Harmony and logical inferentialism

Florian Steinberger
Hughes Hall
University of Cambridge

This dissertation is submitted for the degree of
Doctor of Philosophy

This dissertation is the result of my own work and includes nothing which is the outcome of work done in collaboration except where specifically indicated in the text.

This dissertation does not exceed the word limit set by the Degree Committee of the Faculty of Philosophy.

To Hana

Harmony and logical inferentialism

Florian Steinberger

Abstract

My thesis is an attempt to supply answers to what I take to be the three central questions facing inferentialism about the logical constants (which I call *logical inferentialism*). What are the assumptions about meaning that underpin logical inferentialism? What is the correct formulation of the principle of harmony? And finally: What follows from logical inferentialism? Accordingly, the dissertation falls into three parts.

I begin by laying out the fundamental meaning-theoretic principles that underpin logical inferentialism: it is use-theoretic; it subscribes to the two-aspect theory of meaning; as the name makes plain, it is inferentialist in approach; and it is committed to a weak form of molecularism. Having spelled out its founding assumptions, I defend inferentialism against the charge that, contrary to its defining motto, the meanings of the logical operators are not fully determined by the rules of inference they obey, but also in part by structural assumptions.

The second part offers a comprehensive treatment of the notion of harmony. After giving an analysis of the notion, I review and criticize existing accounts of harmony in the literature. In particular, I show that Michael Dummett's and Stephen Read's theories are unsatisfactory and I present a counterexample to Neil Tennant's principle of harmony. I then advance my own version of harmony, which not only avoids the difficulties that plagued the accounts mentioned, but also boasts additional advantageous features.

In the final part I examine the consequences of these results for so-called proof-theoretic arguments. Such arguments purport to show that the principle of harmony supports broadly intuitionistic revisions of our logic. I argue that, given our inferentialist commitments, a defence of classical logic based on the adoption of multiple-conclusion sequent calculi is misguided. Multiple-conclusion systems, I submit, are illegitimate from an inferentialist point of view. Moreover, I defend the principle of separability against realist attacks.

Acknowledgements

I am deeply obliged to my supervisor, Michael Potter, for his insight, patience and good council in matters dissertation-related and beyond. Also, I am most grateful to my 'shadow' supervisor, Arif Ahmed, who went far beyond his call of duty in giving me extremely discerning comments on extensive sections of various drafts of this dissertation. Special thanks are due to Neil Tennant, who was extraordinarily generous with his time in answering my questions and whose helpful comments on parts of this work led to significant improvements. I also benefited from the input of a number of people who kindly provided detailed comments and criticisms of different sections and drafts of this dissertation. They are Fiorien Bonthuis, Salvatore Florio, Julien Murzi, Peter Smith and Nick Tosh. Needless to say, I bear the sole responsibility for all remaining errors. I also would like to thank my examiners Alex Oliver and Stephen Read for a helpful *viva voce*.

Ole Hjortland's paper 'Proof-theoretic harmony and structural assumptions', to which I responded at the First Cambridge Graduate Conference for the Philosophy of Logic and Mathematics in January 2008, inspired sections 2.6–2.8 (although my conclusions are rather different from his). Volker Halbach and Jeffrey Ketland patiently and competently answered my not-so-competent questions regarding theories of truth, which were pertinent in section 4.7. I want to thank all of them very warmly.

I am also grateful to the Faculty of Philosophy at the University of Cambridge and in particular the very active community of logicians and philosophers of mathematics for providing a stimulating working environment. Very special thanks are due to the members of the RSC for teaching me a great deal about philosophy and for their friendship. The work on my Ph.D. was funded by the AHRC and by the Gates Cambridge Trust, whose support I gratefully acknowledge. For being instrumental in securing my funding and just for being a positive presence in the Faculty, I want to thank Lesley Lancaster.

Finally, the debts I cannot hope to repay. I don't know what I would have done without the advice, help and encouragement—philosophical, editorial and otherwise I received from Yoon Choi. And then there is my family: Sanne, Henry, Chris, Tami, Martin, Nomi and Sami. Without their trust and loving support none of this would have been possible.

Contents

Part I

Introduction

Chapter 1

Introduction

1.1 The really big picture

There are two approaches towards an account of meaning. According to the first, the *truth-conditional* view, meaning is intimately linked to the notion of truth. The meaning of a sentence is explained by the conditions under which a statement made by means of it is true. The meaning of a subsentential expression is then explained by the contribution it makes to those truth-conditions. This explanatory strategy is centred on language-world relations: language's ability to refer to the world is at its core. The other approach explains a sentence's meaning in terms of the use we make of it. Call it the *use-theoretic* view. In order to understand a sentence we must ask under what circumstances its utterance is appropriate. Giving an account of meaning amounts to spelling out the regularities and conventions that govern the use we make of language, the regularities and conventions that we internalize and learn to conform to in acquiring a language. If the former approach prioritizes the explanation of how language latches on to the world, the use-theoretic view emphasizes language-speaker relations.

Both approaches undeniably have a strong *prima facie* appeal. Each one highlights an essential feature of language: language is a means of talking about the world, on the one hand; and a social communicative practice on the other. Nevertheless, there is a tangible tension between these two conceptions. It surfaces in the form of the ambiguity involved in the notion of a sentence's 'correctness'. Plainly, a sentence's correctness may be appraised along two competing axes: a sentence may be correct in virtue of 'telling it how it is', 'being true to the facts', etc.; or it may be correct in the sense that its utterance is appropriate under the circumstances. The speaker has fulfilled his obligations vis-à-vis his audience and so is entitled to perform a speech act with such and such a content. The result is that the utterance is a licit 'move in a language game'.[1] The first axis is concerned with getting the relation between language and non-linguistic reality right; the second depends on whether a speaker's linguistic performance meets a certain normative standard set by his linguistic community. Contrary to the truth-conditional view, in which neither the speaker nor the community of which he is a member plays a role of any significance, the use-theoretic view carries an indelible human stain. Only such cues as are readily recognizable at least in principle, by us, the fellow speakers—e.g. linguistic, perceptual or other—can be taken to be criterial of correct use. Truth-conditions may or may not obtain, quite independently of anybody's capacity to determine which.

Nothing I will say here will, I am afraid, directly contribute to alleviating the tension between truth-conditional and use-theoretic approaches. However, it will surely be a step in the right direction to see how far each type of account can be developed, what obstacles they face and what the theoretical and explanatory costs and benefits are of endorsing them. My contribution to such an enterprise is on the far less developed side of use theories of meaning. In this dissertation I examine the prospects for a use theory for the meanings of the logical constants.[2]

An account of the meanings of the logical constants can be viewed as a component of a larger theory of meaning—we could think of it as a first step towards such an Herculean enterprise. And it would be a particularly important step. After all, the logical constants, or at any rate the operators of propositional logic, afford the means for constructing complex sentences out of simpler ones. Therefore any theory of meaning had better be able to account for them: as Michael Dummett puts it, 'if a theory cannot deal successfully with complex statements, it matters little how good it is at explaining the atomic ones' (Auxier and Hahn 2007, p. 482). From the perspective of the optimistic use-oriented meaning-theorist, the fragment of language

[1] We may ignore possible difficulties arising from different types of propriety. It may be assumed for present purposes that there is a reasonably clear sense in which the utterance of a sentence can be *semantically* (as opposed to other types of propriety) appropriate. We are painting in broad brushstrokes at this stage.

[2] I will for present purposes assume the logical constants to be the standard bunch $\wedge, \vee, \neg, \supset, \forall, \exists$ and $=$. Identity is not treated here, however for an interesting inferentialist take on identity see Read (2004).

comprised of the logical operators is thus an all-important milestone. The questions we face can be stated as follows:

- What form does a use-theoretic account of the meanings of the logical constants take?
- What are its founding assumptions?

However, even if the grand use-theoretic project is found to be unworkable for language in general, it might nevertheless turn out to be our best bet for the logical fragment.[3] Such an account of the meanings of the logical operators would be of great interest in its own right. Moreover, we may reasonably hope that an understanding of how meanings are conferred upon our logical vocabulary would yield valuable insights into the very nature of logic. We may hope to make headway in particular on two of the pivotal questions in the philosophy of logic, namely:

- What distinguishes logical expressions from other types of expressions?
- Which logic (if any) is the right logic?

1.2 Plan

In accordance with my initial questions, the purpose of part one is to sketch the use-theoretic account of meaning that is the object of my inquiry and to make explicit the premises on which it relies. I begin by canvassing a theory of language in general terms. In a second step I take a detailed look at how such a theory may be applied to the fragment of the logical expressions. It is shown that potential problems for the theory do not arise in the restricted context of the logical operators; if there is any region of language that is best accommodated in a use-theoretic framework of the type proposed, it is that of the logical expressions.

Although the account of meaning offered here owes much to Dummett's writings on the matter, my project is not one of exegesis. I borrow also from other sources, especially Robert Brandom and Neil Tennant, in an effort to outline an attractive

[3]This would be so in particular if a systematic theory of meaning of either kind for the whole of language proved to be unattainable. It is less clear whether there could be hybrid accounts that would treat certain fragments truth-conditionally, others use-theoretically. Given how far we are from any form of unified theory, one cannot even speculate as to the legitimacy of such theories.

use theory of meaning. This is the theory I call *logical inferentialism.* I then go on to lay out the key assumptions underlying such a broadly Dummettian approach. The central assumptions can be summarized as follows:

- Meaning is explained in terms of use;
- Use is systematized in the two-aspect model of meaning;
- Inferentialism: The two aspects are given by inferential relations;
- (Minimal) molecularism: At least the logical constants form a semantically autonomous region of language.

While I indicate what motivates some of these assumptions, no attempt is made here to mount full-scale defences of any of them—they are, after all, assumptions. My goal in chapter 2, rather, is to lay bare the pillars on which the edifice of logical inferentialism rests and to show how these underlying assumptions hang together. In the final section 2.6 I consider a potential caveat to logical inferentialism in the form of structural assumptions. Do not global properties of our deducibility relation partially shape the meanings of the logical constants? But if this is so, what becomes of the inferentialist slogan that the rules of inference alone determine the meanings of the constants? I argue (in section 2.8) that while one's choice of structural assumptions affects the output of the system (i.e. the set of theorems provable within it), the meanings of the logical constants remain invariant. The threat to the inferentialist thesis is thus averted.

Having made these posits explicit, I turn to the requirement of harmony in part two. Harmony is presented as both a meaning-theoretic principle that goes hand in hand with molecularism and, in its logic-specific form, a possible defining feature of logic itself. I begin by analysing Nuel Belnap's notion of harmony as conservativeness (3.2). I then go on to briefly consider the view that takes harmony to be a criterion of logicality (3.3). An intuitive notion of harmony, general harmony, is then sketched on the basis of the notion of an equilibrium between the grounds that warrant the introduction of a logical constant and the consequences we are allowed to draw (3.4). The three most prominent existing accounts of harmony, Dummett's (4.1), Tennant's (6) and Stephen Read's (7), are reviewed and demonstrated to be unsatisfactory against the background of our notion of general harmony (8). Drawing in particular on Read's work, I then present a novel version of the principle of harmony. My proposal has the merit not only of giving a precise content to the notion of harmony, but of supplying an effective procedure which enables us, for any permissible rule of inference, to determine its harmoniously matched counterpart. Plus it enables us to test whether or not a given pair of inference rules presented to us does or does not satisfy harmony.[4]

Part three investigates some of the consequences of our meaning-theoretic assumptions, harmony in particular. Anti-realists have put forth proof-theoretic arguments to the effect that classical logic fails to meet meaning-theoretic constraints. Assuming for the purposes of the argument that proof-theoretic arguments succeed in the context of natural deduction systems, I ask to what extent they are sensitive to the proof-theoretic framework in which they are framed. In particular, what are the prospects for a classicist rebuttal to the anti-realist based on a shift to sequent calculi? I begin by showing that the notion of harmony is not only applicable to the sequent setting (10.1), but that standard classical sequent systems are indeed harmonious (10.2). I go on to counter possible objections turning on the idea that sequent systems, because of their meta-theoretical status, prove unsuitable for inferentialist purposes (11). As I show, it is not the sequent format as such that is at fault (11.2–11.3); rather, the problem from an inferentialist perspective resides in the admission of multiple-conclusion systems (11.4). In order to square such systems with ordinary argumentative practice we must read multiple-conclusion proofs disjunctively. However, this undermines the inferentialist project because it presupposes an antecedent knowledge of at least one of the connectives the meaning of which the system was meant to convey (12.1–12.4). An alternative interpretation of multiple-conclusion consequence relations based on the notion of denial is considered, but found equally wanting (12.5). Finally, building on Peter Milne's diagnosis according to which the opposition between classicists and intuitionists is to be understood in terms of their rejection or acceptance, respectively of the principle of separability (13.1–13.2)—the principle that rules of inference may treat of no more than one logical constant at a time—it is shown that separability is a consequence of our meaning-theoretic assumptions, in particular of the two-aspects model of meaning (13.3–13.4).

[4]Note that this last point is not trivial: explaining simply what harmony is is not necessarily enough to establish for any given logical constant presented to us that it cannot be expressed by harmonious inference rules.

Chapter 2

Meaning-theoretic assumptions

2.1 The two-aspect model of meaning

The first step towards formulating a use theory of meaning is to get clearer about what is meant by 'use'. What aspects of use are we after? How can they be integrated into a theoretical framework? And what form ought such a framework take? Now, the ways in which we use most expressions of ordinary language are of course an extremely complex affair. Understanding a sentence will typically involve a capacity to discriminate a certain range of perceptual stimuli as well as an appreciation of inferential connections that link the words occurring within it to other expressions and of appropriate behavioural responses to linguistic acts performed by means of the sentence in question. Is there any hope of moulding these ingredients of meaning into a systematic account?[1]

In our quest for a framework enabling us to systematize the conventions and rules that are determinative of linguistic use we do well to turn to Dummett. Dummett famously proposed what we might call the *two-aspect model of meaning.*

> Crudely expressed, there are always two aspects of the use of a given form of sentence: the conditions under which an utterance of that sentence is appropriate, which include, in the case of an assertoric sentence, what counts as an acceptable ground for asserting it; and the consequences of an utterance of it, which comprise both what the speaker commits himself to by the utterance and the appropriate response on the part of

[1]Note that we are only concerned with *assertoric* use here.

> the hearer, including, in the case of assertion, what he is entitled to infer from it if he accepts it (Dummett 1973, p. 396).

The former category of features of use corresponds to a broadly verificationist conception of meaning—let us therefore call these features of use *V-principles*.[2] The V-principles associated with an expression specify the circumstances under which an assertion of a statement involving that expression is appropriate. The second category of principles represents what we can do with a sentence in virtue of having asserted it. An emphasis on this feature of use is the hallmark of a pragmatist theory of meaning—call such features *P-principles*.

Dummett argues that neither aspect of meaning exhausts the content of a sentence. Suppose I know simply when a particular form of words is appropriately employed, but that I am wholly ignorant of what follows from my having recognized it as true, what I thereby commit myself to, and what it entitles me to say or do. In such a case it would be plain wrong to ascribe to me a suitable grasp of the meaning of the words. Similarly, to understand a sentence I cannot dispense with an appreciation of its assertibility conditions no matter how complete my knowledge of the consequences of asserting it may be. Both of these fundamental features must enter into a use-theoretic account of meaning. Verificationist and pragmatist accounts of meaning should be understood accordingly. Neither doctrine should be

> taken to be so naive as to involve overlooking the fact that there are many other features of the [assertoric and other] use of sentences than that one singled out as being that in which its meaning consists (Dummett 1973, p. 457).

It is true that both doctrines claim that a particular feature—assertibility conditions in the case of the verificationist; practical consequences in the case of the pragmatist—is constitutive of meaning. Yet this is not to be understood as a claim that meaning is reducible to either of these core features taken on its own. The claim is rather that the remaining aspects of a sentence's use can be derived from

[2]Clearly more would have to be said, given that, as Dummett puts it,

> there is a multiplicity within this category, according as we are concerned with when an assertion is *conclusively* established, or with what merely warrants its being made, though defeasibly (Dummett 1991, p. 247).

These subtleties need not detain us here.

these core features. As we will see, at least in the context of the logical constants we may assume that the two features—assertion-conditions and consequences of asserting—are 'alternative' in the sense that

> either is sufficient to determine the meaning of a sentence uniquely; but they are complementary in that both are needed to give an account of the practice of speaking a language (Dummett 1983, p. 142).

In the general case it will be far from easy to state V-principles and P-principles with any degree of precision. The principles regulating the use of most non-logical expressions are likely to involve messy inputs and outputs, many of which may be extra-linguistic—perceptual cues among V-principles and actions and other behavioural consequences as P-principles. It is a wide open question whether a systematic use theory can be formulated along the lines of this bipartite model (or indeed any other model) for the whole of language. For the purposes of the present project, however, all we need to assume is that such an account can be given in the case of the logical constants. And indeed a little reflection shows that this assumption is highly plausible. The meanings of logical expressions are readily thought of as being determined solely by intra-linguistic relations. Logic's characteristic property of *formality* guarantees that the meanings of the logical constants do not depend on the particular contents of the sentences involved, but only (to some extent) on the sentences' logical forms.[3] It is in virtue of this feature that the two-aspects model provides a suitable framework for logic. The V-principles and the P-principles associated with logical expressions can be represented schematically in the familiar form of natural deduction-style introduction and elimination rules.

The familiar picture, then, is this. The meaning of every logical expression is governed by two sets of rules, introduction rules and elimination rules. If \$ is a logical constant (suppose, for simplicity, that \$ is a binary operator), the introduction rules associated with \$ state (perhaps partially) the meaning of that constant by specifying the circumstances in which an assertion of a sentence in which \$ is the main connective is warranted. They inform us what obligations we have to meet in order to be entitled to assert the compound sentence $\$(A, B)$. The role of elimination rules, in turn, is to state which logical inferences we may legitimately draw in virtue

[3] At least in one of the senses of 'formality'. MacFarlane (2000), though still unpublished, has become the standard reference here.

of having endorsed the sentence containing $ in a dominant position. The two-aspects model could not be more neatly represented than in the natural deduction setting.

In a way, of course, this should not come as a surprise. For what is the two-aspect model of meaning but an extrapolation to the whole of language of the insights underlying Gerhard Gentzen's natural deduction systems? In an oft-quoted passage, Gentzen recognizes the meaning-theoretic relevance of introduction and elimination rules.

> The introductions represent, as it were, the 'definitions' of the symbols concerned, and the eliminations are no more, in the final analysis, than the consequences of these definitions (Gentzen 1969a, p. 80).[4]

V-principles and P-principles are simply generalized introduction and elimination rules. We thus arrive at the familiar idea that the meanings of the logical constants are fixed by the rules of inference they obey. Why not simply say that then? Whence the need for such elaborate scene-setting?

One reason is simply to underline the meaning-theoretic significance of the project. The other important motivation, however, is to provide sufficient meaning-theoretic background to explain what it means to say that rules of inference can be determinative of meaning in the first place. We shall turn to this issue in a moment. First, however, we must get clearer on the role of the notion of inference in our use-theoretic account of meaning.

2.2 Inferentialism

We have noted that in the case of non-logical vocabulary, V-principles and P-principles may take a variety of forms. Assertibility conditions may partially consist in a range of perceptual input. Among the meaning-constitutive consequences there may be certain behavioural responses. Nevertheless, a purely mechanical differential response to stimuli could not count as mastery of an expression or of the concept it expresses. A parrot that is taught to screech 'red' in the presence of red things

[4]Note that Gentzen attributes meaning-theoretic priority to introduction rules. There is no need to commit ourselves on this point here.

cannot thereby be said to be a semantically competent user of 'red'.[5] What is involved in a genuine grasp of the meaning of an expression, therefore, is an awareness of at least some of the core inferential relations the expression entertains with other expressions. To know the predicate 'aardvark' a speaker has to know a number of basic inferences *towards* and *away from* sentences involving 'aardvark': that if Helmut is an aardvark, Helmut is an animal, for example. In the Sellarsian tradition such inferential links have been dubbed *material* inferences to distinguish them from formal inferences.

Inferentialism is the view that meaning is at least partially determined by the inferential links among an expression's inputs and outputs. In other words, inference is among the meaning-constitutive core uses of any expression. One can in principle be an inferentialist about certain regions of language and not others, and one can subscribe to different degrees of inferentialism, depending on one's stance concerning the extent to which inference determines meaning. Much depends, of course, on the strength of one's inferentialism. And a healthy dose of it seems in order for many types of expressions, especially those of a more theoretical type—that is, expressions that are removed from the observational outskirts of the Quinian web, separated by a rich body of theory. By contrast, inferentialism becomes more problematic especially in its stronger variants when applied to expressions that are more intimately hooked up with the world because of their content or indeed because of their grammatical category. (Proper names are a case in point.)

