdash \Delta}{\Gamma, A \vdash \Delta}\ \textit{left contraction} \qquad \frac{\Gamma \vdash A, A, \Delta}{\Gamma \vdash A, \Delta}\ \textit{right contraction}$$

$$\frac{\Gamma, A, B \vdash \Delta}{\Gamma, B, A \vdash \Delta}\ \textit{left permutation} \qquad \frac{\Gamma \vdash A, B, \Delta}{\Gamma \vdash B, A, \Delta}\ \textit{right permutation}$$

$$\frac{\Gamma \vdash \Delta}{\Gamma, A \vdash \Delta}\ \textit{left weakening} \qquad \frac{\Gamma \vdash \Delta}{\Gamma \vdash A, \Delta}\ \textit{right weakening}$$

$$\frac{\Gamma, A \vdash \Delta \qquad \Gamma \vdash A, \Delta}{\Gamma \vdash \Delta}\ \textit{cut}$$

Operational Rules:

$$\frac{\Gamma \vdash A, \Delta}{\Gamma, \neg A \vdash \Delta}\ \neg L \qquad \frac{\Gamma, A \vdash \Delta}{\Gamma \vdash \neg A, \Delta}\ \neg R$$

$$\frac{\Gamma, A \vdash \Delta \qquad \Gamma, B \vdash \Delta}{\Gamma, A \vee B \vdash \Delta}\ \vee L \qquad \frac{\Gamma \vdash A, B, \Delta}{\Gamma \vdash A \vee B, \Delta}\ \vee R1 \qquad \frac{\Gamma \vdash A, \Delta}{\Gamma \vdash A \vee B}\ \vee R2 \qquad \frac{\Gamma \vdash B, \Delta}{\Gamma \vdash A \vee B}\ \vee R3$$

$$\frac{\Gamma, A, B \vdash \Delta}{\Gamma, A \wedge B \vdash \Delta}\ L\wedge \qquad \frac{\Gamma \vdash A, \Delta \qquad \Gamma \vdash B, \Delta}{\Gamma \vdash A \wedge B, \Delta}\ R\wedge$$

$$\frac{\Gamma \vdash A, \Delta \qquad \Gamma, B \vdash \Delta}{\Gamma, A \rightarrow B \vdash \Delta}\ L\rightarrow \qquad \frac{\Gamma, A \vdash B, \Delta}{\Gamma \vdash A \rightarrow B, \Delta}\ R\rightarrow$$

$$\frac{\Gamma, A(t/x) \vdash \Delta}{\Gamma, \forall x A \vdash \Delta}\ L\forall \qquad \frac{\Gamma \vdash A(y/x), \Delta}{\Gamma \vdash \forall x A, \Delta}\ R\forall$$

$$\frac{\Gamma, A(y/x) \vdash \Delta}{\Gamma, \exists x A \vdash \Delta}\ L\exists \qquad \frac{\Gamma \vdash A(t/x), \Delta}{\Gamma \vdash \exists x A, \Delta}\ R\exists$$

To be concrete, what I suggest is that multiple conclusions—in particular, contexts on the right-hand side of the turnstile—should be understood as playing a role analogous to the one Dummett allocates to the left-hand side contexts. Just as the left contexts are regarded in the natural deduction reading as a device for keeping track of undischarged assumptions that we make, the right-hand side contexts should be regarded as a device for keeping track of assumptions that we refrain from making. The idea is that in reasoning we often operate not only under some assumptions, but also under refraining from making some other assumptions, e.g., because we believe that the latter are highly unlikely. Thus, at least some conclusions that we draw from our assumptions might be defeasible: they are drawn *unless* we reconsider the status of our unmade assumptions. For example, "Tweety flies" may be drawn from "Tweety is a bird" *unless* we reconsider the status of our unmade assumption that "Tweety is a penguin." It is thus reasonable to keep track of the status of such unmade assumptions. This, I suggest, is what the contexts on the right-hand side do.[32]

Thus, the multiple-conclusion right rules, e.g.,

$$\frac{\Gamma, A \vdash B, \Delta}{\Gamma \vdash A \to B, \Delta}\ {\to}R \qquad \frac{\Gamma \vdash A, \Delta \quad \Gamma \vdash B, \Delta}{\Gamma \vdash A \wedge B, \Delta}\ {\wedge}R$$

are read exactly like the corresponding single-conclusion right rules, namely, as licensing the derivation of $A \to B$ provided that we may infer B from A, and as licensing the derivation of $A \wedge B$ from A and B, respectively. Both left- and right-hand side contexts take no part in those derivations, and are regarded merely as devices for keeping track of undischarged assumptions and unmade assumptions, respectively.