How does the situation present itself in the case we are interested in, the logical constants? In their case, inferentialism is a much less controversial position to hold. The V- and P-principles of logical operators are comprised of *nothing but* inferential links. There are no other features of use that enter into an account of the meaning of logical constants except for the (formal) inferential connections that link them to the associated premises and to the conclusions: the introductions state the 'inferentially sufficient conditions' for asserting a sentence in which the constant in question is the principal operator; the eliminations spell out the 'inferentially necessary consequences of the employment' of such a sentence (Brandom 1994, p. 117).

It is important to note the relevance of this meaning-theoretic background for

[5]This point is made by Dummett (1973, p. 453) and it is at the centre of Brandom's inferentialist project, see e.g. Brandom (2000, p. 17).

our proof-theoretic approach. Our guiding principle that the meanings of the logical constants are given by the rules of inference they obey is compatible in principle with a radical type of conventionalism according to which we could fix the meanings of the logical operators at will. Simply devise any rule you like and you will thereby fix the meanings of the symbols contained in those rules. Nothing prevents us from laying down new rules, of course. The mistake, however, resides in the idea that any formal game incorporating what appear to be inference rules will confer meanings on its logical symbols. Contrary to the Carnapian *dictum* that in logic there are no morals, for the use-theoretic approach I am investigating here, only such deductive systems fit the bill as can be seen to be answerable to the use we put our logical vocabulary to. Only if the deductive system yields an adequate representation of the core features of our (or at least of a possible) deductive inferential practice will a system qualify by our meaning-theoretic standards.[6]

That is not to say, of course, that a system of logic amounts to a mere description of the actual use of the logical expressions. A use-theoretic account of the meanings of the logical constants remains firmly committed to the Fregean doctrine of the normativity of logic. Although our rules of inference must bear witness to our practices, they simultaneously exert normative force. In this respect rules of inference are comparable to a grammar: beginning as a record of how we do in fact do things (or could do things), the grammar attains normative force telling us how we ought to do things.[7]

At the same time our rules of inference are constrained by the general principles that shape our account of meaning. These principles are another source of normativity that weighs on our theory. Importantly, these principles also constitute a corrective for our practice. This has been a point central to Dummett's thoughts about meaning:

> the mere fact that it is established affords no ground for assuming a linguistic practice free from defect. With that, we perceive that our linguistic practice is no more sacrosanct, no more certain to achieve the ends at which it is aimed, no more immune to criticism or proposals of revision, than our social, political, or economic practice (Dummett 1991,

[6] I shall sidestep the avowedly very difficult question of when a formalism can be said to be faithful enough to the practice it represents.

[7] These issues would merit much more careful attention than we are able to give them here. We can do little more than gesture of what will be a future project.

p. 215).

As we will discuss in detail in part two, a given practice may be criticizable not only for harbouring inconsistencies but also for failing to meet the requirement of harmony: the V-principles and P-principles governing the expressions of a language must not be fixed independently of each other, but rather must be suitably matched. However, the demand for harmony is itself a consequence of more fundamental meaning-theoretic principles. Central among them is the assumption of molecularism, to which we must now turn.

2.3 Molecularism

One of our central assumptions is that a Dummettian molecular conception of meaning is correct. Indeed, as we will see, a very restricted form of molecularism will do for our purposes. What exactly does this commitment entail? As Dummett stresses, a molecular theory of meaning is compositional.[8] This, in and of itself, does not say very much. *Any* sensible theory of meaning must explain how the meaning of a sentence depends on its composition—the subsentential expressions of which it is composed and the order in which these are put together. The point on which the molecularist and the holist (as we will understand the position) differ is in their conception of what knowledge is required to understand a sentence. To be sure, to understand the (distinct, non-synonymous) sentences A and B we have to know what sentences they are, i.e. we have to know which expressions they contain and we have to be familiar with the modes of phrase- and sentence-construction that gave rise to them. But in general we have to know more than that. We have to know what the expressions occurring in the sentences in question mean. The holist's claim is that what has to be known in all cases is the same: the entire language. By contrast, the molecularist contends that the knowledge required will in general differ from case to case.

Now, this is not to say that the meanings of sentences can be known in isolation. The molecularist can agree that language is interconnected. But the question is *how* interconnected it is. Contrary to the Augustinian atomistic picture famously attacked by Wittgenstein at the beginning of his *Philosophical investigations*, the

[8] In what follows we are drawing heavily on Dummett's work, in particular (1991, ch. 10).

molecularist holds that in general, understanding an expression will presuppose acquaintance with *some* other expressions and phrases. By contrast, the holist holds that nothing short of the knowledge of the entire language is sufficient to understand fully any expression or phrase in that language.

To illustrate how the molecularist position occupies the middle ground between atomism and holism, Dummett invokes Gareth Evans's 'generality constraint':

> One could not understand the sentence 'That cow is lying down' unless one could also understand other sentences such as 'This cow is standing up', 'That horse is lying down', and so on (Dummett 1991, p. 222).

The molecularist thus allows for what we might call *semantic clusters*. A semantic cluster is a set of expressions or phrases the understanding of any of which necessitates an understanding of all the others in the set. Following Dummett we may call the relation that obtains between two expressions when a grasp of the meaning of one presupposes a grasp of the meaning of the other the relation of *dependence*. Semantic clusters consist of expressions that are mutually dependent. Examples are groups of contrary predicates like colour words or phrases like 'mother of', 'father of', 'child of'. Famously, according to Quine the expressions 'analytic', 'necessary', 'synonymous' also form a cluster.

What the molecularist opposes is the holist notion that all of these clusters collapse into one all-encompassing master cluster, language as a whole. Molecularism can thus be expressed as the claim that the relation of dependence among expressions in a language be by and large asymmetric. Despite the existence of semantic clusters, we can think of the semantic dependencies between expressions as approximating a partial ordering. What prevents them from being a partial ordering properly so-called is the existence of clusters. Nevertheless these clusters, when considered as a single *relatum*, are themselves part of a well-founded dependence relation.

A full treatment of molecularism would call for a much more detailed account of the nature of this dependence relation. Luckily we may dispense with it here.[9] It seems that connected to the problem of characterizing the relation are some of the most fundamental questions concerning the very nature of concepts and our grasp of them. Some expressions will denote sharply delineated concepts for which we can specify necessary and sufficient conditions which themselves are concepts

[9] For a more detailed characterization of the dependence relation see Tennant (1987, p. 56).

or Boolean combinations thereof expressed by other words or phrases. In other cases, the dependence might rather be such that the criteria for the application of a given concept might depend on the applicability of a weighted majority of a set of other concepts. On some even weaker criterialist accounts, concepts may be identified through even looser relations of 'family resemblances'. Fortunately, such fundamental questions may justly, I think, be regarded as being orthogonal not just to my particular project but more generally to the question of the tenability of a molecularist theory of meaning.[10] Indeed, the molecularist need not even be a strict conceptual foundationalist. He is free to hold that the grasp of any one concept always presupposes the grasp of (a limited number of) other concepts; that it is impossible to possess a single concept in isolation. This amounts to the admission that in tracing our dependence relation back to its sources it is possible that our search always culminates in clusters rather than single concepts.[11] So long as these clusters are not too extensive and do not induce circularity at a global level, there is no reason why they should be anathema to the molecularist's position.

2.4 The context principle

Compositionality taken on its own is purely a bottom-up affair: in order to understand a complex expression one must understand its parts and the mode of their combination. But how do we attain an understanding of the constituents of complex expressions? What is the source of the interconnectedness of language discussed in the previous section? Gottlob Frege famously held that one should not ask after the meaning of a word in isolation; *it is only in the context of a sentence that a word has meaning.* Understanding a word thus consists in the capacity to understand and use sentences containing it. But which sentences? Clearly, the molecularist must deny that the knowledge of the meaning of a word presupposes an understanding of literally *all* the possible sentences involving it. Understanding every sentence containing the said expression would presumably require understanding every constituent of every such sentence. By the context principle, this would in

[10] Cf. Tennant (ibid., p. 54).

[11] I am setting aside questions concerning possible divergences between conceptual and linguistic dependence relations. A discussion of this point would lead us too far afield here. Should the reader's particular stance towards this question be in tension with what we have said here, he may simply replace all talk of concepts and the relations between them with talk of linguistic expressions and the relations obtaining between *them*.

turn presuppose an understanding of all the sentences involving *those* expressions, and so on. There would be no stopping short of knowing all of language.

The molecularist wishing to retain the context principle must therefore hold that a speaker's mastery of an expression's meaning requires only an understanding of a certain *representative range* of sentences containing it, a range significantly less extensive than the complete set of sentences involving the said expression. For every expression there is therefore a range of sentences a grasp of which is constitutive of the understanding of the expression in question. All other sentences involving the expression in question then presuppose an antecedent grasp of that expression, whereas a grasp of the sentences within the representative range must precede an understanding of it.

The range is in part determined by the expression's distinctive semantic features, that is, those of its features that tell us what systematic compositional contribution the expression makes to the meaning of a sentence of which it is a constituent. Such features tell us what sort of expression we are dealing with, whether it is a sentence- or a term-forming operation, for example. This aspect is largely schematic.

Another aspect is that which informs us about the content of the expression. This typically involves having the ability to understand other simple sentences relating the expression to perceptual cues or to other expressions already in our repertoire. Take, for instance, Dummett's example of the word 'fragile'. An understanding of 'fragile' requires an understanding of its logico-grammatical role as well as the ability to comprehend

> its use for some simple predications like 'That plate is fragile'; an understanding of such a sentence as 'I'm afraid that I forgot that it was fragile' *builds on* and *requires* an antecedent understanding of the word 'fragile' but is not a condition of understanding it (Dummett 1991, p. 224).

Clearly, for many types of expressions it will not be easy to delineate a representative range of sentences. This does not mean, however, that even in such unfavourable cases we cannot have a rough idea as to what is necessarily involved in understanding an expression, and which more complex sentences containing it can be understood derivatively by an appeal to compositionality. All this shows is that an explicit account for such fragments will be hard to come by. There is no reason to think that the mere existence of fuzzy cases counts in favour of the holist.

Molecularism is, I think, an attractive position. Nevertheless, logical inferentialism does not hinge on so strong an assumption. In fact, all that the inferentialist requires is that the notion of a representative class of sentences can be clearly made out in the case of *one* particular group of expressions particularly dear to our hearts, the logical constants. And indeed, the logical vocabulary once again lends itself particularly well to the molecularist story. The representative range of sentences associated with a logical constant is readily made out.

> The understanding of a logical constant *consists* in the ability to understand any sentence in which it is the principal operator: the understanding of a sentence in which it occurs otherwise than as the principal operator *depends on*, but does not go to constitute, an understanding of the constant (ibid.).

What Dummett says here is not quite accurate. Of course we do not have to get our head around *every* sentence of which a given constant is the principal operator to grasp its meaning. All that is required is a schematic grasp of the bulk of such sentences. An understanding of conjunction involves only the understanding that for any sentences A and B, whenever we are entitled to the assertion of A and B separately we are also entitled to the assertion of ⌜A and B⌝.[12] This guarantees that for any sentences A and B we are in a position to understand, we shall also understand their conjunction, which is rather different from Dummett's claim. But once we have successfully mastered this much, there is nothing more we can learn about the meaning of '$\wedge$' by considering sentences of the form $(A \vee B) \supset (C \wedge D)$ or 'Jim knows that $A \wedge B$, but doesn't care'.

As Tennant rightly points out, this presupposes a 'mastery (albeit implicit) of the general concept of assertion' (Tennant 1987, p. 63).[13] This means that we must also be capable of discerning what Tennant has called the 'first layer of logical form': the ability to 'see through' the (possible) logical complexity of the sub-sentences and to single out the *dominant* operator of the sentence in question (Tennant 1997, p.

[12] Clearly, however, understanding any *particular* conjunction does presuppose an understanding of the two conjuncts as well as an understanding of 'and'.

[13] Note that this does not mean that every instance of, say, $\wedge$-introduction or elimination manipulates sentences to which assertoric force attaches. Inference rules can equally apply in the contexts of supposition and possibly also denial. Nonetheless, there appear to be good reasons for thinking that the other two types of linguistic acts are parasitic on assertions.

313). Once we have grasped this, we dispose of all the prerequisites essential to an understanding of the (logically relevant) meaning of conjunction.[14]

Our minimal molecularist commitment amounts to the claim that there is at least one fragment of language that is self-contained, the logical fragment. It is self-contained in the sense that for each logical constant we can clearly discriminate between those sentences containing it whose grasp is constitutive of our understanding of the expression and those that are not.[15] Furthermore, the representative range of sentence schemata will be very limited. Although a more thoroughgoing molecularism is arguably justified, for current purposes this is all we need.

But why should the logical fragment in particular enjoy such meaning-theoretic autonomy? To see what goes wrong if we deny the logical constants this status, it is instructive to consider the example of quantum logic. It has been suggested by some that many of the most puzzling quantum phenomena could be dealt with by abandoning certain principles of classical logic. In particular, empirical findings appear to be at odds with the law of distributivity $A \wedge (B \vee C) \vdash (A \wedge B) \vee (A \wedge C)$.[16] On this 'more Quinian-than-Quine' approach we could—at least in principle—do away with some of the 'wacky' conclusions quantum-mechanics appears to force upon us at the expense of rejecting certain familiar laws of logic.

To be sure, we can *say* that 'or' now ought to be understood as having precisely the meaning needed to fit the new use dictated by our scientific theory (or, rather, our interpretation thereof). And this is just the holist's reaction faced with such a

[14]Note that we are presupposing that there is a distinction to be made between those aspects of meaning that are logically relevant and those that are not. In Frege's parlance logically irrelevant aspects are part of an expression's *tone* as opposed to its *sense*. The latter but not the former determines the *formal* inferences the connective participates in. For instance, the connoted temporal progression of 'and' in 'Rachel got married and had a child' is part of the expression's tone and so is negligible for the purposes of logical analysis. Likewise, to take Frege's favourite example, 'and' and 'but' may in certain contexts induce appreciable differences in meaning (compare 'Judy is a charming woman and she is married' and 'Judy is a charming woman but she is married'), but this difference is again one of tone, not of sense: the additional inferences the tone may make available to us are part of the expression's conversational implicature and is not logically significant. What we are presupposing, then, is that the logical expressions of our formal language ($\wedge, \vee, \neg, \supset, \forall, \exists$) succeed in capturing at least the logically relevant aspects of the meanings of the corresponding expressions in natural language (leaving it open whether or not this exhausts the meanings of 'and', 'or', etc.).

[15]Note that there could conceivably be other very weak forms of molecularism that would take not the logical constants but some other set or small superset containing the logical constants as the only self-contained fragment alongside a giant network containing the rest of language. Our claim to minimality should therefore not be understood so as to rule out such alternative scenarios.

[16]See for instance Putnam (1975a, p. 184).

scenario. Since the meaning of these expressions (like that of any other expression) is given by the role they play in our overall linguistic practice, it does not matter what changes we wish to subject our practice to: the meanings of our constants will be modified accordingly. They will take on whatever meaning we care to give them.

But what meaning is that? The holist's story fails to address the problem. The problem, as Dummett points out, is that we shall lose any sense of what it is we are saying after such a shift in meaning occurs.

> Can someone, who may never have heard of quantum logic, recognize, or work out, how this individual [i.e. a user of quantum logic] does understand it [disjunction]? He attaches a weaker meaning to ⌜A or B⌝ than we do since, for him, it does not have all the implications that it has for us. What meaning *does* he attach to it, then? If he tells you, 'Either deforestation is halted within the next ten years, or human life will be extinct before the end of the next century', and you are disposed to think he knows what he is talking about, how alarmed should you be? Just *what* is he asserting? (Dummett 1991, p. 206)

It is therefore not enough to say that the meaning of the disjunction operator so constrained just is its contribution to some configuration of our total linguistic practice. We would not, on that basis, be able to partake in rational deliberation and debate. We would not, for instance, as in Dummett's example, be capable of assessing the cogency of arguments that are presented to us. In order to participate, to use one of Brandom's favourite phrases, in 'the game of giving and asking for reasons' we must have a secure (even if tacit) grasp of the deductive properties of logical expressions. And this we would not have even if quantum logic were imposed upon us by a federal decree and every member of the linguistic community became accustomed to using 'or' in accordance with quantum logic. For even if we had successfully internalized the appropriate rules, we would still be unable to fathom the possibly disruptive effects the altered rule would have on our overall linguistic practice. From the molecularist point of view, this is not surprising. For what reason do we have to expect that, for any mould carved out by a possible linguistic practice, there must be a stable meaning for our logical operators to fit it? A snappier way of restating our conclusion would be to say that *logic is not empirical*.[17]

[17]We will pursue this train of thought in section 3.3.

Dummett's argument delivers strong grounds for adopting our minimal molecularist assumption, namely that *if* we assume that all-out holism is mistaken, and we assume, hence, that there is at least one self-sufficient language fragment, *then* logic must form such a meaning-theoretic island. It also constitutes a powerful argument against holism *tout court*. Nevertheless, an adequate treatment would require us to say more about both issues than can be accommodated by this introduction. Therefore, at least for official purposes, both our rejection of holism and the conditional claim should be regarded as having the status of posits. However, the discussion of the previous paragraphs should at the very least have indicated a promising line of argument in favour of our assumptions.

Let us explicitly state our assumption for future reference:

> **Assumption of minimal molecularism**: In what follows we will assume that there is at least one semantically autonomous language fragment. Among these fragments is the set of the logical constants.

2.5 Immediate and mediate inferential transitions

There is a further assumption we need to make. What we are after is in a sense already brought out by Dummett's example of how the context principle applies to a word like 'fragile'. Though it proved difficult to determine the exact extension of the representative range for such an expression, it was still obvious that the range did not include sentences of any structural complexity. An understanding of 'fragile' does not require us to know how it behaves in belief contexts or in logically complex sentences.

The same thought applies to the logical constants. (Although, again, it can be given a crisper formulation in this case.) Thus far we have been concerned exclusively with the question of what range of sentences a speaker must master in order to understand a constant. We found that the representative range in question is the set of sentences in which the operator has a dominant occurrence. Moreover, we said that for an understanding of the constant it is sufficient to have a schematic grasp of the licit inferences involving sentences in the representative range. Understanding a constant is knowing the core inferences that are determinative of its meaning. But what are the core inferences?

Our choice is constrained by the commitments we have already incurred. The two-aspect model of meaning allows only inferences that take the form of introduction and elimination rules to enter into consideration. Also, as molecularists, we must deny that it is necessary to grasp the *entire* net of deductive inferences to which a constant contributes. Having mastered a manageable set of core inferences, our understanding of $\wedge$, say, would not be improved if we were to be told that some barely surveyable formula is a remote consequence of a conjunction $A \wedge B$. Our understanding cannot depend on our knowing *every* one of the infinitely many deductive consequences of $A \wedge B$. Clearly, such a demand would be exorbitant. In the following we may thus safely assume that the meaning of any constant can be displayed by a finite number of rules that recursively encapsulate that constant's meaning-constitutive core use. For each logical constant there must be a finite set of basic rules that are sufficient to generate all of the other deductive links to and from sentences containing the operator in question in a dominant position. Tennant dubs this requirement the principle of analytic systematicity.

> It reflects the fact that we have to codify our own logical competence finitely and schematically in order to exhibit the very possibility of our understanding what it is upon which validity of argument in general rests (Tennant 1997, p. 309).[18]

The crucial point here is that on this view it will be incumbent upon us to justify all other logical laws on the basis of these elementary ones; we have to show that any logical law is already, as it were, contained within the basic laws in embryonic form. The task of justification is a hallmark of molecularism; it does not arise for the holist, nor could it.

The question thus arises which are the salient inferential connections from a meaning-theoretic point of view. Here it will be useful to distinguish between direct and indirect inferential transitions. Clearly there are a number of ways to establish a statement of the form $A \wedge B$. Most straightforwardly, it can be inferred *directly* once we have recognized its immediate constituents A and B as true. A direct (or canonical) way of establishing a statement is thus

> one which proceeds in accordance with the composition of the sentence by means of which it is expressed (Dummett 1991, p. 229).

[18] As we will see, this is consonant with our assumption of minimal molecularity in the previous section and its consequence, the principle of autonomy discussed in section 3.3.

All the standard natural deduction introduction rules proceed in this way: they inform us about what a speaker has to do in order to be entitled to assert the more complex sentence involving the logical constant in question. For any n-ary logical constant \$, the introduction rule leading to a statement of the form $\$(A_1, \ldots, A_n)$ will specify a list of i $(1 \leq i \leq n)$ constituents of the sentence in question that are sufficient conditions for its assertion. Inferences in accordance with inference rules of this type are what we call immediate inferential transitions.[19]

However, $A \wedge B$ may also be established *indirectly*. For instance, in classical logic, $A \wedge B$ can be seen to follow from $\neg(\neg A \vee \neg B)$. Here we are presented with a mediate inferential transition. The inferential move from $\neg(\neg A \vee \neg B)$ to $A \wedge B$ is not basic; it requires filling in so as to expose the underlying chain of basic inferential moves. The same distinction between mediate and immediate inferential transitions applies to elimination rules.[20]

Of course, in intuitionistic logic $\neg(\neg A \vee \neg B) \vdash A \wedge B$ does not hold and thus cannot be unpacked into more basic inferential transitions. The question therefore arises whether it is permissible in such situations simply to add the corresponding inference rule to our repertoire, thereby simply *treating* it as an immediate inferential transition. On the face of it, the rule

$$\wedge\text{-I}\ \frac{\neg(\neg A \vee \neg B)}{A \wedge B}$$

distinguishes itself in a number of ways from Gentzen's natural deduction rules. Most notably, it differs from our customary rules in that the logical complexity of the premise exceeds that of the conclusion.[21] Should we allow for rules of this type?

The answer is that we should not. Compositionality requires that the premises featuring in introduction rules should be of lower complexity than their conclusions.

[19]We are assuming here that, in the case of quantifier introduction rules, the substitution instance of a sentence occurring within the scope of a quantifier is a subsentence of the quantificational sentence taken as a whole; i.e. that $A[t/x]$, which is the result of substituting t for all free occurrences of x in A, is a subsentence of $QxA(x)$, where Q is a variable-binding first-order quantifier. Note also that not every immediate inferential transition constitutes a permissible inference rule according to our definition. The question which inference rules are permissible will occupy us at various stages in the course of this dissertation. See especially section 8.5 and chapter 13.

[20]This does not hold true in the same way of the standard $\vee$- and $\exists$-elimination rules. Nevertheless, the immediate constituents of the major premises do also play a significant role as assumptions initiating the subproofs in their respective rules. The same applies to the so-called 'generalized elimination rules' we will encounter in chapter 7.

[21]Throughout, by 'logical complexity' I will simply mean the number of logical constants occurring in a sentence.

We may thus adopt the principle that it must be possible to grasp the meaning of any logically complex sentence in terms of its logically *less complex* subsentences. Following Dummett we shall call this the *complexity condition*:[22]

> the minimal demand we should make on an introduction rule [...] is that its form be such as to guarantee that, in any application of it, the conclusion will be of higher logical complexity than any of the premises and than any discharged hypothesis (Dummett 1991, p. 258).[23]

We may thus assume, at least for the time being, that only inference rules expressing immediate inferential transitions are instrumental in determining the meanings of the constants. And it is understood that only rules satisfying the complexity condition can express such transitions.[24]

2.6 The role of structural assumptions

So far we have represented logical inferentialism as holding that the meanings of the logical constants are given by the rules of inference that govern them. Tacit in our presentation was the assumption that the meanings of the constants are *exhaustively* determined by the rules they obey. More specifically, as we have seen in the previous section, it is the direct rules of inference—the rules that express immediate inferential relations—that we are after. And the direct rules are taken to be operational rules; that is, rules that act directly on a specific connective. That being said, surely we must acknowledge that operational rules do not function in a vacuum, but always against the backdrop of certain structural assumptions concerning the relation of logical consequence of which the rules are only partially constitutive. Depending on what structural properties we ascribe to our relation of logical consequence, we end up with different logical systems with different sets of theorems. Indeed, the

[22]We will discuss this issue in more depth in section 13.3. For the time being we may endorse the complexity condition as a plausible assumption.

[23]The complexity condition of course carries over straightforwardly to elimination rules: the major premise must be logically more complex than any of the minor premises including possibly discharged assumptions occurring therein. In cases where there are no minor premises as in the case of the elimination rules for conjunction and the universal quantifier, the complexity condition demands that the conclusion be of lower complexity than the major premise.

[24]Note that our definition does not rule out the possibility of interdependencies between the meanings of two or more logical operators. However, it is argued in chapter 13 that we need to restrict our notion of meaning-giving inference rules to those that treat one operator at a time.

exploration of this simple fact has given rise to the field of substructural logics. As is well known, a number of intrinsically interesting systems of logic fall into the class of substructural logics. They can be viewed as sharing the same operational rules, differing only in the structural assumptions they impose. Alan Anderson and Nuel Belnap's systems of relevant logic can illuminatingly be viewed from this angle, as can Jean-Yves Girard's systems of linear logic and (in part) Tennant's systems of relevant logic, to name but a few examples.

To give two simple illustrations of how different structural assumptions can lead to distinct systems of logic with distinct sets of theorems, consider the following derivations:

$$\supset\text{-I, 1}\;\dfrac{\supset\text{-I}\;\dfrac{[A]^1}{B \supset A}}{A \supset (B \supset A)}$$

$$\supset\text{-I, 2}\;\dfrac{\supset\text{-I, 1}\;\dfrac{\supset\text{-E}\;\dfrac{\supset\text{-E}\;\dfrac{[A \supset (A \supset B)]^2 \qquad [A]^1}{A \supset B} \qquad [A]^1}{B}}{A \supset B}}{(A \supset (A \supset B)) \supset (A \supset B)}$$

The first of the two proofs incorporates an instance of a vacuous discharge (introducing a formula different from the one we started off with). Allowing for such a discharge policy is tantamount to adopting the structural rules of thinning (or weakening). The second proof discharges two tokens of the formula A simultaneously, which corresponds to an application of the structural rule of contraction. Restricting our attention momentarily to the rules governing the conditional, we find that depending on which discharge policy (and hence structural rules) we adopt, we can generate four different systems:

- Allow thinning and contraction $\Longrightarrow$ intuitionistic implicational fragment
- Drop thinning, allow contraction $\Longrightarrow$ relevant implicational fragment (in the sense of Anderson and Belnap's system **R**)
- Allow thinning, drop contraction $\Longrightarrow$ affine implicational fragment
- Drop thinning, drop contraction $\Longrightarrow$ linear implicational fragment

This raises an important question for the logical inferentialist: Are the meanings of the logical operators sensitive to our choice of structural assumptions? If so, this would suggest that the meanings of the logical constants are not governed solely by the corresponding inference rules, but also by the global properties of the system in which they operate.

Let us restrict our attention to the intuitionistic and the relevant implicational fragment. In the standard natural deduction system structural assumptions are built into the introduction rule for conditionals in the form of discharge policies. We can make these assumptions explicit in the form of separate rules by reformulating our systems in a sequent setting. Note that the systems in question are still *natural deduction* systems with elimination rules as well as introduction rules, as opposed to sequent calculi properly so called, which are characterized by introduction rules operating on formulas on both sides of the sequent sign.[25] In the sequent formulation of natural deduction the structural rules of thinning and contraction are excised from within the operational rules and explicitly stated.

$$\text{weakening}\ \frac{\Gamma : C}{\Gamma, A : C}$$

$$\text{contraction}\ \frac{\Gamma, A, A : C}{\Gamma, A : C}$$

We will be taking the sequent sign ':' to relate a set of formulas on the left with a formula on the right. In part three we will also consider multiple-succedent calculi where the right-hand side *relatum* is also a set. Consequently, we may dispense with the structural rule of contraction. In the sequent setting the first of the above derivations takes the following form:

$$\supset\text{-I}\ \cfrac{\supset\text{-I}\ \cfrac{\text{weakening}\ \cfrac{A : A}{A, B : A}}{A : B \supset A}}{: A \supset (B \supset A)}$$

Although the standard natural deduction system and its sequent cousin are co-extensive, the difference in mode of presentation is not entirely innocent. On the standard natural deduction presentation, where the structural rules are integral to

[25]The sequent formulation of natural deduction systems, like the original systems of natural deduction and sequent calculus, are due to Gentzen (see Gentzen (1969b, p. 152) for the former; Gentzen (1969a, p. 77) for the latter).

certain operational rules (in the present case, the rules for the conditional), we have the impression that the meanings of the two connectives are distinct. For given that the structural assumption of weakening manifests itself in the form of discharge policies on this presentation, the intuitionistic conditional $\supset$ and the relevantist connective $\rightarrow$ are governed by different introduction rules. Hence, judging by the logical inferentialist's standards, it appears that the two connectives must be distinct: there are circumstances under which the intuitionistic conditional but not the relevantist's conditional may be asserted.

The sequent formulations of the respective systems offer a different perspective. Viewed from the angle of the sequent presentation the introduction rules appear to be identical. Consequently, for the logical inferentialist it looks as though we are dealing with the same connective placed in two different deductive contexts: one in which the consequence relation is taken to tolerate thinning, another in which it is not. This way of putting things suggests that there is a stable meaning conferred on the constants by their operational rules. Where the two systems diverge, rather, is in the properties they ascribe to their respective deducibility relations. The sequent presentation suggests that these two issues—the meanings of the logical constants as given by their operator-specific introduction and elimination rules and the global properties attributed to the consequence relation by the structural rules adopted—are separate.

But which of these modes of presentation is the correct one? Do we find ourselves in a standoff between two equally acceptable modes of presentation? Or are there any grounds for choosing between them? It would be a mistake, I think, to present things as if the question were simply a matter of taste. The sequent presentation does get things right. Although the standard natural deduction presentation is arguably more faithful to ordinary practice, it clouds the distinction between *local* operator-specific rules of inference and *global* properties of the system expressed by structural rules. The fact that there are various ingenious ways of redistributing structural information across operational rules does not tell against the robustness of the distinction. Repackaging such global features as parts of operational rules or as operational rules in their own right simply masks their true status.

This might raise the question of why only certain rules and not others carry structural information in standard natural deduction formulations. If the information encoded by structural rules is indeed information about the system as a whole and if natural deduction really does illicitly disguise such information by integrating it into local rules, should not all local rules be equally affected? And, therefore, does not the fact that only certain rules are affected indicate that the information in question pertains to only some constants (namely, those defined by the affected rules) and not others? The answer to both of these questions is no. After all, the intimate link between discharge policies in natural deduction in the standard format and the explicitly stated structural rules of the natural deduction systems in their sequent formulation is no coincidence. Discharge policies, by definition, affect only so-called *improper* inference rules; that is, inference rules like $\supset$-I, $\vee$-E and $\exists$-E whose statement involves the mention of subproofs as well as the possible discharge of hypotheses.[26] Rules of this type can be distinguished from *proper* inference rules whose characteristic feature is to mark immediate inferential transitions that do not call upon other (sub-)derivations from assumptions as their premises, but rather operate directly on statements. The customary rules $\wedge$-I, $\wedge$-E, $\vee$-I, $\forall$-E, etc., are all examples of proper inference rules.[27] Standard natural deduction systems, unlike sequent-style presentations, do not explicitly represent the deducibility relation in the object language. Consequently, structural rules cannot be explicitly stated in the form of manipulations on sequents. The only way, therefore, in which we can fix the structural properties of the system is by injecting them locally, as it were, where the deducing is going on: in improper rules that take subdeductions as premises. Though undeniably compact, such rules misrepresent the roles of (operational) inference rules. By contrast, the sequent formulation with its explicit distinction between operational and structural rules seems to capture the natural divide between local operator-specific rules and global properties of the relation of deducibility—properties that a deducibility relation may be taken to have quite independently of the presence of any particular (or indeed of any) logical vocabulary. A structural feature simply is one that *can* be represented as a global property of a system; an operational feature cannot be so represented.

[26] The terminology is borrowed from Dag Prawitz (1965, p. 23).

[27] See section 11.3 for a more detailed discussion of this distinction.

2.7 *Ex falso* as a structural rule

There remains an inference rule that sits uneasily within the classification of rules into operational and structural rules: the *ex falso sequitur quodlibet* rule (*ex falso* or *EFQ* for short).

$$\text{EFQ}\ \frac{\bot}{A}$$

In this section I criticize the standard accounts of *ex falso*. Based on Tennant's work concerning the nature of $\bot$ and its role in negation, I propose an unorthodox reading of *ex falso*, classifying it as a *structural rule*.

There is surprisingly little agreement about the exact status of *EFQ* in the literature. Gentzen groups the *ex falso* rule alongside the other operational rules (1969, p. 77). Prawitz follows this practice. He treats $\bot$ as a '0-place sentential operation' (Prawitz 1978, p. 38), one 'for which there is no canonical proof' (Prawitz 1977, p. 26). *EFQ* is then understood as the elimination rule for $\bot$. Others, including Dummett, present $\bot$ as an elimination rule for $\neg$ (e.g. Dummett 1991, p. 291). I contend that both of these approaches are mistaken.

Beginning with the latter approach, it suffices to point out that we already have an elimination rule for negation (one that is, as we will see, harmonious with respect to the standard introduction rule).

$$\neg\text{-E}\ \frac{A \qquad \neg A}{\bot}$$

Contrary to ¬-E, which corresponds to Gentzen's original rule, Dummett's use of *ex falso* as an elimination rule turns out not to be a harmonious match for the standard introduction rule (as we will see). I therefore see no reason to prefer *ex falso* over ¬-E. Nevertheless, it is not hard to see why *ex falso* is frequently assimilated to negation. The structural rule of thinning on the left is easily recognisable as a global rule acting on the deducibility relation rather than on any particular constant. By contrast *ex falso* tells us that anything follows from a contradiction. And contradictions are thought to arise only in the presence of negation. Hence, *ex falso* and negation appear to be intimately linked. This idea seems to find additional support from the fact that *ex falso* is generally justified by appeal to the law of disjunctive syllogism. If *ex falso* holds because disjunctive syllogism holds, must there not be a close tie with the meanings of the operators $\neg$ and $\vee$?

Let us begin with the latter point. While *EFQ* can of course be justified by appeal to disjunctive syllogism, it should not be overlooked that disjunctive syllogism is itself only provable in the presence of *EFQ*.[28] The rules for $\vee$ and $\neg$ alone are insufficient to bestow a meaning upon them that would enable us to prove disjunctive syllogism. Moreover, seeing that it is preferable to have a specification of the meanings of the logical constants that respects separability, adopting disjunctive syllogism as an axiom or as a primitive rule of inference is not an option. Hence, it is *ex falso* that is naturally treated as primitive, not the principle of disjunctive syllogism. Therefore, rather than thinking of *ex falso* as the product of the meanings of disjunction and negation, we must treat disjunctive syllogism as a product of the *ex falso* rule. This suggests that the law of disjunctive syllogism is not a consequence only of the meanings of the constants involved, but also hinges on one's stance concerning the general principle expressed by *EFQ*.

This leaves us with the first point. Undeniably, we take it that contradictions occur when we have derived two contradictories, a statement and its negation. But there is no reason in principle why *ex falso* could not also play a role in systems devoid of a negation operator. A 'contradiction' in this sense is not necessarily the result of asserting two contradictories. For example, we could conceive of systems in which $A, B \vdash \bot$ would hold, where A and B are contraries. So long as our language contains antonyms like hot/cold, soft/hard, weak/strong, etc. or contraries like red/green, triangle/circle, there will be cases where various sentences fail to be jointly assertible. This we register with the sign '$\bot$'. In all such cases, we have a choice of whether or not to adopt *EFQ*. The fact that there is such a choice, that the decision for or against *EFQ* is not predetermined by the meanings of the logical constants in the system—for a system with *EFQ* need not contain any constants at all—shows that *EFQ* is indeed a structural rule. As such it expresses a general policy regarding our consequence relation, namely that we allow any statement whatsoever to follow from a contradiction.

Now, let us return to the first approach, the approach favoured by Prawitz, according to which *ex falso* is to be understood as an elimination rule for $\bot$. As we

[28] At least this is so in standard systems of natural deduction. Tennant's systems of relevant intuitionistic logic (**IR**) and relevant classical logic (**CR**) are exceptions. They are distinguished by their liberalized $\vee$-elimination rule, which enables one to derive $A \vee B, \neg A \vdash B$ in the absence of *EFQ*. The downside of Tennant's systems is that in order to prevent fallacies of relevance from creeping back in, Tennant is forced to sacrifice the full transitivity of the deducibility relation. Section 5.3 offers a slightly more elaborate discussion of Tennant's rule.

have seen, in natural deduction systems $\bot$ marks the occurrence of a contradiction.

$$\neg\text{-E}\ \frac{A \qquad \neg A}{\bot}$$

The thought then is that '$\bot$' in fact denotes a particular contradiction or a generic 'absurdity constant' for which *EFQ* is taken to be the elimination rule.

$$\neg\text{-E}\ \frac{A \qquad \neg A}{\mathit{EFQ}\ \dfrac{\bot}{B}}$$

As we have seen, this is Prawitz's position: $\bot$ has *EFQ* as its elimination rule, but no corresponding introduction rule (see also Read 2000, p. 139). This raises a number of questions. Is $\bot$ a particular constant? If so, which one? And how could such a constant be discourse-independent, as it must be if it is to be a genuine logical (sentential) constant? If $\bot$ is not taken to be a particular constant it must be a generic constant. But what exactly is a generic constant and how could such a constant be embedded in complex sentences.

In his (2004a) Tennant demonstrates that these questions admit of no satisfactory answers. Rather than treating $\bot$ as a sentential constant, we do better to understand it as a 'structural punctuation marker within deductions' (ibid., p. 1). As such its job is to register the occurrence of an absurdity in the course of a derivation. The upshot is that $\bot$ is not embeddable in compound sentences and hence that negation has to be taken as primitive, rather than as a mere shorthand for $A \supset \bot$.[29]

But what are we to make of the idea that *ex falso* functions as an elimination rule? Must not $\bot$ be a logical constant of sorts given that it stands in need of an elimination rule? The real question is why we should consider *ex falso* as an elimination rule in the first place. Prawitz's reason for treating it so appears to be this. According to Gentzen the role of an elimination rule is just to unpack the content of a statement made by means of a sentence containing the constant in question in a dominant position—the content that the corresponding introduction rule had invested in the statement. To make things more concrete consider a binary logical constant $ governed by a single introduction rule, and suppose $ is associated with the following elimination rule.

$$\$\text{-E}\ \frac{\$(A, B)}{C}$$

[29] For a lucid discussion of these points I refer the reader to (Tennant 2004a, p. 5).

For Prawitz, Gentzen's *dictum* then translates into the requirement that there be a procedure that transforms any canonical (direct) proof of $\$(A, B)$ into a canonical proof of C. If this requirement is fulfilled, the constant is in good working order and the elimination rule is justified relative to the introduction rule. Prawitz seeks to extend this mode of justification to *ex falso*. And in a sense this is easily done. Suppose Π is a canonical proof of $\bot$. But there is no canonical proof of $\bot$; $\bot$ has no introduction rule. The procedure needed to transform Π into a proof of A for any A just consists of 'appending A to (the non-existent) Π'. Hence, the result of appending A to Π is a canonical proof of A. In other words, the argument

> Π is a canonical proof of $\bot$; therefore, the proof obtained by appending A to Π is a canonical proof of A

is valid, though vacuously so.[30] *EFQ* is thus 'justified' under the mantle of being the elimination rule for $\bot$, which is understood as a kind of limiting case of a logical constant. This account is clearly driven by a concern for theoretic systematicity, rather than by philosophical motivations. The resulting artificiality of Prawitz's move seems to me to be undeniable, especially in the light of our aforementioned doubts about the nature of $\bot$ as a 'sentential constant'.

I submit that these difficulties can be overcome simply by treating $\bot$ as a punctuation marker, as Tennant suggests. Of course this means that *EFQ* is left without a logical constant which it serves to eliminate. But on the view I am proposing no such constant is needed because *EFQ* is not an operational rule at all. Its role is best accounted for by according it the status of a *structural rule*. It simply is the structural rule that tells us that any sentence whatsoever follows from a contradiction. Such a license is *not* specific to any logical constants, but amounts to a blanket policy. Our reclassification of *EFQ* thus fits neatly with our characterization of structural rules as global rules that assign properties to our deducibility relation.

The structural character of *ex falso* is brought out more clearly still when we consider the analogous role played by structural of weakening on right in the sequent calculus (not natural deduction presented in a sequent framework, but the sequent calculus properly so called). In the sequent calculus an inconsistency is denoted by the sequent without succedent:

$$A, \neg A :$$

[30] Cf. Tennant 2004a, p. 21.

'Adopting *EFQ*' then simply amounts to adopting the structural rule of weakening *on the right*.

$$\text{weakening right}\ \frac{A, \neg A :}{A, \neg A : B}$$

The empty space to the right of the sequent sign reminds us that we have reached a logical dead-end in the same way in which we use '$\bot$' in natural deduction systems to express the fact that we have entangled ourselves in a contradiction. Given that we have reached such a dead-end, an instance of the structural rule of weakening on the right enables us to infer any statement whatsoever from the same premises.[31] I want to suggest that the sequent calculus represents things accurately. It reveals the true nature of *EFQ*. In the natural deduction setting this is obscured by our use of $\bot$. As Tennant puts it,

> $\bot$ was introduced by modern logicians in the natural deduction context in very much the way that the ancient Hindus introduced the symbol 0 into arithmetic. Rather than writing nothing, we indicate that it's nothing that we intend, by writing something in particular, which is to stand for the nothing that we intend (Tennant 2004a, p. 8).