The right structural rules are read analogously to the left ones, namely, as expressing certain presuppositions about the very notion of unmade assumptions. Thus, the right contraction and permutation rules

$$\frac{\Gamma \vdash A, A, \Delta}{\Gamma \vdash A, \Delta} \qquad \frac{\Gamma \vdash B, A, \Delta}{\Gamma \vdash A, B, \Delta}$$

with multiple-conclusion derivations, namely, aspects of the behaviour of the connectives that determine multiple-conclusion derivations. A detailed discussion of this theory goes beyond the scope of the present paper. Suffice it to point out once more that such a manoeuvre, in the context of a used-based theory, makes that theory much less attractive.

32. I leave here open the question of whether refraining from making an assumption amounts to rejecting that assumption. On the one hand, it seems that to draw "Tweety flies" from "Tweety is a bird" one need only maintain that "Tweety is a penguin" is not the default case, rather than outright reject it. On the other hand, "maintaining that it is *not* the default case" seems to involve the concept of negation, which is out of place here. As a referee for this journal pointed out, we might have to go bilateralist along the lines of Restall (2005, 2020) and Ripley (2013) at this point, acknowledging some weak form of rejection as an independent speech act.

express, respectively, the presuppositions that an unmade assumption may be referred to as many times as we like, and that the order of these assumptions does not matter.

Likewise, recall the convention that an empty right-hand side is a placeholder for every sentence A. In a multiple-conclusion setting, the same convention is expressed by the right weakening rule

$$\frac{\Gamma \vdash \Delta}{\Gamma \vdash A, \Delta}$$

since no formula is active in the premise sequent. In other words, the premise sequent is read as an inference from Γ to any conclusion A, where the the conclusions in Δ are merely regarded as unmade assumptions.[33]

Therefore, the negation multiple-conclusion rules

$$\frac{\Gamma \vdash A, \Delta}{\Gamma, \neg A \vdash \Delta}\ \neg L \qquad \frac{\Gamma, A \vdash \Delta}{\Gamma \vdash \neg A, \Delta}\ \neg R$$

are read just like the corresponding single-conclusion rules, since the conclusions in Δ are regarded merely as unmade assumptions. Namely, the left rule allows us to infer every conclusion from $\neg A$, provided that we may infer A, and the right rule allows us to infer $\neg A$, provided that every sentence follows from A.

Unfortunately, not every multiple-conclusion rule can be given such a natural deduction reading. Hence, my strategy of defending such rules is not unlimited in scope. Most notably, the right disjunction rule $\vee R1$

$$\frac{\Gamma \vdash A, B, \Delta}{\Gamma \vdash A \vee B, \Delta}\ \vee R1$$

resists the above reading. This is because a rule cannot be given a natural deduction reading unless each of its sequents contains at most one active conclusion. By contrast, both A and B are active in the premise sequent of $\vee R1$. On the other hand, the alternative right rules $\vee R2$ and $\vee R3$

$$\frac{\Gamma \vdash A, \Delta}{\Gamma \vdash A \vee B, \Delta}\ \vee R2 \qquad \frac{\Gamma \vdash B, \Delta}{\Gamma \vdash A \vee B, \Delta}\ \vee R3$$

can clearly be given a natural deduction reading.

33. Hence, one can also make sure that the multiple-conclusion cut rules—both multiplicative and the additive

$$\frac{\Gamma \vdash A, \Delta \quad \Gamma, A \vdash \Delta}{\Gamma \vdash \Delta} \qquad \frac{\Sigma \vdash A, \Pi \quad \Gamma, A \vdash \Delta}{\Sigma, \Gamma \vdash \Delta, \Pi}$$

express the same presupposition that the single conclusion cut rule expresses, namely, that the underlying notion of inference is transitive.

In effect, $\lor R1$ is interderivable with $\lor R2$ and $\lor R3$, given the (right) weakening rule, and either contraction or cut.[34] Hence, if weakening and either contraction or cut are available, we may also admit the $\lor R1$ rule, but only derivatively, as a convenient abbreviation of both $\lor R2$ and $\lor R3$. As it cannot be given a natural deduction reading, this rule cannot be regarded as governing our actual inferential practice, and so it is ruled out by the principle of answerability.[35]

Let us now consider the following derivation of the law of excluded middle:

$$\cfrac{\cfrac{\cfrac{\cfrac{A \vdash A}{A \vdash A \lor \neg A}\;\lor R2}{\vdash \neg A, A \lor \neg A}\;\neg R}{\vdash A \lor \neg A, A \lor \neg A}\;\lor R3}{\vdash A \lor \neg A}\;\textit{contraction}$$

As is well known, this derivation essentially draws on the presence of multiple conclusions. However, there is only one active conclusion in each step of this derivation, and so each step can be given a natural deduction reading. Hence, multiple-conclusion rules both make a difference, and are justifiable on meaning-theoretic grounds.