What was introduced as a mere placeholder, a punctuation mark within a derivation, has taken on a life of its own. The result has been confusion about the role of $\bot$. It has led some to believe that it is a logically significant expression that stands in need of inference rules to determine its meaning. It has led others to confuse the structural property of admitting weakening on the right with the inferential properties of the negation operator. And it has led many to see the *ex falso* rule as occupying an 'anomalous position [...] inside the scheme of introduction/elimination rules', tarnishing the neat symmetry of the intuitionistic system (Weir 1986, p. 461). I submit that all of these difficulties can be overcome simply by applying the lesson from the sequent calculus to the natural deduction setting. If we, bucking convention, simply denote the existence of a contradiction by an empty space in natural deduction,

$$\neg\text{-E}\ \frac{A \qquad \neg A}{}$$

[31] Standard intuitionistic sequent calculi require that the rule of weakening on the right be restricted to empty succedents.

we thereby expose *ex falso* for the structural rule it is; we reveal its intimate relation to the structural rule of thinning on the right.

$$\text{weakening right}\ \dfrac{\neg\text{-E}\ \dfrac{A \qquad \neg A}{}}{B}$$

This shows that the standard natural deduction formulations of intuitionistic logic are misleading. They present *ex falso* as an operational rule that appears to introduce an awkward asymmetry into the system. In fact, however, this is just a superficial feature of the standard system. The move from minimal logic to intuitionistic logic, represented by the adoption of the *ex falso* rule, turns on structural assumptions rather than on a shift in the meanings of the logical constants. Likewise, relevant logic can tidily be characterized as the substructural logic obtained by dropping the structural rule of weakening on the left from minimal logic.

2.8 Operational meaning

Thus far we have argued that the illusion created by the natural deduction formulation—that the intuitionist's $\supset$ and the relevantist's $\rightarrow$ are governed by different *operational* rules—is just that, an illusion. What separates the intuitionist from the proponent of relevant logic are their structural assumptions. And it is the sequent presentation of natural deduction systems that best brings to light the division between operational rules and structural rules. Moreover, we have argued that, contrary to standard presentations, the *ex falso* rule should be thought of as a structural rule closely related to the rule of weakening on the right in sequent presentations.

But the inferentialist is not yet home free. Indeed, we have not so much as addressed the central question. The question, let us recall, was whether the meanings of the constant are in part determined by our structural assumptions. Our question is of considerable importance for the inferentialist. For if it turns out that structural rules really do have a say in determining the meanings of the logical constants, this would undermine the logical inferentialist's foremost slogan that the meanings of the constants are *fully* determined by the rules of inference that govern them.

Now, we have seen that even if we are careful to separate operational rules from structural rules, one's choice of structural rules has an impact on the output of the system. As the sequent presentation makes plain, the only thing that separates the intuitionistic fragment $\{\supset\}$ from the relevantist fragment $\{\rightarrow\}$ is that the former but not the latter allows for the structural rule of (non-trivial) thinning on the left. It is this difference that accounts for the fact that $\vdash_I A \supset (B \supset A)$ but $\nvdash_R A \rightarrow (B \rightarrow A)$. But do these structurally-induced differences in the outputs of the systems (i.e. in the sets of sequents derivable in the respective systems) betray genuine differences *in the meanings of the connectives*? In particular, to take our above example again, does the difference in extension between **I** and **R** bear witness to a difference in the meanings of the conditionals in the systems, as our notational conventions suggest?

Let us suppose they do. Since we have argued for a sharp distinction between the contribution made by operational rules to the meanings of the constants and the global contribution made by structural rules towards the system's overall deductive power, we might expect there to be a parallel distinction to be made at the level of an operator's meaning. That is, we might be able to distinguish between a constant's 'operational meaning', which is determined by its operational rules, and its 'structural meaning' given by the structural rules in the system.[32] On this view, the operational meaning would constitute the invariant core of the operator's meaning. But the basic operational ingredient would need to be supplemented by a structural component with which it jointly determines the overall meaning of the constant. According to this position, the intuitionist's $\supset$ and the relevantist's $\rightarrow$ would share the same operational meaning; the difference in meaning between the two would reside solely in the structural aspects of their meanings.

Accrediting structural rules with meaning-theoretic power thus leads us to postulate an overall meaning for each constant, a meaning that is the joint product of the operational and the structural contributions. The impact of the structural rules is then immediately felt. Adopt different structural assumptions and you modify the overall meanings of all the constants in the system. But what basis is there for postulating such an overall meaning in the first place? The underlying thought, presumably, is that any change in the set of derivable sequents involving a constant must be indicative of a change in the way we use that constant. Hence, the meaning of a constant is thought to be answerable to every aspect of its global inferential behaviour. To take our example again, the fact that we have $\vdash_I A \supset (B \supset A)$ but $\nvdash_R A \rightarrow (B \rightarrow A)$ is not a theorem in **R** shows that $\supset$ and $\rightarrow$ differ in their

[32]Francesco Paoli proposes such a distinction in his (2003).

overall meaning. But this presupposes that the *entire* web of inferential connections in which the constant in question partakes is determinative of its *overall* meaning.

But why should we assume such a thing? The inferentialist account we are proposing does not commit us to taking every possible inferential connection into account. That would be a kind of holism, not an all-out holism encompassing the whole of the language, but a kind of inner-logical holism nonetheless.[33] In order to understand fully the meaning of a constant a speaker would have to master all the deductive inferences involving it, each and every one of the consequences, however remote, of any sentence containing it. We rejected this position as indefensible above (see section 2.5). Rather, we argued that knowledge of a constant's meaning consists in a grasp of its core inferential use—the immediate inferential transitions expressed by the operational rules. It is they that tell us everything we need to know about the correct use and hence the meaning of the constant. Given that the operational rules provide us with all the constant-specific information we need, why should we suppose that there is any more to the meaning of a constant?

The upshot is that operational meaning is all we need. It is certainly true that the operational meanings of a constant, the conditional, say, are insufficient to determine all the theorems and implications involving it; there is a gap between the purely operational input into the system and the system's overall output. It is also correct to say that what closes the gap are structural assumptions. Operational meanings and structural assumptions do fully determine the set of sequents provable in a system. Where I part ways with Paoli is in the idea that the overall output of a system specific to a given constant (all the sequents containing $\supset$, say) should be considered part of the meaning of $\supset$, let alone that it should constitute its 'overall meaning'. And since there is no overall meaning of a constant, a constant has no structural meaning either.

There should not be anything surprising about this conclusion. What logically follows from what is not a matter of the meanings of the logical operators alone. Structural assumptions set the scene in that they encapsulate our assumptions concerning the global properties of our deducibility relation. The considerations that weigh on our choice of structural assumptions are rather different from those that shape the meanings of the logical constants. Our motivation for adopting such and such a structural rule may depend on our conception of what we take a valid argu-

[33] We will pursue this idea in section 13.3.

ment to be (e.g. on whether our conception of validity supports monotonicity and so on). Also, as in Tennant's case, our choice may be driven by a concern to maximise the epistemic gain we can reap from our arguments (see e.g. (Tennant 1997, p. 322)). These considerations are orthogonal to those that we might call upon in a dispute over the meaning of a logical constant. In the latter case our decision is going to turn on the correct use of the logical constant which in turn will hinge on our meaning-theoretic principles and ultimately on our conception of what an adequate account of meaning should look like. It follows that controversies surrounding the question which logic is the correct one can take either of these two forms. It can take the form of the debate between, say, advocates of intuitionistic logic and advocates of relevant logic, in which case it will be a debate over structural assumptions; or it can be like the debate between classicists and intuitionists, in which case it is the operational meaning of some of the logical constants that is being questioned.[34]

[34]We will be concerned with the dispute between intuitionists and classicists in part three of this dissertation.

Part II

Harmony

Chapter 3

Harmony: Its nature and purpose

3.1 Introduction

Having outlined the meaning-theoretic background to our investigation, we may now turn to the notion of harmony. Before asking how the notion of harmony is best characterized, we would do well to get clear about the purpose it is to serve. We shall begin, therefore, by surveying the ends for which it has been put forth. As a closer look reveals, harmony has been advanced in an attempt to solve a number of distinct (though often more or less interrelated) problems. It has, in various guises, been proposed as

1. a response to Arthur Prior's famous '`tonk` challenge';
2. a criterion of logicality, namely as the idea that all (and perhaps only) those expressions whose meaning can be exhaustively specified in terms of harmonious inference rules count as properly logical;
3. a component of a proof-theoretic justification of the laws of logic;
4. a *desideratum* in the dispute between revisionary anti-realists and realists defending classical logic;[1]
5. a 'precondition for the possibility of a compositional meaning-theory' (Dummett 1991, p. 247). For a language as a whole (or a fragment thereof) to

[1] Anti-realists press for broadly intuitionistic reforms of logic. On the basis of 2. and/or 3. they argue that because classical logic inevitably has to appeal to disharmonious principles of inference it strictly speaking ceases to be logic.

be in good working order, the principles governing the use of the expressions contained therein must be harmonious.

My ambition is not—nor could it be—to offer an exhaustive treatment of all of the above points. 1. will be discussed in section 3.2. I will very briefly address 2. in section 3.3. Questions relevant to 4. will be considered in part three. 5. was briefly touched upon in section 2.3, where we emphasized the interdependence between harmony and molecularism.

Let us begin, however, by looking at the challenge to logical inferentialism issued by Prior with his infamous spoof connective `tonk` and the role the notion of harmony has played in the writings of those who have tried to meet the challenge.

3.2 Harmony as a cure for tonkitis

Inferentialism, we had said, holds that the meanings of the logical constants are fixed by the rules of inference they obey. This does not mean, however, that we may lay down any rules of inference we please. Indeed, it is with this erroneous idea—that we could arbitrarily lay down logical laws and thereby stipulate meaningful logical constants into existence—that Arthur Prior famously takes issue in his well-known article 'A runabout inference ticket' (1960). Prior therein introduces a putative logical connective, the infamous `tonk`. `Tonk` is in effect a hybrid obeying one of the introduction rules ordinarily associated with disjunction

$$\frac{A}{A \text{ tonk } B}$$

and one of the elimination rules for conjunction

$$\frac{A \text{ tonk } B}{B}$$

A constant governed by these laws would immediately plunge any system into inconsistency, inasmuch as it would enable us to draw any conclusion from any premise. `Tonk` thus constitutes a counterexample to the *laissez-faire* inferentialism mentioned above. According to Prior, the moral of the story is that the possibility of freak constants like `tonk` proves the inferentialist approach as a whole to be fundamentally misguided.

While we may agree that `tonk` is an unacceptable constant and that the inferentialist has some explaining to do, it seems equally clear that not all logical inferentialism is *laissez-faire* inferentialism. A more accurate conclusion for Prior to draw would be that he has in fact shown that not just *any* set of inference rules will confer a coherent meaning on the logical constants involved; i.e. that *laissez-faire* inferentialism is wrong. But this does not show that the notion of meaning-conferring inference rules is *per se* absurd. All that it shows is that a set of logical laws must meet certain conditions if it is to determine the meanings of the logical expressions occurring within them. The challenge raised by Prior is thus that of formulating suitable constraints on sets of inference rules the satisfaction of which would ensure that the laws in question are fit to perform their semantic duty. It should be clear, moreover, that these constraints must form part of a general account of the meanings of the logical constants; mere *ad hoc*ery—however effective at blocking roguish connectives like `tonk`—will not do.

Belnap, in his (1962) reply to Prior, offers both a diagnosis of the problem and a proposal for its solution. *Au fond* what is objectionable about Prior's connective, according to Belnap, is that it perturbs our existing deductive practice. Indeed its effect on any existing practice is so deleterious that the entire system is reduced to triviality. What is needed, therefore, is a requirement to the effect that the introduction of a new logical operator not interfere with our established logical practice.

Belnap proposes to adapt for this purpose the notion of a conservative extension familiar from the study of formal theories. Let us briefly elaborate. Let T and T' be theories formulated in the languages $\mathcal{L}$ and $\mathcal{L}'$ respectively, such that $T \subseteq T'$ and $\mathcal{L} \subseteq \mathcal{L}'$. T' can then be said to extend T conservatively if T' allows us to prove no hitherto unprovable sentences expressible in the restricted language $\mathcal{L}$; any newly provable theorems must be sentences involving new vocabulary. In other words, for every sentence $A \in \mathcal{L}$, if $T' \vdash A$, then $T \vdash A$. Call this *theoretical conservativeness.*

How does this apply to systems of logic? Let us first get clear on the notions involved here. A *system* of logic is determined by the set of inference rules (operational and structural) it permits. Note that structural rules need not be explicitly stated as such, but may, as in the case of natural deduction systems, be incorporated into the rules of inference or their interaction, as is the case, for instance, with the discharge policies for rules involving $\supset$. On our definition, there may be any number of distinct yet extensionally equivalent systems. For example, there may be a number of different systems with different rules of inference that all still prove all classically valid sequents $\Gamma : A$. In the following we shall at times refer to individual systems, at times to equivalence classes of extensionally equivalent systems. Where the context leaves room for doubt, we shall be explicit about which is meant. A system comes with its associated language. When we speak of 'extensions of a system', what is meant is that one or more logical constants are added to the language and correspondingly that the associated pairs of (sets of) inference rules are added to the system.

Belnap's proposal, then, is that in the context of systems of logic, conservativeness should amount to the proviso that the introduction of one or more novel logical operators (by laying down logical laws that govern it/them) is legitimate only if it does not result in the derivability of new sequents involving operators drawn from the restricted language. Let us call this variant *systematic conservativeness*. More precisely, suppose we have a base system S and its extension S' with the associated respective languages $\mathcal{L}$ and $\mathcal{L}'$, where $S \subseteq S'$ and $\mathcal{L} \subseteq \mathcal{L}'$. Then Belnap's conservativeness constraint requires that for every $A \in \mathcal{L}$, $\Gamma \vdash_{S'} A$ only if $\Gamma \vdash_S A$.[2]

The difference between theoretic conservativeness and Belnap's revamped notion of systematic conservativeness is that in the former case the deductive machinery, the logical laws, remains untouched by the extension. In the case of formal theories, we are providing our relation of logical consequence with more fodder in the form of non-logical axioms, but crucially no new logical moves are added to the repertoire. In the case of systematic conservativeness, by contrast, the system of logic—the logical axioms and/or set of inference rules—is itself extended. The demand for conservativeness so understood thus limits the purely *logical* expressive power of the resulting system relative to the restricted system. Logical extensions of this kind can be thought of as parts of a series of consecutive steps towards the construction of a relation of logical consequence 'from below': we begin with a relation of logical consequence determined solely by a number of structural properties (e.g. transitivity, monotonicity, etc.) in order then to introduce logical constants progressively by way of the inference rules that govern them; at each step we take the closure of our relation of logical consequence under the newly available primitive inferences and

[2]In general A will also be provable from different sets of hypotheses involving new vocabulary in S'. This need not concern us so long as the constraint as we formulated it above is satisfied.

combinations thereof.[3]

It should be noted that Belnap's account is beset by some ambiguity. It may be that Belnap's original proposal should be understood as a weaker requirement limited to the special case where S *contains only structural assumptions*. On this view, what is required for a set of inference rules to determine a coherent meaning of a given constant is that their introduction should not impede on the antecedently given *structural* assumptions.[4] The question is whether the requirement of 'consistency with antecedent assumptions' (Belnap 1962, p. 131) for every newly introduced constant is to be understood as extending to other logical constants already in S, as in our formulation above, or whether, rather, Belnap places himself in a system S devoid of logical vocabulary. On the latter interpretation, Belnap's requirement says that, given the structural assumptions made on S, the newly introduced constant ought not generate a non-conservative extension of the remainder of language.[5]

For convenience we shall stick to our formulation.[6] However, it should be noted that the two formulations do not amount to the same thing. As we will see in more detail in section 4.6, the addition of classical negation to the implicational fragment produces a non-conservative extension. Classical negation can thus become problematic on our proposal, whereas it is perfectly compatible with Belnap's weaker requirement. How might we justify our stronger version of the conservativeness constraint? The underlying inferentialist intuition is that if the meaning of a constant is fully determined by its rules of inference, the introduction of further connectives should add nothing to that meaning. But then the addition of new logical vocabulary should not make available new theorems involving only the old vocabulary. This is because logical truths should only depend on the meanings of the constants occurring within them and the structural properties of the system. Therefore, if new logical theorems can be proven in the presence of novel vocabulary, this demonstrates that the meanings of the old vocabulary *have* been affected. And this implies that the meanings of the old vocabulary were not fully determined by their inference rules after all; i.e. that there were ingredients of their meanings that were only

[3]In fact Belnap introduces the further constraint of uniqueness: a set of logical laws must *uniquely* determine the meaning of the logical operator governed by them.

[4]See Belnap 1962, p. 132 for passages that strongly support this reading.

[5]So understood the conservativeness constraint just amounts to a special case of our principle of innocence, which we will encounter in the next section.

[6]Our formulation is also in line with the trend in recent literature; see e.g. Dummett (1991, p. 250) and Tennant (forthcoming, p. 15).

brought out by the presence of certain other constants.

The tacit presupposition here is that the meanings of the logical constants can always be specified for one operator at a time. Put another way, we are assuming that the inference rules associated with a particular connective need mention only that connective; they need not appeal to the meaning of any other connective. This is the content of the principle of separability. We may take it on faith for the time being. We will address the question of separability in detail in section 13.2, where it will be shown to follow from our meaning-theoretic assumptions.

There is a further moot point. As Belnap points out, whether or not a given constant is acceptable depends on whether or not the rules we lay down for it are consistent with our prior assumptions. Hence a constant, even if it is acceptable, is always acceptable only 'relative to our characterization of deducibility' (Belnap 1962, p. 133). A similar context-dependence arises also on certain readings of our stronger conservativeness constraint. One and the same logical constant may result in conservative extensions of some systems, but lead to non-conservativeness when adjoined to others. The question therefore arises whether the notion of harmony, like Belnap's notion, should be dependent on context.[7] This would have the effect of making it in effect a property ascribable to pairs $(S, (\$\text{-I}, \$\text{-E}))$ where S is a base system and ($-I, $-E) is a pair of inference rules (or possibly a pair of sets of inference rules) associated with the logical constant that conservatively extends $\mathcal{L}_S$. But this is not quite what we are after. Inasmuch as we intend harmony to be a property of pairs of inference rules alone, conservativeness in and of itself is not a sufficient candidate for explicating harmony.

We could try to reformulate our conservativeness constraint so as to rid it of its context-sensitivity. However, the only conceivable proposal that fits the bill seems to be this: the rules ($-I, $-E) for a logical constant $ are harmonious if and only if the addition of these rules to *any* base system S (and the adjunction of the corresponding constant to $\mathcal{L}_S$) yields a systematically conservative extension. So formulated, the conservativeness constraint does indeed apply to pairs of inference rules alone. As we will see in section 4.4, the downside is that so strong a requirement is hardly defensible: there are some intuitively irreproachable principles of inference that nonetheless produce non-conservative extensions when introduced into a defective

[7] In Belnap's case (on the weaker reading), the context is set by the base system; on our strict reading, it is determined by our structural assumptions and the non-logical expressions of the language.

base system. We should not therefore expect harmonious pairs of rules to be able to extend conservatively *any* system whatsoever. The conservativeness constraint ought to apply only to 'reasonable' systems. But once we impose constraints on the base system, our conservativeness requirement again becomes sensitive to context. Since a reasonable demand for conservativeness is inescapably context-dependent conservativeness cannot be an adequate interpretation of the constraint of harmony as such.

This is not to say that conservativeness has no place in an account of harmony. As a necessary condition for harmony, it may continue to serve us as a test. Indeed, as a test for harmony it can be shown to be extremely effective. Since any conservative extension of a consistent system is itself consistent, Belnap's constraint is sure to rule out `tonk` and other inconsistency-inducing connectives, provided that our base system is consistent.[8] But it also reliably detects less flagrant instances of non-conservativeness.

We shall return to the notion of conservativeness in the form of Dummett's concept of 'total harmony' (see section 4.2). For now, let us briefly turn to item 2. on our list above and see how harmony has been advanced as a means of delineating logical from non-logical vocabulary. We will find that this approach too is underpinned by a form of conservativeness constraint.

3.3 Harmony as a criterion of logicality

It is sometimes held that harmony, in the specific form in which it applies to logical expressions, affords a means of demarcating logical expressions from other types of expressions. Such accounts hold either that the concept of harmony has no significant role to play in the case of non-logical vocabulary (Brandom 2000, p. 71, Tennant 1997, p. 296, fn. 24) or that it *is* applicable to all types of expressions albeit, in the non-logical regions of language, in a different, more relaxed form. On the latter view, properly logical expressions distinguish themselves by the fact that their meanings can be specified by means of inference rules satisfying a particularly crisp, logic-specific requirement of harmony. Call this the *inferentialist approach* to logicality. This approach arguably provides us with a necessary condition for

[8]If this were not so, the extended system would prove any proposition whatsoever, in particular $A \wedge \neg A$ for any A in the restricted language $\mathcal{L}$. *Ex hypothesi* '$A \wedge \neg A$' is provable in the base system, which contradicts our assumption of consistency.

logicality,[9] but it is far from clear that it is also sufficient.[10] We will not pursue this issue here. However, as we will point out along the way, some of our findings in the following chapters will be directly relevant for anyone who hopes to delimit the realm of the logical based on the notion of harmony. In particular, it will be useful here, for future reference, to unearth two fundamental assumptions that underpin not only the inferentialist approach to logicality but also the logical inferentialist's view more generally.

- **The principle of innocence**: logic alone should not be a source of new information. That is, it should not be possible, solely by engaging in deductive reasoning, to discover hitherto unknown (atomic) truths about the world that we would have been incapable of discovering (at least in principle) independently of logic.[11]

- **The principle of autonomy**: the logical fragment is self-contained. That is, the meaning of a logical constant cannot depend on the meanings of non-logical expressions.

There is an obvious connection between the principle of innocence and the considerations that occupied the previous section. The principle of innocence implies that our language plus logical vocabulary has to be a conservative extension of the remainder of our language. This corresponds to the special case of the conservativeness constraint—indeed to the strict interpretation of Belnap's constraint—where the base system is devoid of logical constants. Suppose we begin with a set $\mathcal{A}$ of atomic sentences and a relation of consequence $\vdash_{\mathcal{A}}$ enabling us to derive atomic sentences from other atomic sentences on the basis of certain non-logical rules of inference. The principle of innocence then demands that any expansion of $\vdash_{\mathcal{A}}$ by introducing logical vocabulary $\mathcal{L}$ along with suitable rules of inference leading to $\vdash_{\mathcal{A}+\mathcal{L}}$ should be conservative over $\vdash_{\mathcal{A}}$. The principle of innocence guarantees that

[9] Though even this is not uncontroversial since the inferentialist approach excludes the truth predicate and modal operators from the realm of logical vocabulary; see sections 4.7 and 5.5.

[10] Some people, like Tennant, choose to bite the bullet, holding that the fact that certain expressions commonly classified as non-logical reveal their true (logical) colours on the inferentialist approach to logicality; 'erstwhile "mathematical"
expressions for which we can furnish suitable [i.e. harmonious] introduction and elimination rules will thereby have been revealed as logical' (Tennant 1997, p. 296).

[11] The principle of innocence has a well-known flipside: if logic really does not deliver any new knowledge, then how are we to explain its usefulness? Tempting though it is to wrestle with this question, I will make no attempt to do so here.

the meanings of the non-logical expressions would remain unperturbed by the introduction of logical operators into the language. It ensures that deductive reasoning does not affect the V- or P-principles of other expressions and thereby modify their meanings. In the absence of the principle of innocence, nothing would prevent us from introducing 'constants' into the language that license illegitimate inferential transitions allowing us to assert sentences under novel circumstances. Some of these newly available inferences will be partially constitutive of the meanings of the expressions contained in them. The result would be a shift in the meanings of those expressions; a shift that will in turn have an impact on the meanings of other expressions and so on.[12]

Conversely, the principle of autonomy guarantees that there is no need to appeal to the content of any non-logical expressions in order to determine the meaning of any logical expression. If the principle of innocence shields non-logical expressions from distortions caused by roguish logical expressions, the principle of autonomy ensures that the logical constants are semantically isolable from the rest of language. It states that the realm of the logical constants forms a self-contained language fragment. We can see that the principle of autonomy is a straightforward consequence of our assumption of minimal molecularism.[13]

It should be noted that the principles of innocence and autonomy are closely related. For what guarantees the innocence of a set of logical principles? In order for the principles to be such that they do not disturb the functioning of non-logical regions of language they have to display a delicate balance. As we will see, it is precisely the requirement of harmony that ensures that this balance is maintained. The connection between our two principles consists in the fact that—in general—breaches of the principle of autonomy will lead to an imbalance which will in turn cause violations of the principle of innocence. If the principles of logic are subject to change as a result of global conceptual changes, they will not satisfy harmony and therefore will not display the balance that ensures their innocence. But, going

[12]The role of harmony, as we will see, is precisely to protect the innocence of logic.

[13]Though the molecularist usually holds that there are language fragments that could in principle function as languages in their own right ('in principle', because they need not, of course, ever actually have been in use), it would be absurd to claim that the logical fragment is 'autonomous' in *this* sense. The logical fragment is of necessity parasitic on other fragments of language because it requires an antecedent understanding of the general category of a sentence and of the act of assertion that can be effected by means of it. Hence, although logic does not presuppose the prior existence of any particular region of language, it does presuppose the existence of *some* language.

full circle, violations of the principle of innocence are likely to perturb our linguistic practice so much that the repercussions would be felt at the level of the logical constants, thus contravening the principle of autonomy.

3.4 Harmony: The intuitive notion

We have touched on the content of the notion of harmony at various points. Let us now lay out the informal notion of harmony more systematically. This intuitive notion is to function as a kind of blueprint or guide to our attempts of formulating a suitable, rigorously defined principle of harmony. In section 2.1 we introduced the two-aspect model of meaning: the idea that the meaning of any expression is determined by two aspects of its use, its V-principles and its P-principles. We had argued, moreover, that in the case of the logical constants, these two determinants are naturally expressed by introduction and elimination rules. For a language fragment to be in good working order, these two aspects of meaning—V-principles and P-principles—have to be in *harmony* for all the expressions contained within it. Dummett introduces harmony as an informal notion which he admits is 'difficult to make precise' although it is 'intuitively compelling' (Dummett 1991, p. 215; 1973, p. 397). Let us refer to this intuitive notion of harmony as *general harmony*. The key idea is that the two types of meaning-governing principles cannot be determined independently of one another. More precisely, between the two aspects of meaning a kind of equilibrium ought to obtain: the grounds for asserting a sentence S containing an expression E should be appropriately counter-balanced by the consequences of having asserted it. Given that V-principles state the obligations incurred by a speaker who wishes to assert S containing E, and the P-principles specify the entitlements enjoyed by a listener who understands and accepts the speaker's assertion, harmony ensures that P-principles do not authorize assertions or actions on the part of the listener that would not have been warranted on the basis of the grounds the speaker had for asserting S. Nor, on the other hand, should the P-principles be so restrictive as to disallow assertions or actions (compatible with S) that the listener would have been in a situation to assert or perform prior to the speaker's assertion of S.[14] Either way, the disharmony caused by ill-suited P-principles perturbs the meanings of other expressions by directly affecting the V-principles governing

[14] The last formulation is inspired by Tennant (forthcoming, p. 15).

other expressions—adding to them in the former case and limiting them in the latter thereby threatening to violate the principle of innocence.

Disharmony can thus arise in one of two ways: P-principles may either be too weak or too strong for their corresponding V-principles. In the former case, we find ourselves in the awkward situation where the P-principles governing E do not allow the assertion of sentences (not containing E) that the grounds for asserting S would have licensed. Invoking E thus results in a net loss of information: its unduly restrictive P-principles prevent us from making assertions to which we are in fact entitled. Let us call this type of disharmony *P-weak disharmony* (or *V-strong disharmony*). For a simple example, consider the putative connective \$ whose meaning is determined by the ordinary introduction rules for conjunction:

$$\dfrac{\begin{matrix}\Gamma \\ \vdots \\ A\end{matrix} \quad \begin{matrix}\Gamma' \\ \vdots \\ B\end{matrix}}{\$(A, B)}$$

but equipped with only one of the usual elimination rules, say

$$\dfrac{\begin{matrix}\Gamma \\ \vdots \\ \$(A, B)\end{matrix}}{A}$$

but not

$$\dfrac{\begin{matrix}\Gamma \\ \vdots \\ \$(A, B)\end{matrix}}{B}$$

Dummett illustrates the same point by means of the restricted $\vee$-elimination rule characteristic of quantum logic:

$$\text{Q}\vee\text{-E}, i\ \dfrac{\begin{matrix}\Gamma \\ \vdots \\ A \vee B\end{matrix} \quad \begin{matrix}[A]^i \\ \vdots \\ C\end{matrix} \quad \begin{matrix}[B]^i \\ \vdots \\ C\end{matrix}}{C}$$

Contrary to the standard rule, the restricted ∨-elimination rule does not allow for side premises or assumptions alongside A and B.

$$\dfrac{\begin{matrix}\Gamma_0 \\ \vdots \\ A \vee B\end{matrix} \quad \begin{matrix}\Gamma_1, [A]^i \\ \vdots \\ C\end{matrix} \quad \begin{matrix}\Gamma_2, [B]^i \\ \vdots \\ C\end{matrix}}{C}\ \vee\text{-E}, i$$

Provided that the standard rules governing disjunction are indeed in harmony, we may suspect that the restricted elimination rule is in breach of the requirement of harmony: it fails to deliver all the consequences warranted by the corresponding V-principles (i.e. the introduction rules), namely those deducible from A or B *together* with collateral hypotheses. To investigate these and allied matters with the desirable precision, however, we shall have to wait until we avail ourselves of a more refined notion of harmony for the logical constants. In section 4.4 we will be taking a closer look at the quantum-logical disjunction operator.

The other, perhaps more flagrant violation of harmony occurs if the P-principles enable us to make assertions or legitimize actions that we are not entitled to make or do directly under the circumstances that warranted the introduction of the expression in question. While P-weak disharmony results in a net loss of information, this form of disharmony, *P-strong disharmony* (or *V-weak disharmony*), yields an illicit net gain: the introduction of the expression E makes available inferences that we would not have been in a position to make had we not introduced E via one of the V-principles. We have already seen an extreme example in Prior's `tonk`. As it is sometimes put, the conventional consequences of asserting a statement involving E outrun the circumstances under which its assertion is justified.

Let us briefly take stock of the past chapter. We have, departing from Belnap's response to the `tonk` challenge, developed the notion of a systematic conservative extension. We came to the conclusion that while systematic conservativeness cannot stand in for a principle of harmony as such, it does represent a necessary condition for harmony and so may serve as a test for harmony. In the section that followed, we considered the view that seeks to utilize harmony as a means of delineating logical from non-logical vocabulary. In particular, we focused on two principles—the principles of innocence and of autonomy—that encapsulate the conception of logic that underpins logical inferentialism. We showed how the principles relate to our meaning-theoretic assumptions on the one hand, and how they relate to the notion of harmony. With these coordinates in place we were then able to provide an informal characterization of the notion of harmony on the other. Our task now is to examine the accounts of harmony currently available on the market. We shall begin by assessing Dummett's conception of harmony.

Chapter 4

Dummett on harmony

4.1 Dummett on harmony

So far, we have treated general harmony as a principle applicable possibly to all expressions of a language but also to the logical constants in particular. This is very much in line with Dummett's view. He takes harmony to be a universal principle that regulates all language use. Moreover, he assumes that it is best made precise in terms of Belnap's notion of a conservative extension. The case of `tonk`, on this view, is just a special (and rather spectacular) instance of the general phenomenon of disharmony—an imbalance between V- and P-principles—as it may arise in any region of language. Belnap, recall, proposed to adapt the notion of conservativeness familiar from the context of formal theories (theoretic conservativeness) to systems of logic (systematic conservativeness). In order to align harmony with conservativeness not just in the case of the logical constants but in other regions of language as well, Dummett has to find a way of further broadening the concept of conservativeness, of bringing the notion of conservativeness to bear on languages in general as opposed to only logical expressions in deductive systems. Let us call this notion *linguistic conservativeness*.[1]

In the context of formal theories we defined conservativeness in terms of provability: in the extended theory we should not be able to *prove* any theorems expressible in the language of the base theory. The natural language analogue of formal

[1] I shall explicitly distinguish the three notions of conservativeness—theoretic, systematic and linguistic—only in circumstances where these matter. Also, I shall refrain from doing so in cases where the context makes it clear which type of conservativeness is at stake.

provability for a given expression is all the means language puts at our disposal for justifying an assertion containing the expression. Introducing a new expression into a language thus yields a linguistically conservative extension so long as it does not legitimate us to make assertions (or perform actions) that we should not have been entitled to otherwise (i.e. without appeal to the newly introduced expression). To put it another way, suppose a novel expression does not lead to a conservative extension. Invoking the new expression would now enable us to assert sentences comprising only old vocabulary in new ways, i.e. in circumstances under which their assertion would not have been warranted prior to the introduction of the new expression. Let S be such a sentence. What this means is that the V-principles governing the expressions in S have been altered; the range of circumstances under which we may assert sentences involving (some of) the expressions in S has been extended. How could adopting a new expression into our language have such consequences on the behaviour of preexistent vocabulary? This could only happen if the principles regulating the use of that expression are disharmonious.[2]

According to Dummett conservativeness so understood

> offers us at least a provisional method of saying more precisely what we understand by 'harmony': namely, that there is harmony between the two aspects of the use of any given expression if the language as a whole is, in this adapted sense, a conservative extension of what remains of the language when that expression is subtracted from it (Dummett 1991, p. 219).

Dummett thus identifies conservativeness as the way to cash out general harmony. As such, conservativeness applies across the board to all areas of language and so also to the logical constants. This approach has the virtue of affording a *principled* (i.e. decidedly non-*ad hoc*) way of parrying Prior's challenge: it is a direct consequence of general harmony, which, as we have seen, is required to ensure molecularism. In other words, it flows directly from our meaning-theoretic commitments. And since, moreover, Dummett contends that we have a better handle on the idea of conservativeness than on the intuitively compelling but rather fluffy concept of harmony, he proposes to take the former as a provisional explication of the latter.

[2]The case we have been describing is one of P-strong disharmony, however it is not difficult to imagine the analogous case of P-weak disharmony.

But Dummett's contention that linguistic conservativeness aligns with general harmony can't be quite right. As we have seen in the previous section, general harmony implies conservativeness. However, the converse does not hold: an extension of a language may be conservative, yet the principles governing the newly introduced expression may be disharmonious. This would be the case if the principles were flawed on account of *P-weak disharmony.* Conservativeness guards against extensions of our language that enable us to say *more* about the old vocabulary; it offers no protection against extensions that result in our being able to say *less* about the old vocabulary. This thought may lead one to suspect that the introduction of an expression governed by disproportionately weak P-principles is not only compatible with conservativeness, but that it even entails conservativeness. For suppose the meaning of an expression is given by a set of principles that displays P-weak disharmony. In that case the V-principles will invest the expression with a meaning that the P-principles are too weak fully to exploit. The introduction of an expression governed by such principles will thus license *fewer*, not more assertions of statements not involving the expression in question. It would then seem to follow that the addition of an expression governed by P-weakly disharmonious principles will result in a conservative extension of the base system (which we assume to be in good order). However, this reasoning is marred by the assumption that P-weak and P-strong disharmony are mutually exclusive. But this is not so. For assume the connective $ is governed by the standard introduction rules for conjunctions, but that it enables us to infer that aardvarks are nocturnal. The rules for $ would then be *both* P-weakly and P-strongly disharmonious. Consequently, P-weak disharmony does not entail conservativeness. Interestingly, it does not entail conservativeness *even if* we assume the system *not* to be P-strongly disharmonious. There are cases where adding a perfectly harmonious logical constant to certain deductive systems that contain P-weakly disharmonious (though not P-strongly disharmonious) expressions results in a non-conservative extension. We shall encounter such a system involving the aforementioned restricted or-elimination rule in section 4.4.

But let us return to the problems we found in Dummett's hasty assimilation of general harmony with conservativeness. It is not clear what Dummett has in mind here. At times it seems as if Dummett anticipates the arguments of the previous section, as when he tentatively claims that cases of P-weak disharmony do not 'produce so deleterious an effect' (Dummett 1991, p. 218). The interference with prior linguistic practice caused by P-weak disharmony—although a violation of general harmony—is presented as negligible, or at least in some sense less pernicious than the effects wrought by P-strong disharmony. This might have some plausibility in the case of the logical constants: it is not possible to concoct a P-weakly disharmonious constant that can match `tonk` in sheer destructive capacity, i.e. a P-weak constant that by itself can reduce an entire (otherwise harmonious) system to triviality. But even in the realm of the logical constants P-weak disharmony does not guard against non-conservativeness as Dummett's own discussion of the quantum or-elimination rule (Dummett 1991, p. 290) shows (see section 4.4). Moreover, he explicitly argues that also P-weak disharmony is no lesser an evil than P-strong disharmony also in other contexts (Dummett 1991, p. 206). In both cases, P-strong as well P-weak disharmony, the existing practice is disturbed in inestimable ways—the speaker loses a clear sense of what it is he is saying.[3] What is particularly puzzling is that Dummett eventually does react to the threat of P-weak disharmony in the case of the logical constants by introducing the notion of stability (of which more later, see section 6.1), but he makes no attempt to amend his characterization of harmony in the general case accordingly.

The project of devising a precise principle of harmony is more tractable for the logical constants than for other types of expressions for reasons given earlier when we were explaining why the logical operators are particularly amenable to inferentialist accounts based on the two-aspect model of meaning. First: logic is formal; therefore the meanings of the logical operators can be specified in the form of schematic inference rules (see section 2.1). Secondly and relatedly, V- and P-principles for the logical constants are given by purely formal immediate inferences, facilitating the implementation of the idea of a balance between the two components of meaning. Since the logical fragment of language is our focus, we may set aside the problem of formulating harmony for the remainder of language.[4] It is Dummett's treatment of harmony for the logical constants that will occupy us in the following sections.

---

[3]See also our discussion of this point in section 2.4.

[4]It is far from obvious how one might go about formulating a universal principle of harmony for the whole of language—a formidable task, as even optimists will concede. As I mentioned above, some, most explicitly Brandom (2000, p. 71), have rejected the very idea that harmony should find any application outside of the confines of logical vocabulary. Such a position is compatible with our assumption of minimal molecularism. Fortunately, however, we need not take a stance on this matter here.

4.2 Levelling local peaks and normalizability

The problem now is to give a precise characterization of Dummett's notion of harmony as it relates specifically to introduction and elimination rules for the logical constants. First, though, let us specify the form our target notion, general harmony, takes in this context. Let $ be a logical constant: harmony reigns when its introduction rules are matched by corresponding elimination rules in such a way that nothing more and nothing less may be inferred from a statement containing $ in a dominant position than is warranted by the premises of the introduction rules. In other words, the rules for $ should be balanced, in the sense that the elimination rules ought to exploit *all* and *only* the inferential powers the introduction rule has invested it with.

In the course of his discussion of harmony as it applies to the logical constants, Dummett offers three distinct notions of harmony: total harmony, intrinsic harmony and stability. In the end only two of them will survive, as intrinsic harmony is replaced by a strengthened version of itself, stability.[5] What Dummett calls *total harmony* is a redeployment of the conservativeness constraint within the logical fragment. It corresponds exactly to our original formulation of systematic conservativeness (see section 3.2); i.e. it amounts to the requirement that the introduction of a new logical constant should result in a conservative extension of the logical system to which it is added. Note that total harmony is specific to the domain of logic and so differs from the global conservativeness constraint Dummett equates 'provisionally' with harmony *tout court* for language as a whole. The former is a requirement of systematic conservativeness dealing with the extension of systems of logic rather than of languages or fragments thereof, while the latter is one of linguistic conservativeness. The two are related in that any violation of total harmony is *ipso facto* a violation of Dummett's generalized conservativeness constraint; the former is a localized special case of the latter.

Total harmony, which is 'in a high degree relative to the context', is contrasted with *intrinsic harmony*, 'a property of the rules governing the logical constant in question' (Dummett 1991, p. 250). Intrinsic harmony is based on what Prawitz called the *inversion principle*, an elaboration of Gentzen's famous programmatic remarks. In Prawitz's words,

[5]We will deal with Dummett's account of stability separately in the section 6.1.

> an elimination rule is, in a sense, the inverse of the corresponding introduction rule: by an elimination rule one essentially only restores what had already been established by the major premiss of the application of an introduction rule (Prawitz 1965, p. 33).

If an elimination rule is really just a device for 'undoing' a primitive inferential move effected by an application of an introduction rule (and vice versa), it should not be possible, simply by introducing and subsequently eliminating a logical constant, to arrive at new conclusions about the world (nor should it be possible to disallow any conclusions warranted by its introduction rules). And such a parity should obtain when the rules of inference are harmonious in the sense of our intuitive notion of general harmony.

This general idea has been cashed out by Prawitz as follows: for any pair of purportedly harmonious inference rules for a constant \$, there must be a procedure enabling us to transform any proof from the set of hypotheses Γ to a conclusion C in the course of which \$ is introduced and subsequently eliminated into a deduction that reaches the same conclusion, but without the superfluous detour via \$. Though all this is familiar ground, it is worth fixing our terminology. Let us call any procedure for removing detours in this sense a *reduction procedure*.

Where the introduction rule for \$ is *immediately* followed by a \$-elimination, we speak of a *local peak* (with respect to \$). The sentence containing \$ as its main connective, which serves simultaneously as the conclusion of the \$-introduction rule and the major premise of the corresponding elimination rule in the local peak, we shall call a (\$-)*maximum*; 'maximum' because such formulas are logically more complex than the sentences in the immediate vicinity on the same deductive path. We shall, following Dummett's metaphorically apt formulation, refer to reduction procedures that ensure the dispensability of maxima as *levelling* (of local peaks) or as the *elimination* (of local peaks or of maxima). It is this procedure of levelling that serves as the basis for Dummett's account of *intrinsic* (or *local*) harmony for the introduction and elimination rules: intrinsic harmony can be 'provisionally identified [...] with the possibility of carrying out this procedure, which we have called the levelling of local peaks' (Dummett 1991, p. 250).