Admittedly, the reading of the above derivation is nonuniform. In particular, notice that we first derive $A \vdash A \lor \neg A$ from $A \vdash A$ by $\lor R2$, which is construed in the natural deduction reading as a derivation of $A \lor \neg A$ from A, but then we immediately read that same sequent as an inference from A to every conclusion, where $A \lor \neg A$ is regarded merely as an unmade assumption. It is only because of this second reading that we may proceed in the derivation, applying $\neg R$ so

34. In the first direction, one may derive both $\lor R2$ and $\lor R3$ from $\lor R1$ with one application of weakening:

$$\cfrac{\cfrac{\Gamma \vdash A, \Delta}{\Gamma \vdash A, B, \Delta}\;\textit{weakening}}{\Gamma \vdash A \lor B, \Delta}\;\lor R1 \qquad \cfrac{\cfrac{\Gamma \vdash B, \Delta}{\Gamma \vdash A, B, \Delta}\;\textit{weakening}}{\Gamma \vdash A \lor B, \Delta}\;\lor R1$$

In the other direction, if contraction is available we reason as follows:

$$\cfrac{\cfrac{\cfrac{\Gamma \vdash A, B, \Delta}{\Gamma \vdash A \lor B, B, \Delta}\;\lor R2}{\Gamma \vdash A \lor B, A \lor B, \Delta}\;\lor R3}{\Gamma \vdash A \lor B, \Delta}\;\textit{contraction}$$

and if cut is available we reason as follows:

$$\cfrac{\cfrac{\Gamma, B \vdash B, \Delta}{\Gamma, B \vdash A \lor B, \Delta}\;\lor R3 \qquad \cfrac{\Gamma \vdash A, B, \Delta}{\Gamma \vdash A \lor B, B, \Delta}\;\lor R2}{\Gamma \vdash A \lor B, \Delta}\;\textit{cut}$$

35. These requirements are by no means trivial. See, e.g., Golan and Hlobil (2022) for logics in which the right weakening rule fails.

as to get $\vdash \neg A, A \vee \neg A$. Be that as it may, I already argued on independent grounds (toward the end of Section 4.2) that the natural deduction reading need not be uniform, and that such nonuniformity does not threaten its cogency. Consequently, multiple-conclusion rules can be accommodated in the inferential theory I outlined in this paper.

6. Conclusion

To be sure, my strategy of accommodating sequent rules with unintuitive features is limited in scope. Thus, the disjunction rule $\vee R1$ can be given a natural deduction reading only in the presence of structural rules that make it interderivable with $\vee R2$ and $\vee R3$. Moreover, my strategy only works for rules whose unintuitive features are restricted to the structure of contexts. It remains here an open question whether an extended version of this strategy can be applied to, say, hypersequent rules, or display calculi rules. That said, I believe that the strategy is sufficiently broad, and that the theoretical principles on which it rests—in particular, the extended context principle and the two readings of sequent rules—should be taken seriously by any inferentialist.

Acknowledgement

I presented materials from this paper at a conference at Concordia University and a colloquium at Haifa University. My thanks to the audience in both talks for fruitful discussions. Special thanks to Robert Brandom, Ulf Hlobil, Ryan Simonelli, and Shuhei Shimamura for extensive comments on earlier drafts of the paper. I would also like to thank two referees for this journal for insightful suggestions. Funding for this project was provided by the Israel Science Foundation, grant 1765/25.

References

Anscombe, Gertrude Elizabeth Margaret. (2000). *Intention*. Harvard University Press.

Anderson, Alan R., and Belnap, Nuel. (1975). *Entailment: The Logic of Relevance and Necessity*. Princeton University Press.

Belnap, Nuel. (1962). Tonk, plonk and plink. *Analysis*, 22(6), 130–134.

Brandom, Robert. (1983). Asserting. *Noûs*, 17(4), 637-650.

Brandom, Robert. (1994). *Making It Explicit: Reasoning, Representing, and Discursive Commitment*. Harvard University Press.

Brandom, Robert. (2000). *Articulating Reasons: An Introduction to Inferentialism*.