Let us briefly pause to illustrate intrinsic harmony with the aid of two standard examples. Consider first the case of $\supset$ with its familiar introduction and elimination rules,

$$\dfrac{\begin{matrix}\Gamma, [A]^i \\ \vdots \\ B\end{matrix}}{A \supset B}\ {\scriptstyle \supset\text{-I},\, i}$$

and

$$\dfrac{\begin{matrix}\Gamma_0 & \Gamma_1 \\ \vdots & \vdots \\ A \supset B & A\end{matrix}}{B}\ {\scriptstyle \supset\text{-E}}$$

A local $\supset$-peak consequently has the following form:

$$\dfrac{\dfrac{\begin{matrix}\Gamma_0, [A]^i \\ \Pi_0 \\ B\end{matrix}}{A \supset B}\ {\scriptstyle \supset\text{-I},\, i} \quad \begin{matrix}\Gamma_1 \\ \Pi_1 \\ A\end{matrix}}{B}\ {\scriptstyle \supset\text{-E}}$$

Levelling in this case amounts to the simple cut-and-paste job of appending the proof Π_0 to the end of the proof Π_1 of A (keeping in mind that the proof Π_0 may require additional hypotheses, Γ_0, alongside A). We obtain the following proof of B directly from Γ_1:

$$\begin{matrix} & \Gamma_1 \\ & \Pi_1 \\ \underbrace{\Gamma_0, A} & \\ \Pi_0 & \\ B & \end{matrix}$$

Because local $\supset$-peaks can be levelled, Dummett would grant that $\supset$-rules are in intrinsically harmonious.

For a second example consider the case of disjunction with its familiar introduction rules (where $j = 0$ or 1)

$$\dfrac{\begin{matrix}\Gamma \\ \vdots \\ A_j\end{matrix}}{A_0 \vee A_1}\ {\scriptstyle \vee\text{-I}}$$

and the aforementioned corresponding elimination rule

$$\dfrac{\begin{array}{c}\Gamma_0\\ \vdots\\ A \vee B\end{array} \quad \begin{array}{c}\Gamma_1, [A]^i\\ \vdots\\ C\end{array} \quad \begin{array}{c}\Gamma_2, [B]^i\\ \vdots\\ C\end{array}}{C}\ \vee\text{-E}, i$$

Suppose we have a local peak featuring $\vee$ of the following form

$$\dfrac{\dfrac{\begin{array}{c}\Gamma_0\\ \Pi_0\\ A\end{array}}{A \vee B}\ \vee\text{-I} \quad \begin{array}{c}\Gamma_1, [A]^i\\ \Pi_1\\ C\end{array} \quad \begin{array}{c}\Gamma_2, [B]^i\\ \Pi_2\\ C\end{array}}{C}\ \vee\text{-E}, i$$

Our proof is then straightforwardly transformed into one that avoids the detour through the introduction of $A \vee B$: we do this by concatenating the proof Π_0 of A with the proof Π_1 of C from A (and similarly for the other case)

$$\begin{array}{c}\Gamma_0\\ \Pi_0\\ \underbrace{\Gamma_1, A}\\ \Pi_1\\ C\end{array}$$

The rules for $\vee$ thus also satisfy intrinsic harmony.

The levelling of local peaks plays a crucial role in the proof of the normalization theorem, Prawitz's natural deduction incarnation of Gentzen's *Hauptsatz*. The idea underlying the normalization theorem is that any proof that proceeds via detours can be converted into a normal form where there is a direct deductive route joining hypotheses and conclusion. Normalized proofs enjoy a sub-formula property akin to that of cut-free sequent calculi. Roughly, any formula occurring in a deduction in normal form of A from hypotheses Γ is either a sub-formula of A or a sub-formula of at least one of the formulas contained in Γ.[6]

Levelling procedures guarantee the crucial inductive steps in the proof of normalization. This is *not* to say, however, that there is no more to normalization than the levelling of local peaks. What is required is not only a demonstration of the eliminability of maxima; it is necessary, moreover, to show that proofs in the course of which formulas are introduced and subsequently eliminated, but where the succession of introductions and eliminations is not immediate, can be permuted in

[6] Things are not quite as neat in classical logic, see e.g. Prawitz (1965, p. 42).

such a way as to create a maximum. This standardly requires a rearranging of the order of application of the inference rules involved in the proof. Once a local peak has been created in this way, it can be dealt with in the familiar fashion. Reduction procedures other than levelling thus generally have the function of manipulating the order of application of inference rules in order to show that detours are avoidable. A proof can be said to be in *normal form* if no further procedures (neither levelling nor the auxiliary procedures just mentioned) can be applied to it. A system is *normalizable* if any proof within it can be converted into normal form.

4.3 Other reduction procedures

Let us briefly dwell on the notions of intrinsic harmony and of normalizability. Given that Dummett is careful—as we have been—to distinguish the *global* property of normalizability (a property of a system) from the *local* property of pairs (of sets) of inference rules consisting in the possibility of eliminating local peaks (Dummett 1991, p. 250), one is surprised to find that an accomplished logician like Stephen Read—referring to the very same page in Dummett—repeatedly attributes to Dummett the view that *intrinsic* or local harmony should be identified with normalizability.[7] One is led to suspect that Read uses 'normalizability' in a non-standard way. And indeed, a number of passages suggest that he uses 'normalizability' interchangeably with the possibility of 'levelling local peaks' (Read 2000, p. 128, and forthcoming, p. 13). So it appears that when Read speaks of 'normalizability' we should simply take him to be talking about what we have called 'the levelling of local peaks' (or equivalently, 'the eliminability of maxima'); similarly, his talk of a 'normalization theorem' corresponds exactly to what Prawitz calls an *inversion theorem*:

> If $\Gamma \vdash A$ then there is a deduction of A from Γ in which no formula occurrence is both the consequence of an application of an I[introduction]-rule and major premise of an application of an E[elimination]-rule (Prawitz 1965, p. 34).

So far this seems rather harmless: Read is of course free to use the term as he sees fit. One problem, however, is that he appears to be unaware of the deviance of

[7]'Dummett distinguishes, in fact, between total harmony, which he equates with conservativeness, and local [i.e. intrinsic] harmony, which he equates with normalization' (Read 2000, p. 126, see also Read (2008, p. 13).

his use of the term; he apparently fails to notice that his definition is at odds with the one given by Dummett on the aforementioned page 250 of *The logical basis of metaphysics*. But a more critical problem is that Read's own use of 'normalization' appears to be inconsistent. He conflates the two uses—normalization as levelling and normalization in our stronger sense—when he claims that 'normalization, in Dummett's helpful metaphor, "levels local peaks": maximal formulae can be removed, and proofs obey the sub-formula property' (Read 2000, p. 128).[8] So Read in fact mistakenly takes the eliminability of local peaks to be sufficient for normalization in the strong sense.

Given that the use of 'normalization' and kindred terms tends to fluctuate significantly in the literature, and given that, as we have seen, these fluctuations are a source of genuine confusion, it is worth dwelling on normalization for a moment. To get clear on what, in addition to the levelling of local peaks, is needed to establish normalization, let us give a brief sketch of the remaining reduction procedures usually required. We shall content ourselves with informal illustrations of procedures for the familiar operators and their rules.

As we have seen, normalization, as we conceive of it, obtains when any proof within a system can be transformed into a normal form. A proof is in normal form if *all* detours (as opposed to only local peaks) can be removed. Roughly, any proof in normal form falls into two parts divided by a kind of proof-theoretic equator. North of the equator is the realm of eliminations: all instances of elimination rules (if any) apply to initial assumptions and their consequences until we end up—entering now into the southern hemisphere—only with sub-formulas of the final conclusion to which introduction rules are applied. As Tennant fittingly puts it,

> the striking thing about a proof in normal form is that the reasoning it represents 'dismantles' the premises into their constituents and 'reassembles' the information therein so as to form the conclusion' (Tennant 2005b, p. 630).

The first additional type of reduction procedure required—we have already alluded to it above—may be called *permutative reductions*, borrowing Dummett's

[8] See also p. 124 and p. 130 where Read equates the possibility of levelling local peaks with normalization and the latter with cut-elimination. Cut-elimination corresponds to normalization in the full standard sense of the term. This suggests that the passage quoted above is not a mere slip of the pen (Read 2000).

terminology (Dummett 1977, p. 112). To illustrate the necessity of permutative reductions consider the following deduction of C from the hypotheses $(A \wedge B) \vee (B \wedge E)$ and $C \wedge D$

$$
{\scriptstyle \wedge\text{-E}}\,\dfrac{{\scriptstyle \vee\text{-E},\,1}\,\dfrac{(A \wedge B) \vee (B \wedge E) \qquad {\scriptstyle \wedge\text{-I}}\,\dfrac{{\scriptstyle \wedge\text{-E}}\,\dfrac{[A \wedge B]^1}{B} \quad {\scriptstyle \wedge\text{-E}}\,\dfrac{C \wedge D}{C}}{B \wedge C} \qquad {\scriptstyle \wedge\text{-I}}\,\dfrac{{\scriptstyle \wedge\text{-E}}\,\dfrac{[B \wedge E]^1}{B} \quad {\scriptstyle \wedge\text{-E}}\,\dfrac{C \wedge D}{C}}{B \wedge C}}{B \wedge C}}{C}
$$

Note that the proof contains no local peak and hence no maximal formula according to our definitions. Although $B \wedge C$ has been introduced and subsequently eliminated in the course of the proof, the succession is not immediate: it is delayed by an instance of $\vee$-elimination. We therefore cannot apply our levelling procedure directly. Yet the proof is not in normal form since it does not satisfy the sub-formula property: $B \wedge C$ is neither a subformula of any of the hypotheses nor of the conclusion. Though $B \wedge C$ does not constitute a maximum, its introduction was clearly an unnecessary detour. The situation may be remedied by the use of the permutative reduction procedure. With its help we can reshuffle the order of the proof in such a way as to bring the applications of the $\wedge$-introduction and $\wedge$-elimination rules into immediate succession, thus creating a maximum. This can be achieved by permuting the order of application of the $\vee$-elimination rule and the $\wedge$-elimination rule.

$$
{\scriptstyle \vee\text{-E},\,1}\,\dfrac{(A \wedge B) \vee (B \wedge E) \qquad {\scriptstyle \wedge\text{-E}}\,\dfrac{{\scriptstyle \wedge\text{-I}}\,\dfrac{{\scriptstyle \wedge\text{-E}}\,\dfrac{[A \wedge B]^1}{B} \quad {\scriptstyle \wedge\text{-E}}\,\dfrac{C \wedge D}{C}}{B \wedge C}}{C} \qquad {\scriptstyle \wedge\text{-E}}\,\dfrac{{\scriptstyle \wedge\text{-I}}\,\dfrac{{\scriptstyle \wedge\text{-E}}\,\dfrac{[B \wedge E]^1}{B} \quad {\scriptstyle \wedge\text{-E}}\,\dfrac{C \wedge D}{C}}{B \wedge C}}{C}}{C}
$$

This being done, we may proceed as usual with a straightforward application of our levelling procedure to the minor premises of the $\vee$-elimination rule. Permutative reductions—though sometimes requiring iterated application—provide us with the resources necessary to reduce stretched peaks—call them *plateaux*—to local peaks. As we noted, plateaux occur when the elimination of a constant previously introduced is delayed by one or more applications of particular types of elimination rules (usually $\vee$- or $\exists$-elimination). Permutative reductions thus become necessary whenever our system contains elimination rules that involve sub-deductions from hypotheses as in the case of the elimination rules for $\vee$ and $\exists$. We usually have to devise specific permutative reduction procedures for all such elimination rules. Because of this, permutative reductions, unlike levelling (which can be carried out in the presence of any operational rules so long as the consequence relation is transitive), is sensitive to the types of rules contained in the system. Certain configurations of rules may create obstacles for the elimination of detours.

Aside from levelling and permutative reductions, we also need to ensure that applications of ∨- or ∃-elimination rules are not redundant. An instance of an ∨-elimination (respectively ∃-elimination) rule is redundant; i.e. if the assumptions on which the minor premise (or minor premises) depends are not left undischarged. In the case of ∨-elimination, its application is redundant if the set of hypotheses on which the conclusion depends wholly contains the hypotheses of one of the minor premises. In such cases an application of, say, an ∨-elimination rule

$$\dfrac{\begin{matrix}\Gamma_0 \\ \Pi_0 \\ A \vee B\end{matrix} \qquad \begin{matrix}\Gamma_1, [A]^1 \\ \Pi_1 \\ C\end{matrix} \qquad \begin{matrix}\Gamma_2, [B]^1 \\ \Pi_2 \\ C\end{matrix}}{C}\ {\scriptstyle \vee\text{-E, } 1}$$

may simply be converted (taking the case where it is the first minor premise that is undischarged) to

$$\begin{matrix}\Gamma_1, [A] \\ \Pi_1 \\ C\end{matrix}$$

The case of ∃-elimination is similar.

4.4 Locality and globality

So much, then, for normalization and the ancillary reduction procedures it requires.[9] The most important point of the preceding sections is that normalization is *a property of a proof system as a whole*. The mere inspection of a set of inference rules without any knowledge of the properties of the deductive system of which they are a part is not in general sufficient to determine whether the system normalizes. In short, normalization is a *global* property of a system. Levelling, on the other hand, is *a structural feature* of pairs of introduction and elimination rules; it is, as we have already noted, a *local* property of the rules.

[9]Further procedures may be needed depending on the particular language and the structural and operational rules adopted.

Dummett gives an illuminating illustration of the globality of normalizability. Let us call this example the *Q-example* for future reference. Consider again the restricted disjunction elimination rule familiar from quantum logic, $Q\vee$-E

$$Q\vee\text{-E}, i\ \dfrac{\begin{matrix}\Gamma\\ \vdots\\ A\vee B\end{matrix}\qquad \begin{matrix}[A]^i\\ \vdots\\ C\end{matrix}\qquad \begin{matrix}[B]^i\\ \vdots\\ C\end{matrix}}{C}$$

As we noted, the quantum-disjunction elimination rule distinguishes itself from the standard disjunction elimination rule by disallowing collateral hypotheses in the minor premises. Let us denote the quantum-logical disjunction operator whose meaning is given by the usual introduction rules and the restricted elimination rule by '$\sqcup$'. Since the procedure for eliminating local peaks we gave for standard disjunctions in section 4.2 applies equally in the case of $\sqcup$, the rules for $\sqcup$ are intrinsically harmonious by Dummett's standards. However, as the Q-example shows, intrinsic harmony is insufficient to guarantee general harmony. Even if all the pairs of rules in a system satisfy intrinsic harmony, the system as a whole might still fail to normalize. Suppose we situate ourselves in a system containing only, say, $\wedge$ and $\sqcup$. Dummett now observes that the addition of $\vee$ with its usual unrestricted elimination rule collapses quantum-logical disjunction into standard disjunction.

$$\sqcup\text{-E}, i\ \dfrac{\begin{matrix}\Gamma_0\\ \vdots\\ A\sqcup B\end{matrix}\qquad \vee\text{-I}\ \dfrac{[A]^i}{A\vee B}\qquad \vee\text{-I}\ \dfrac{[B]^i}{A\vee B}}{A\vee B}$$

We have thus derived the full $\vee$-elimination rule by appeal to $\sqcup$ alone; any grounds sufficient to warrant the assertion of $A\sqcup B$ will be *ipso facto* sufficient to assert $A\vee B$, and will so justify the application of the unrestricted elimination rule. What this shows is that the rules governing $\sqcup$ do not succeed in conferring on it a stable meaning. This is also reflected in the fact that the resulting system ($\{\wedge,\sqcup,\vee\}$) is not a conservative extension of the base system: in particular, the law of distributivity $A\wedge(B\sqcup C)\vdash(A\wedge B)\sqcup(A\wedge C)$ becomes readily derivable for $\sqcup$.

What is more, in the process of deriving $\vee$ from $\sqcup$ we have created a plateau: the instance of $\sqcup$-elimination following the application of the $\vee$-introduction rule delays the subsequent $\vee$-elimination. However, when we now try to apply our permutative reduction procedure we obtain the following

$$\sqcup\text{-E, 3}\ \dfrac{\begin{matrix}\Gamma_0\\ \vdots\\ A \sqcup B\end{matrix} \qquad \vee\text{-E, 1}\ \dfrac{\vee\text{-I}\ \dfrac{[A]^3}{A \vee B} \quad \begin{matrix}\Gamma_1, [A]^1\\ \vdots\\ C\end{matrix} \quad \begin{matrix}\Gamma_2, [B]^1\\ \vdots\\ C\end{matrix}}{C} \qquad \vee\text{-E, 2}\ \dfrac{\vee\text{-I}\ \dfrac{[B]^3}{A \vee B} \quad \begin{matrix}\Gamma_3, [A]^2\\ \vdots\\ C\end{matrix} \quad \begin{matrix}\Gamma_4, [B]^2\\ \vdots\\ C\end{matrix}}{C}}{C}$$

But the final application of $\sqcup$-elimination is not in general permissible—it is legitimate only in special cases where Γ_{1-4} all happen to be empty. Therefore, in our system the application of reduction procedures may lead us from genuine deductions to ill-formed ones.

Summing up, Dummett has produced a system composed exclusively of intrinsically harmonious pairs of sets of inference rules that is nonetheless not normalizable and does not display total harmony. This shows that the P-weak inference rules for $\sqcup$ failed to fix its meaning. We may conclude that the notion of intrinsic harmony is an inadequate characterization of harmony.

Dummett concludes from this that a strengthening of the notion of intrinsic harmony is needed. He conjectures that an amended version of intrinsic harmony, which he dubs *stability*—intrinsic harmony supplemented by a provision for ruling out P-weak disharmony—will, if reigning over all the logical constants of the system, entail total harmony.[10]

4.5 Why go local?

We have seen Dummett's distinction between intrinsic harmony (which, recall, Dummett had 'provisionally identified' (Dummett 1991, p. 250) with the possibility of levelling local peaks) and total harmony. But we have yet to get clear about Dummett's motivation for making the distinction in the first place. Given that the Q-example has revealed the insufficiency of intrinsic harmony as a way of cashing out general harmony, Dummett goes on to develop the notion of stability. Intrinsic harmony is only half of the story. The new and improved concept of stability must in addition ward off P-weak disharmony (ibid., p. 287). As it is intended, at least, stability promises to be the most faithful match for our intuitive notion of general

[10] I will use the term 'stability' more loosely to cover any proposal that seeks to improve on the notion of intrinsic harmony by barring P-weak disharmony. On my understanding of the term both Dummett's account of stability and Tennant's account of harmony (see chapter 6) are to be regarded as attempts to formulate a principle of stability.

harmony. One would expect, therefore, that stability would take centre stage in Dummett's account.

Surprisingly this is not so. Rather than focusing on stability, which, like the notion of intrinsic harmony (and also like our target notion of general harmony), acts locally, Dummett privileges a global conception of harmony. We have already seen that he makes no attempt at articulating a local harmony constraint that could apply to all expressions of a language, but rather equates harmony with the global property of linguistic conservativeness (Dummett 1991, p. 214). The same, apparently, holds at the level of the logical constants, according to Dummett. He tells us that we have

> proceeded as though intrinsic harmony was all that mattered; but it is total harmony that must prevail if the *point* of the requirement of harmony is to be attained, namely, that, for every logical constant, its addition to the fragment of the language containing only the other logical constants should produce a conservative extension of the fragment (Dummett 1991, p. 290).

Stability, on this account, would be a sufficient but not necessary condition for total harmony to obtain, and so for the 'point of harmony' to be met.

But the 'point of harmony', surely, is that *general* harmony obtains. As we have argued above in section 3.4, conservativeness, both at the level of the logical constants and for non-logical expressions, will always fall short of fully capturing the notion of general harmony, which requires the prevention of not just P-strong but also of P-weak disharmony. Dummett, after wavering on the issue, thus comes down on the wrong side and privileges total harmony over stability.

The mistake is not only that Dummett overlooks the inadequacy of the notion of conservativeness (which he himself helps to expose to some extent in his discussions of quantum logic). Rather, the very idea of equating harmony with a principle that acts globally is misguided. The target notion, our *formalisandum*, is general harmony. And general harmony is a local constraint (or at least we may assume that this holds true for the logical constants): it obtains between the V-principles and the P-principles governing a given expression when these principles are appropriately balanced. The crucial point being that whether or not this balance obtains should depend only on the principles being assessed, not on the larger context of the system in which they inhere. This is of course not to deny that the *effects* of harmony (or better of disharmony) are felt on a global level. After all, for Dummett the chief motivation behind the notion of harmony, its *raison d'être*, is to ensure the good functioning of a language as a whole. The task must be to formulate local principles in such such a way as to avert global disaster.