Harvard University Press.
Cobreros, Pablo, Égré, Paul, Ripley, David, and van Rooij, Robert. (2013). Reaching transparent truth. *Mind*, 122(488), 841–866.
Da Costa, Newton, & Alves, Elias. (1977). A semantical analysis of the calculi Cn. *Notre Dame Journal of Formal Logic*, 18(4), 621-630.
Davidson, Donald. (1977). Reality without reference. *Dialectica* 31(3-4), 247-258.
Dicher, Bogdan. (2020). Hopeful monsters: A note on multiple conclusions. *Erkenntnis*, 85(1), 77-98.
Dicher, Bogdan, and Paoli, Francesco. (2019). ST, LP and tolerant metainferences. Başkent C., & Ferguson T., editors, *Graham Priest on Dialetheism and Paraconsistency*, pp., 383–407. Springer.
Dicher, Bogdan, and Paoli, Francesco. (2021). The original sin of proof-theoretic semantics. *Synthese* 198(1), 615-640.
Dicher, Bogdan, and Paoli, Francesco. Logical Metainferentialism. *Ergo*, forthcoming.
Dummett, Michael. (1973). *Frege: Philosophy of language*. Duckworth.
Dummett, Michael. (1991). *The Logical Basis of Metaphysics*. Harvard University Press.
Dummett, Michael. (1993). *The Seas of Language*. Clarendon Press.
Dunn, J. Michael. (1974). A 'Gentzen' system for positive relevant implication. *Journal of Symbolic Logic*, 38, 356–357.
Evans, Gareth, and McDowell, John. (eds.) (1976), *Truth and Meaning: Essays in Semantics*, Oxford University Press.
Francez, Nissim. (2015) *Proof-Theoretic Semantics*. College Publishers.
Frege, Gottlob. (1953). *Foundations of Arithmetic*. Basil Blackwell.
Gentzen, Gerhard. (1964). Investigations into logical deduction. *American Philosophical Quarterly*, 1(4), 288-306.
Golan, Rea, & Hlobil, Ulf. (2022). Minimally nonstandard K3 and FDE. *Australasian Journal of Logic*, 19(5), 182–213.
Hacking, Ian. (1979) What is logic? *Journal of Philosophy*, 76(6): 285–319.
Hlobil, Ulf, and Brandom, Robert B. (2024). *Reasons for Logic, Logic for Reasons: Pragmatics, Semantics, and Conceptual Roles*. Routledge.
Kneale, William. (1956). The Province of Logic. In H. D. Lewis, editor, *Contemporary British Philosophy: Personal Statements*, Allen and Unwin, pp. 235-261.
Kremer, Michael. (1988). The philosophical significance of the sequent calculus. *Mind*, 47, 50–72.
Kripke, Saul. (1959). Distinguished Constituents, Semantical Analysis of Modal Logic, and the Problem of Entailment. *Journal of Symbolic Logic*, 24(4), 312-326.
McDowell, John. (1998). *Meaning, Knowledge, and Reality*. Harvard University Press.
Milne, Peter . (1986). Frege's context principle. *Mind*, 95(380), 491–495.
Negri, Sara, and von Plato, Jan . (2001). *Structural Proof Theory*. Cambridge Uni-

versity Press.
Peacocke, Christopher. (1976). What is a logical constant?. *The Journal of Philosophy*, 73(9), 221-240.
Peregrin, Jaroslav. (2014). *Inferentialism: Why Rules Matter*. Springer.
Prawitz, Dag. (1965). *Natural Deduction: A Proof Theoretical Study*, Almqvist and Wiksell.
Priest, Graham. (2005). *Doubt Truth to Be a Liar*. Oxford University Press.
Priest, Graham. (2006). *In Contradiction*. Oxford University Press.
Prior, Arthur. (1960). The runabout inference ticket. *Analysis*, 21(2), 38–9.
Restall, Greg. (2005). Multiple conclusions. In Hájek, P., Valdés-Villanueva, L., and Westerståhl, D., editors, *Logic, Methodology, and Philosophy of Science: Proceedings of the Twelfth International Congress*, pp. 189–205. College Publications.
Restall, Greg. (2020). Speech acts & the question for a natural account of classical proof. Article in progress, available online at https://consequently.org/writing/speech-acts-for-classical-proofs/
Ripley, Ellie. (2013). Paradoxes and failures of cut. *Australasian Journal of Philosophy*, 91, 139–164.
Ryle, Gilbert. (2009). *Collected Essays: 1929-1968* (Volume 2). Routledge.
Schroeder-Heister, Peter. (2012) The categorical and the hypothetical: A critique of some fundamental assumptions of standard semantics. *Synthese,* 187, 925–942.
Sellars, Wilfrid . (1953). Inference and meaning. *Mind*, 62(247): 313–338.
Shoesmith, D., and Smiley, Timothy. (1978). *Multiple-Conclusion Logic*. Cambridge University Press.
Steinberger, Florian. (2011). Why conclusions should remain single. *Journal of Philosophical Logic* 40, 333-355.
Tennant, Neil. 1997. *The Taming of the True*. Clarendon Press.
von Plato, Jan. (2003). Translations from natural deduction to sequent calculus. *Mathematical Logic Quarterly*, 49(5), 435-443.
Wansing, Heinrich. (1994). Sequent calculi for normal modal propositional logics. *Journal of Logic and Computation*, 4(2), 125-142.