This, presumably, is also the motivation that drives Dummett to introduce the notion of intrinsic harmony and the new and improved version of it, stability. But why then does he abandon these local principles in favour of the global concept of conservativeness? One can do little more than speculate. As we shall see in section 6.1, Dummett never goes on to develop the principle of stability to a tolerable degree of precision. Moreover, even if a precise notion of stability were at hand, it still remains entirely unclear what form an analogous local principle for the rest of language could take. Perhaps reasons like these led Dummett to conclude that stability is a dead end. The notion of conservativeness, on the other hand, appears to tick both boxes: it holds the promise of being clearly statable and it is applicable both in the restricted context of the logical constants and to language in general (cf. Dummett 1991, p. 287).

Our situation is thus the following. We have shown conservativeness to be unsatisfactory *qua* definition of harmony. As our sole focus here is the notion of harmony within the context of the logical constants, the problem of finding principles for the logical fragment that have analogues within more extensive regions of language need not faze us. Rather, we must ask ourselves whether we are right to insist that our requirement of harmony ought to be *local*? Why should this be the case? After all, we could imagine scenarios in which local disharmony between V- and P-principles could be compensated for at a global level. But how global are we talking here? We have said that for the molecularist the class of expressions a speaker needs to understand in order to grasp the meaning of any given expression must always be a proper subset of the class of all expressions in the language. The relevant dependence relations must not mushroom out to *all* other expressions. Does it follow that the set of expressions we need to look to for possible compensation for local disharmonies is equally bound? Perhaps this would be a reasonable assumption to make if we were concerned with language in general (although it does not obviously follow from molecularist principles).

Fortunately no such additional assumption is required in the case of the logical operators. We may distinguish two degrees of globality here. Compensation can in principle occur at each of these levels: it could occur in the realm of non-logical expressions, thus contradicting our principle of autonomy (see section 3.3); or it could occur closer to home, within the same class, i.e. the class of the logical expressions, thus contradicting the aforementioned principle of separability (see section 2.5). But it will be recalled that we have not yet defended the principle of separability; this issue occupies us in section 13.2 where we will mount a defence of separability.

Having decided to leave this problem for later, our task in the sections to come is to advance our own version of a local principle of harmony. Before we turn to this large task, however, it will be useful to attain a better understanding of the role the oft-mentioned notions of conservativeness and of normalizability are to play in an account of harmony. If harmony is to be a locally acting principle that applies to inference rules (as we are assuming), then neither conservativeness nor normalizability can be a sufficient condition for harmony. We have seen this in the case of conservativeness. Matters are less obvious in the case of normalizability, but there are normalizable systems incorporating constants that are not regulated by intuitively harmonious rules of inference (according to the standards set by our meaning-theoretic principles). Classical logic is such a system. Of course, the staunch defender of classical logic is likely to take this to be a *reductio ad absurdum* of our meaning-theoretic assumptions. But this should not deter us. Our aim is to explore the logical inferentialist's principles, following them wherever they may lead us without prejudice as to which logic ought to be validated at the end of the day.

Let us then investigate how these notions—conservativeness and normalizability—relate to the notion of general harmony and to one another. The Q-example showed us that intrinsic harmony does not ensure total harmony, or systematic conservativeness. However, Dummett conjectures that the stronger local requirement of stability *is* a sufficient condition for systematic conservativeness, i.e. that stability entails total harmony (Dummett 1991, p. 290). Let us dub this *Dummett's conjecture*. It may be stated as follows:

> **Dummett's conjecture:** If a system S is composed exclusively of stable rules of inference and a new constant c also governed by stable rules is added to S to form S', then S' is systematically conservative over S.

Note again that, as we discussed in section 3.2, total harmony is not a property of pairs of inference rules, but depends also on the base system. In section 4.7 we shall discuss Dummett's conjecture and a number of objections that have been levelled against it.

By contrast, it is not immediately clear what role normalizability plays in Dummett's account. At times, Dummett argues that normalizability entails conservativeness where conservativeness is now understood in the stronger sense that 'for each logical constant c, the full language is a conservative extension from that obtained by omitting c from its vocabulary' (Dummett 1991, p. 250). Call this property of a system *full conservativeness*. Of course full conservativeness, like normalizability, is a global property of a system. But Dummett also argues that the failure of the system explored in section 4.4 involving both $\vee$ and $\sqcup$ to normalize is 'the reason' (ibid.) for the non-conservativeness (and hence *a fortiori* of non-full conservativeness) of the extended system over the system $\{\wedge, \sqcup\}$. This suggests, contrapositively, that the converse implication—from full conservativeness to normalizability—holds also. Let us call these two claims combined *Dummett's claim*. It can be put as follows:

> **Dummett's claim:** Let S be a system of logic containing only constants governed by stable inference rules and c a logical constant. If the rules governing c are stable, the addition of c to S will result in a normalizable system S' just in case S' is fully conservative (and hence S' is a conservative extension of S).

It is to this latter claim that we turn first.

4.6 Normalizability and conservativeness

It is enough to examine the argument advanced in support of the first half of Dummett's claim, namely that normalizability entails full conservativeness. Suppose a system S is normalizable. Take any constant c of S:

> if we have a proof whose final sequent does not contain c, any sentence [occurring in the proof] containing c must first have been introduced by an introduction rule, and then eliminated by an elimination rule; hence, by normalization, we can obtain a proof not involving that sentence (Dummett 1991, p. 250).

Therefore, it may seem that taking any constant c of S and adding it to $S - \{c\}$ will result in a conservative extension. Dummett takes this to demonstrate that

S is fully conservative. But in this he is wrong. It is not difficult to produce a counterexample drawing on any of the well-known instances of non-conservativeness in classical logic. Here we will construct a counterexample employing Peirce's Law: $((A \supset B) \supset A) \supset A$.[11]

Consider the fragment of classical logic $S = \{\wedge, \supset\}$ with the familiar introduction and elimination rules for these connectives. The normalizability of S is obvious. In the absence of the elimination rules for $\vee$ and $\exists$ that necessitate additional reduction procedures, all detours must be instances of local maxima and so we can cheerfully level them away. Now take $S' = S \cup \{\neg\}$, which is obtained by adding introduction and elimination rules for $\neg$,

$$\dfrac{\begin{array}{c}\Gamma, [A]^i \\ \vdots \\ \bot\end{array}}{\neg A}\ \neg\text{-I}, i$$

$$\dfrac{\begin{array}{c}\Gamma \\ \vdots \\ A\end{array} \qquad \begin{array}{c}\Gamma' \\ \vdots \\ \neg A\end{array}}{\bot}\ \neg\text{-E}$$

and the classical *reductio ad absurdum* rule ('*CRAA*'):

$$\dfrac{\begin{array}{c}\Gamma, [\neg A]^i \\ \vdots \\ \bot\end{array}}{A}\ CRAA, i$$

S' is also normalizable. Despite the fact that the classical rules for negation do not admit of a levelling procedure, it can still be shown that all proofs invoking negation can be converted into maxima- and plateaux-free ones. This is done by providing a levelling procedure for the $\neg$-I and $\neg$-E rules and by showing that all applications of *CRAA* can be reduced to instances where the conclusions are atomic sentences. To start with the levelling procedure for $\neg$, consider a $\neg$-maximum:

[11] I am of course not suggesting that Dummett is unaware of examples of non-conservativeness like that of Peirce's Law. I am simply pointing out that he must have failed to take them into account when formulating the above argument for the implication of full conservativeness by normalizability.

$$\neg\text{-E}\ \dfrac{\begin{array}{c}\Gamma\\ \Pi\\ A\end{array} \qquad \neg\text{-I},\, i\ \dfrac{\begin{array}{c}\Gamma', [A]^i\\ \Pi'\\ \bot\end{array}}{\neg A}}{\bot}$$

Such a local peak can be levelled as follows:

$$\begin{array}{c}\Gamma\\ \Pi\\ \underbrace{\Gamma', A}\\ \Pi'\\ \bot\end{array}$$

Turning now to the reduction procedures required to show that any application of *CRAA* can be transformed into one acting only on atomic formulas, assume we have an instance of *CRAA* applied to a complex formula of the form $\neg(A \wedge B)$

$$CRAA, i\ \dfrac{\begin{array}{c}\Gamma, [\neg(A \wedge B)]^i\\ \Pi\\ \bot\end{array}}{A \wedge B}$$

The appropriate reduction procedure for the case of conjunction is this:

$$\wedge\text{-I}\ \dfrac{CRAA, 2\ \dfrac{\begin{array}{c}\neg\text{-I}, 1\ \dfrac{\dfrac{\wedge\text{-E}\ \dfrac{\Gamma, [(A \wedge B)]^1}{A} \qquad [\neg A]^2}{\bot}}{\neg(A \wedge B)}\\ \Pi\\ \bot\end{array}}{A} \qquad CRAA, 4\ \dfrac{\begin{array}{c}\neg\text{-I}, 3\ \dfrac{\dfrac{\wedge\text{-E}\ \dfrac{\Gamma, [(A \wedge B)]^3}{B} \qquad [\neg B]^4}{\bot}}{\neg(A \wedge B)}\\ \Pi\\ \bot\end{array}}{B}}{\begin{array}{c}A \wedge B\\ \vdots\end{array}}$$

A similar procedure is readily available for the remaining case in which the complex formula has the form $\neg(A \supset B)$.[12] An inspection of the rules in this system reveals that any obstacle to the subformula property would have to be either

1. a local peak involving a $\wedge$-maximum;

2. a local peak involving a $\supset$-maximum;

[12] For details see e.g. Prawitz (1965, p. 40).

3. a local peak involving a ¬-maximum;

4. or the result of a non-atomic application of *CRAA*.

Cases 1.–3. are taken care of by corresponding levelling procedures. 4. is adequately dealt with by the reduction procedure demonstrated above. It follows that S' is normalizable.

According to Dummett's claim, S' should also thereby be fully conservative. In particular, S' should be conservative over S. But we know that this is not so. For Peirce's Law is demonstrably not provable in S, though it can be proved in S':

$$
\begin{prooftree}
\AxiomC{$[A]^2$}
\AxiomC{$[\neg B]^1$}
\LeftLabel{\scriptsize $\wedge$-I}
\BinaryInfC{$A \wedge \neg B$}
\LeftLabel{\scriptsize $\wedge$-E}
\UnaryInfC{A}
\AxiomC{$[\neg A]^3$}
\LeftLabel{\scriptsize $\neg$-E}
\BinaryInfC{$\bot$}
\LeftLabel{\scriptsize $CRAA$, 1}
\UnaryInfC{B}
\LeftLabel{\scriptsize $\supset$-I, 2}
\UnaryInfC{$A \supset B$}
\AxiomC{$[(A \supset B) \supset A]^4$}
\LeftLabel{\scriptsize $\supset$-E}
\BinaryInfC{A}
\AxiomC{$[\neg A]^3$}
\LeftLabel{\scriptsize $\neg$-E}
\BinaryInfC{$\bot$}
\LeftLabel{\scriptsize $CRAA$, 3}
\UnaryInfC{A}
\LeftLabel{\scriptsize $\supset$-I, 4}
\UnaryInfC{$((A \supset B) \supset A) \supset A$}
\end{prooftree}
$$

This disproves Dummett's claim.

Or does it? Interestingly the introduction and immediate elimination of $A \wedge \neg B$ in the proof above generates a local peak that cannot be levelled. Perhaps, the system is not normalizable after all.

However, the local peak disappears if we recast our proof in a sequent format (of the same natural deduction system) with explicit structural rules. In such a setting we can simply bypass the apparent difficulty by an application of the weakening rule. We would replace the first three lines of the proof with the following:

$$
\begin{prooftree}
\AxiomC{$A : A$}
\LeftLabel{\scriptsize weakening}
\UnaryInfC{$A, \neg B : A$}
\AxiomC{$\neg A : \neg A$}
\LeftLabel{\scriptsize $\neg$-E}
\BinaryInfC{$A, \neg A, \neg B : \bot$}
\LeftLabel{\scriptsize $CRAA$}
\UnaryInfC{$A, \neg A : B$}
\noLine
\UnaryInfC{$\vdots$}
\end{prooftree}
$$

Alternatively, we could stick to our customary system and introduce the rule of *ex falso* (*EFQ*):

$$\text{EFQ}\ \frac{\bot}{A}$$

We would then need to show—as we did with *CRAA*—that any derivation in which *EFQ* is used to infer a complex formulas can be reduced to a derivation in which all applications lead to atomic conclusions only. But this is a routine procedure (see e.g. (van Dalen 1997, p. 207)). The requisite reduction procedures are comparable in style and complexity to the examples we have considered above in the case of *CRAA*. Thus the system S' does indeed turn out to be a normalizable system and yet a non-conservative extension of S—hence a counterexample to Dummett's claim.

The reason Dummett's seemingly plausible argument fails is because it does not take into account the classical *reductio* rule. It is this rule that smuggles in non-conservativeness by allowing us to eliminate occurrences of the negation operator. At the same time it stays under the radar of normalizability because it induces neither local peaks nor plateaux.

We have thus established that normalizability is not sufficient for full conservativeness, and so refuted Dummett's claim. Have we also disproved Dummett's conjecture that *stability* is a sufficient condition for conservativeness (and hence total harmony)? Everything depends on whether the rules in S' qualify as stable. However, pending a precise notion of stability, we cannot pronounce on this question. While we may expect that any notion of stability will count in the rules governing $\wedge$ and $\supset$, it is bound to be a matter of controversy (to say the least) whether the strictly classical rules (*CRAA*, double negation elimination, etc.) will make the cut. Therefore, Dummett's conjecture that ubiquitous stability (the strengthened version of intrinsic harmony) implies systematic conservativeness (total harmony) is not necessarily imperilled.

4.7 Does stability entail total harmony?

What, then, are we to make of Dummett's conjecture to the effect that stability—when prevailing throughout—implies total harmony? Both Prawitz (1994, p. 374) and Read (2000, p. 127) attack Dummett on this point, arguing that a ready-made counterexample can be found in Jeffrey Ketland's and Stewart Shapiro's well-known demonstration that adding the Tarskian truth-theory to Peano arithmetic (**PA**) (and allowing the truth predicate to occur in instances of the induction schema) yields a non-conservative extension.[13] Take the following form of the T-schema given by the following straightforward rules:[14]

$$Tr\text{-I}\ \frac{A}{Tr(\ulcorner A \urcorner)}$$

$$Tr\text{-E}\ \frac{Tr(\ulcorner A \urcorner)}{A}$$

Prawitz and Read assume that no matter what the exact details of our account of stability are, if any rules are to count as stable, then the rules for the truth predicate (as just presented) are. Hence, if the truth predicate can be found to induce non-conservativeness when adjoined to a good-functioning system, Dummett's conjecture will be refuted.

A first noteworthy point is that Prawitz and Read fail to appreciate the distinction between *theoretic* conservativeness which fails in the case of the truth-theoretic extension of **PA**, and *systematic* conservativeness, which is what is at issue in Dummett's conjecture. The non-conservativeness observed by Ketland and Shapiro, as it presumably lies outside of the scope of logic, is of the former kind. Consequently, the non-conservativeness result in question cannot serve as a counterexample to Dummett's conjecture, which is concerned solely with total harmony (i.e. systematic conservativeness). That being said, it would nevertheless be devastating for Dummett's conjecture if a stable inference rule, when added to an orderly theory, were to perturb non-logical regions of the language (even if it did not affect the logical fragment), for this would constitute a violation of the principle of innocence.

However, a closer look reveals that Prawitz's and Read's counterexample has no teeth.[15] The truth predicate (along with the inference rules it obeys) does not *in itself* produce a non-conservative extension of the language of arithmetic. Adding the truth predicate to the language of arithmetic with the above inference rules results in a *conservative* extension of **PA**, *even if* we allow the truth predicate to occur in instances of the induction schema (see Ketland, op. cit., p. 76, Halbach

[13]See e.g. Ketland 1999 and Shapiro 1998b for details.

[14]I take it that appropriate precautionary measures have been taken to ward off paradox.

[15]It should be noted that both authors also suggest that 'higher-order concepts' (Read 2000, p. 127) might be a possible source of non-conservativeness. The argument to follow, insofar as it does not deal with this possibility, may therefore only supply a partial defence of Dummett's conjecture. My hunch is that the case of higher-order quantification can be dealt with along similar lines as the truth predicate.

2005). It is only when we add the entire Tarskian truth-theory with its compositional axioms (*and* when we allow the truth-predicate to occur in instances of the induction schema) that we obtain a non-conservative extension of **PA**. The reason, roughly, is because the Tarskian theory enables us to prove that all of **PA**'s rules of inference are truth-preserving and hence that all of **PA**'s theorems are true, thereby proving the consistency of **PA**. Of course, by Gödel's second incompleteness theorem, **PA** cannot, on pain of inconsistency, prove its own consistency. This shows that **PA** augmented by the Tarskian truth theory is not conservative over **PA**. The T-schema inference rule when added on its own—and this is all that is at issue here—does not generate non-conservativeness. The result, interesting though it may be in its own right, thus bears no relation to the question of harmony and its relation to conservativeness. The purported 'counterexample' poses no threat to Dummett's conjecture.

How does Dummett's conjecture fare within the narrower confines of the realm of logic? Again, a superficial look may suggest that it does not fare very well: in the case of first-order logic the addition of the truth predicate again appears to give rise to a non-conservative extension, as can be seen by the following reasoning.[16] Let A be a logical truth and B a contradiction.

1. $\vdash Tr(\ulcorner A \urcorner) \leftrightarrow A$ (T-schema)

2. $\vdash A$ (logical truth)

3. $\vdash Tr(\ulcorner A \urcorner)$ (from 1. and 2.)

4. $\vdash Tr(\ulcorner B \urcorner) \leftrightarrow B$ (T-schema)

5. $\vdash \neg B$ (logical truth)

6. $\vdash \neg Tr(\ulcorner B \urcorner)$ (from 4. and 5.)

7. $\vdash Tr(\ulcorner A \urcorner) \leftrightarrow \neg Tr(\ulcorner B \urcorner)$ (from 3. and 6.)

8. $\vdash \forall x \forall y (x = y \supset Tr(x) \leftrightarrow Tr(y))$ (logical truth)

9. $\vdash \forall x \forall y (\neg (Tr(x) \leftrightarrow Tr(y)) \supset x \neq y)$ (contraposition)

[16] The proof holds for intuitionistic and classical logic alike. Since the proof is somewhat unwieldy in Gentzen-Prawitz-style format, we here content ourselves with a slightly abridged linear version.

10. $\vdash \neg(Tr(\ulcorner A \urcorner) \leftrightarrow Tr(\ulcorner B \urcorner))$ (from 7.)

11. $\vdash \ulcorner A \urcorner \neq \ulcorner B \urcorner$ (from 9. and 10.)

12. $\vdash \exists x \exists y\ x \neq y$ (from 11.)

We just proved—with the help of the truth predicate—that there are at least two things; and this result, we may assume, is not a logical truth.[17] But this result should not surprise us. After all, the truth predicate presupposes the existence of resources in our language that allow us to name all the sentences to which it may be applied. It would make little sense to introduce '$Tr(x)$' into our language in the absence of a term-forming operator (e.g. a quotation operator). As before, in the case of the adjunction of the Tarskian theory of truth to **PA**, the fault does not lie with the inferential behaviour of the truth predicate itself but with the additional apparatus required to put it to work. In a sense, the presence of the term-forming quotation operator denoted by '$\ulcorner \urcorner$' is in itself sufficient to generate non-conservativeness.[18]

1. $\vdash A$ (logical truth)

2. $\vdash \neg B$ (logical truth)

3. $\vdash \ulcorner A \urcorner \neq \ulcorner B \urcorner$ (from 1. and 2. by injectivity of $\ulcorner \urcorner$)

4. $\vdash \exists x \exists y\ x \neq y$ (from 3.)

Here we have simply imported lines 2., 5., 11., and 12. from the previous proof and derived the same conclusion. Evidently 'truth' has had no say in this.

We have to acknowledge, however, that the second 'proof' is a proof in the metatheory only. The reason is that we have no standard way of expressing the injectivity of the term-forming operator in the object language. The second proof tacitly relies on the fact that two formulas that are not logically equivalent cannot be the same formula, i.e. that they are constituted by different strings of symbols and therefore are assigned different names. We might try to approximate this thought as follows

[17] This is just an instance of the commonplace view that logic should not have existential import. While non-free logic does commit us to the existence of something it should not commit us to the existence of any minimal number of things, nor to the existence of an indeterminate number of things.

[18] I am here making the straightforward assumption that the term-forming operator in question is injective: i.e. if $A \neq B$, then $\ulcorner A \urcorner \neq \ulcorner B \urcorner$.

1. $\vdash \neg(A \leftrightarrow B) \supset A \neq B$ (postulate)

2. $\vdash A \neq B \supset \ulcorner A \urcorner \neq \ulcorner B \urcorner$ (by 1. and injectivity)

3. $\vdash \neg(A \leftrightarrow B) \supset \ulcorner A \urcorner \neq \ulcorner B \urcorner$ (from 1. and 2.)

4. $\vdash \neg(A \leftrightarrow B)$ (by assumption)

5. $\vdash \ulcorner A \urcorner \neq \ulcorner B \urcorner$ (from 3. and 4.)

6. $\vdash \exists x \exists y\ x \neq y$ (from 5.)

The problem is apparent already in our postulate in line 1. The consequent of the conditional is ill-formed because '=' may only be correctly inserted between names, not between formulas. (The same is obviously true for all subsequent occurrences.)

But though this piece of reasoning is intuitively obvious—we do not hesitate for a moment to endorse it in the meta-theory—we have no way of reproducing it in the object language. So we cannot prove within our system that the term-forming operator alone is the culprit. Does this mean that our diagnosis of the sources of non-conservativeness is mistaken? I do not think so. The fact that we need to resort to the metalanguage in order to express the fact that if A and B are not equivalent, they must be distinct does not make it any less obvious that it is the presence of the term-forming operator, not the truth predicate that induces non-conservativeness. This is clearly brought out by the fact that we can re-run the first proof above, replacing the truth predicate with any other one-place predicate P for which we can prove $\vdash \neg(P(\ulcorner A \urcorner) \leftrightarrow P(\ulcorner B \urcorner))$. Hence, virtually any number of predicates can be a catalyst for non-conservativeness: nothing hinges on the specific deductive properties of the truth predicate.

Even more fundamentally, because the truth predicate only has a place in a system that contains a name-forming device, there can be no question of introducing it into a system of pure logic. It can only ever sensibly be introduced into a system that has already been adulterated by the presence of a name-forming operator. Recall that Dummett's conjecture states that intrinsic harmony implies total harmony in a 'context where stability prevails' (Dummett 1991, p. 250). This should be read as, 'Take a system in which all of the rules of inference are stable. Adding a further stable inference rule to such a system will result in a systematically conservative extension'.[19] But in the case of the truth predicate the former condition can never be satisfied: the base system must contain a term-forming operator; it is therefore not composed solely of well-behaved logical expressions—unless, that is, a suitable quotation operator could be articulated and shown to be governed by stable inference rules.

Is such a stable quotation operator likely to be found? It is hard to see what introduction and elimination rules for such an operator would look like, let alone harmoniously matched ones. One place to start looking would be to inquire what principles of inference would have to be in place for the previous 'proof' of $\ulcorner\urcorner$-induced non-conservativeness to be an acceptable proof in the object language. The following rule of inference captures the previously unavailable meta-theoretical information and brings it within the reach of the object language

$$\frac{\ulcorner A \urcorner = \ulcorner B \urcorner}{A \leftrightarrow B}$$

It is not hard to see how, with the help of this principle, we obtain a proper proof essentially along the same lines as the previous putative proof.[20] However, it is clear that this rule of inference qualifies neither as an introduction rule, nor as an elimination rule; it conflicts with the assumption of the two-aspect model of meaning.[21]

More problematically still, it is difficult to see how one could even begin to go about stating inference rules that would capture the meaning of this quotation operator (let alone harmoniously matched ones). In his (2004b) Tennant proposes a general inferentialist account of abstraction operators, showing how such expressions can be viewed as obeying introduction and elimination rules. But abstraction operators of the kind Tennant considers differ from quotation operators in that they essentially involve a relation. In the case of quotation operators the only plausible relation intrinsic to its meaning is 'being the same string of symbols', which

[19]This is also what Dummett must have had in mind, given that he advances his conjecture in the context of the Q-example (1991, p. 290).

[20]Even someone like Tennant, who rejects the 'dogma' that logic must be free of existential commitments (since logic does, according to him, commit us to necessary existents) and so is not moved by our proof, will be swayed by the fact of non-conservativeness. His solution, as someone who equates logicality with being governed by a 'harmoniously balanced pair of introduction and elimination rules' (1997, p. 296, fn. 24), is to deny the truth predicate the status of a logical constant (ibid. p. 294). Given Tennant's criterion for logicality, our analysis thus concurs with his.

[21]On this count it fares no better than the classical *reductio* rule, which cannot obviously be assigned to either of those categories either.

is meta-theoretical in nature and so again not expressible in the object language. Only if the quotation operator can be domesticated (perhaps in a fashion analogous to Tennant's treatment of abstraction operators) would the non-conservativeness of the system augmented by the truth predicate plus quotation operator pose a threat to Dummett's conjecture.

But is this enough to exculpate the truth predicate and thus to vindicate Dummett's conjecture? Not yet. For: without a term-forming operator no truth predicate, and with a term-forming operator no conservativeness. There thus remains a worrisomely intimate connection between the truth predicate and conservativeness, since we cannot have the truth predicate without a name-forming operator. But on the other hand, the rules for the truth predicate as we have presented them certainly appear to be stable. Something remains to be explained.

The rules we gave are indeed stable, but they do not fully capture the meaning of the truth predicate. What these rules describe is the inferential behaviour of a *sentential operator* t (as opposed to the predicate T that attaches to terms):

$$t\text{-I}\ \frac{A}{t(A)}$$

$$t\text{-E}\ \frac{t(A)}{A}$$

Clearly *this* operator cannot cause non-conservativeness, nor can anyone who espouses the view that logicality is proof-theoretically determined challenge t's claim to being a constant (though we may be excused for omitting it on account of its dullness). But clearly t is not the truth predicate, since the truth predicate *essentially* involves a name-forming operator. Therefore, Shapiro is wrong when he claims that inferentialists—in what he calls the 'Dummett-Tennant-Hacking view'—are necessarily committed to the logicality of the truth predicate (Shapiro 1998a, p. 617). Our discussion has shown that even if we choose to define logicality in terms of expressibility by means of harmonious inference rules, the truth predicate does not qualify as logical.

Dummett's conjecture is thus not affected by these alleged 'counterexamples'. Clearly, however, we have not ruled out the existence of genuine counterexamples. Nevertheless, we have shown Prawitz's and Read's original misgivings to be unwarranted at least inasmuch as the truth predicate is concerned, and our best attempt at constructing a counterexample by improving on their model has also missed its mark.

4.8 Summary

Let us briefly recapitulate the conclusions we have reached in this chapter. We began by introducing Dummett's distinction between total harmony, which corresponds to systematic conservativeness, and intrinsic harmony, which Dummett equates initially with the levelling of local peaks and later with the notion of stability. Having clarified the notions of levelling procedures and normalizability, we went on to consider the Q-example, which demonstrates the insufficiency of the notion of intrinsic harmony. We pointed out Dummett's mistake of identifying harmony for the logical constants with the global property of total harmony rather than with the much more suitable principle of stability. In the last part of the chapter we examined the relations between conservativeness and normalizability and the respective relations these bear to harmony. We disproved Dummett's claim to the effect that a system comprising only stable inference rules is normalizable just in case it is conservative. Finally, we defended Dummett's conjecture that full conservativeness (and hence conservativeness *tout court*) is a necessary condition for stability.

In the next chapter, we turn to an idea that was implicit in all our talk of harmony, namely, the notion that there is a functional relationship between V-principles and P-principles, i.e. between introduction rules and elimination rules. It is the tenability of this underlying assumption that occupies us in the next chapter.

Chapter 5

Interlude: The principle of functionality

5.1 Introducing the principle of functionality

Let us return once again to the Q-example. We have seen that Dummett's notion of intrinsic harmony, understood in terms of the existence of a levelling procedure, is insufficient both for normalization and for conservativeness. Can we get a better handle on what it is that stands in the way of obtaining these entailments in cases like the system containing the quantum-logical or-elimination rule? In other words, can we come up with additional constraints which, perhaps jointly with the requirement of intrinsic harmony, would constitute a well-motivated constraint sufficient for conservativeness and normalizability? How might we go about devising such an improved notion of intrinsic harmony, which—following Dummett again—we can call stability?

Well, we could begin by checking the rules of inference in the Q-system for possible defects. Naturally, the co-occurrence of two distinct disjunction operators involving the same introduction rules but different elimination rules (the restricted elimination rule for quantum-or and the unrestricted elimination rule for standard disjunction) immediately arouses our suspicion. It is this feature—that the same set of introduction rules is coupled with two distinct, non-equivalent elimination rules—that seems to be the root of all evil in our example. We exploited precisely this feature to demonstrate the failure both of normalization and of conservativeness. The coexistence of the two competing disjunction operators violates what we may call the *principle of functionality*. It is first stated by Gentzen in the following form:

> it should be possible to display the E-inferences [i.e. elimination rules] as unique functions of their corresponding I-inferences [i.e. introduction rules] on the basis of certain requirements (Gentzen 1969a, p. 81).

According to Gentzen and many of his followers, a scenario like the above should be disallowed. Introduction rules should uniquely determine harmoniously matched elimination rules. Supposing this is correct, one of the two elimination rules for the disjunction-introduction rules has to go. But which one? Given that both rules are intrinsically harmonious we may suspect that there is no danger of the $\vee$-elimination rule being too strong. And this would be enough to disqualify the restricted rule. The availability of a stronger yet intrinsically harmonious elimination rule entails that the weaker rule is too weak: knowing that the stronger rule does not say too much, we can conclude that the restricted rule does not allow us to say enough. In other words, the restricted rule generates P-weak disharmony. Of course this does not show that the $\vee$-elimination rule does not suffer the same defect. For all we know there might be an even stronger, intrinsically harmonious elimination rule to be had. Short of possessing a reliable criterion that assesses whether some set of rules is stable, we shall be in no position to rule out the existence of such a stronger rule. A notion of stability will have to provide precisely such a criterion. We will be concerned with stability in the following chapters. But before we move on it is necessary to dwell on the notion of functionality a while longer.

It is worth noting that Gentzen, like many of his successors, considered introduction rules to be constitutive of the meanings of the constants whose deductive behaviour is regulated by them. This bias in favour of introduction rules can be seen in Gentzen's formulation of the principle of functionality quoted above: the function Gentzen describes takes us from (permissible) introduction rules that invest the constant with meaning to the adequate elimination rules whose role is to state the inferential consequences of those meanings and so to reverse their effect. So far, however, we have no reason to accord meaning-theoretic primacy to introduction rules. Why should not elimination rules in some or all cases be at least partially determinative of the meanings of the logical operators? If we do, contrary to Gentzen, grant elimination rules a meaning-determinative role at least in some cases, then, given the aforementioned affiliations between an 'introductions first'-view and a verificationist view of meaning, on the one hand, and an 'eliminations first'-view and a pragmatist view of meaning, on the other hand, we could find ourselves espousing a hybrid approach towards the meanings of the logical operators—verificationist about some expressions, pragmatist about others. Alternatively, it might turn out that the two theories coincide, at least for the domain of the logical particles.[1] Either way, so far we have encountered nothing that should dissuade us from thinking that elimination rules too might at least in some cases give the meanings of the constants they govern.

Moreover, there are arguably reasons for privileging eliminations in certain circumstances. It has been argued that in the case of the conditional (Rumfitt 2000, p. 790), the universal quantifier (Dummett 1991, p. 275) and the disjunction operator (Tennant 1987, p. 90), it is in some sense more natural to treat elimination rules as the primary determinants of meaning. Flexibility on this point is even more essential if one hopes to extend an inferentialist account to the modal operators. The standard systems of modal logic distinguish themselves by the strength of their necessity and/or possibility operators. But the strength of these operators is in turn usually fixed by the restrictions that weigh on their introduction rules (in the case of $\Box$) or elimination rules (in the case of $\Diamond$). Thus in standard natural deduction representations of modal logic, different possibility operators, share for example, the same introduction rule and differ only in their elimination rules. An introductions-first approach would therefore be incapable of distinguishing between, say, the **S4**- and the **S5**-strong $\Diamond$ operator. Conversely, an eliminations-first approach would be impotent to account for the meaning of $\Box$.

On the basis of these considerations we may advance the following neutral version of the functionality principle:

> **Principle of functionality**: for any permissible set of introduction/elimination rules for a constant c there can only be one set of elimination/introduction rules for c that harmoniously matches it.

[1] The account we advance here does indeed hold that the verificationist and the pragmatist approaches coincide as we will see in section 5.3.

5.2 Some critical considerations concerning functionality

Recall that what motivated our functionality principle was the Q-example. It was because both $\vee$ and $\sqcup$ shared the same introduction rules that the latter collapsed into the former. We may ask whether this is an instance of a more general phenomenon. In other words, is it always the case that failure of functionality results in non-conservativeness and/or non-normalizability? It may seem as if the trick that generated the non-conservativeness and failure of normalizability in the Q-example can be readily replicated in analogous contexts. The method that suggests itself is this: take two sets of elimination rules that are each intrinsically harmonious with respect to the same set of introduction rules, but where at least one of the two sets of elimination rules induces P-weak disharmony. Adding the connective governed by the stronger elimination rule to a system containing the connective governed by the weaker elimination rule will then result in the collapse of the latter into the former.

So is this a general recipe for generating non-normalizable and/or non-conservative systems? Not so, at least if we allow for modal operators. For consider an example featuring the modal operator $\Diamond$ whose meaning is given by the following rules:

$$\Diamond\text{-I}\ \frac{\begin{matrix}\Gamma\\ \vdots\\ A\end{matrix}}{\Diamond A}$$

$$\Diamond\text{-E},\ i\ \frac{\begin{matrix}\Gamma\\ \vdots\\ \Diamond A\end{matrix}\qquad \begin{matrix}\Gamma', [A]^i\\ \vdots\\ C\end{matrix}}{C}$$

Crucially certain restrictions apply to the elimination rule: in the minor premise all the hypotheses on which C depends apart from A have to be *modal* (i.e. be of the form $\Box B$, $\neg\Diamond B$ or $\bot$), and C itself has to be *co-modal* (i.e. be of the form $\Diamond B$, $\neg\Box B$ or $\neg\bot$). With these restrictions in place we obtain the possibility operator for **S4**. These rules can be relaxed to yield the corresponding **S5** operator by counting $\neg\Box B$ as modal (and so conflating the modal/co-modal distinction). In accordance with our putative recipe for constructing non-normalizable, non-conservative systems, we imagine a system S containing only the **S4**-strong possibility operator (which we will denote by '$\Diamond$' as before) and the operators and rules for the propositional calculus. The next move should consist in the creation of S' by introducing the **S5**-strong possibility operator, which we will denote '$\blacklozenge$'. *Prima facie* it again seems that the weaker operator will be absorbed into the stronger one, for again it seems that we can derive the stronger operator from the weaker

$$\Diamond\text{-E}, i\ \dfrac{\begin{matrix}\vdots\\ \Diamond A\end{matrix} \qquad \blacklozenge\text{-I}\ \dfrac{[A]^i}{\blacklozenge A}}{\blacklozenge A}$$

Similarly, it again looks as if the system obtained by adding $\blacklozenge$ is not normalizable as the following derivation shows

$$\blacklozenge\text{-E}, 2\ \dfrac{\Diamond\text{-E}, 1\ \dfrac{\begin{matrix}\Gamma_0\\ \vdots\\ \Diamond A\end{matrix} \qquad \blacklozenge\text{-I}\ \dfrac{[A]^1}{\blacklozenge A}}{\blacklozenge A} \qquad \begin{matrix}\Gamma_1, [A]^2\\ \vdots\\ C\end{matrix}}{C}$$

Any attempt at removing the plateau created by the introduction and elimination of $\blacklozenge$ by applying our usual reduction procedures (i.e. in this case permutative reductions) is, again, doomed to fail.

$$\Diamond\text{-E}, 2\ \dfrac{\begin{matrix}\Gamma_0\\ \vdots\\ \Diamond A\end{matrix} \qquad \blacklozenge\text{-E}, 1\ \dfrac{\blacklozenge\text{-I}\ \dfrac{[A]^2}{\blacklozenge A} \qquad \begin{matrix}\Gamma_1, [A]^1\\ \vdots\\ C\end{matrix}}{C}}{C}$$

For although we have created a local peak—$\blacklozenge A$ has been turned into a maximum formula—there is no guarantee that the final application of the $\Diamond$-elimination rule is correct; we have no guarantee that Γ_1 does not contain any formulas of the form '$\neg\Box B$' (i.e. formulas that are, as we have seen, countenanced as modal by the **S5**-strong $\blacklozenge$, but not by its weaker **S4** counterpart $\Diamond$).

However, everything we have said so far hinges on whether the initial application of the $\Diamond$-elimination rule is correct:

$$\Diamond\text{-E}, i\ \dfrac{\begin{matrix}\vdots\\ \Diamond A\end{matrix} \qquad \blacklozenge\text{-I}\ \dfrac{[A]^i}{\blacklozenge A}}{\blacklozenge A}$$

Recall that one of the restrictions placed on the $\Diamond$-elimination rule is that it can only operate on premises that are co-modal. The question then is whether we are to count formulas of the form '$\blacklozenge A$' as co-modal. This is certainly not licensed by standard specifications of the $\Diamond$-elimination rule. While it may with some plausibility be argued that there are good grounds for classifying well-formed formulas prefixed by $\blacklozenge$ as co-modal, there can be no doubt that this involves a modification of the original rule. The breach of the principle of functionality therefore does not seem to have the same deleterious effect in this particular case.

Where does this leave us vis-à-vis our principle of functionality? What the above example shows is that a lack of functionality is not a sufficient condition for a lack of conservativeness and normalizability. In other words, functionality is not a necessary condition for either conservativeness or normalizability. And, one might add, this is not so surprising. After all, what is so aberrant about the idea that two logical expressions should be introducible under the same circumstances, but that they should have distinct meanings on account of their respective (distinct) elimination rules? Moreover, as we suggested above, the principle of functionality and the modal operators cannot both be comfortably accommodated in a single account of harmony. More precisely, if we wish to account for modal operators and yet hold on to functionality, our choice as to which rules—introduction rules or elimination rules—have meaning-theoretic priority is strongly constrained. In order to be able to make sense of modal operators proof-theoretically we must allow that both introduction and elimination rules can be constitutive of meaning. In order to distinguish necessity operators of varying strength, for example, we must assume that the introduction rules are meaning-determinative; in order to distinguish among possibility operators we must take the elimination rule to be meaning-determinative. These considerations do not necessarily count against functionality. They may, however, be sufficient to deter someone who believes that modal operators ought to have a place in an account of harmony and who simultaneously has verificationist or pragmatist commitments.

5.3 Functionality defended

How seriously are we to take the doubts concerning the principle of functionality voiced in the previous section? What grounds do we have for upholding functionality apart from the fact that it rules out some pathological cases, such as that of the system containing the two disjunction operators? The first point to make is that the principle of functionality is not simply an *ad hoc* device tailored for such purposes. It stems, ultimately, from the principle of innocence (see section 3.3), which states that logic may neither add nor subtract information. An elimination rule should therefore be able to 'undo' the effect of introducing a logical constant. What the above example involving the quantum-logical disjunction operator shows is that this undoing must be exactly apportioned. The elimination rule should not license any new inferences that were not permissible on the grounds upon which the constant in question was introduced in the first place; that is, P-strong disharmony should be prohibited. But nor should the elimination rule usurp information by disallowing inferences that we were in a position to make prior to the introduction of the constant; i.e. P-weak disharmony is equally to be avoided. Clearly, either of these two constraints taken on its own may be satisfied by a range of elimination rules, but there can only be one rule satisfying both constraints. This, at least, is the thought underlying the principle of functionality: once we have fixed the conditions under which the introduction of a given logical operator is permissible, the two constraints—guarding against P-weak and P-strong disharmony—jointly determine a unique choice.

Here we are presupposing that elimination rules distinguish themselves solely by the deductive consequences they enable us to draw. Accordingly, we may, at least for present purposes, treat two elimination rules as equivalent if, given the same premises, they give rise to the same set of consequences. In the light of our foregoing considerations concerning the notion of disharmony, we would do well to give a somewhat more precise characterization of the notion of relative strength of elimination rules. This is easily done by defining an equivalence relation between elimination rules, as the relation which obtains between two rules if they give rise to the same set of consequences. We can then define the notion of strength in terms of an ordering on equivalence classes of elimination rules. Let $\mathsf{C}(\$\text{-E})$ be the set of statements that may be inferred directly or indirectly from the elimination rule \$-E from a given set of undischarged premises Γ.[2] The equivalence relation R can now be expressed as follows: we have $R(\$\text{-E}, \$\text{-E}')$ if and only if, for any set of premises

[2]The importance of taking a rule's *indirect* consequences into account is illustrated by the following alternative $\wedge$-elimination rule:

Γ, $\mathsf{C}(\$\text{-E}) = \mathsf{C}(\$\text{-E}')$. We can then define the partial ordering $\prec$ holding between the equivalence classes B and A if and only if A is stronger than B in the sense that for every member \$-E of A and every member \$-E′ of B, $\mathsf{C}(\$\text{-E}') \subsetneq \mathsf{C}(\$\text{-E})$. That is, if the set of formulas that members of B allow us to infer is a proper subset of the set generated by the members of A. In practice we will speak loosely of the respective strength of *rules* rather than of equivalence classes.

We are assuming that structural assumptions and the other rules are held fixed across different systems in each such comparison. This is crucial because the relative strength of two elimination rules can be altered by altering structural assumptions. Take a standard system without *EFQ*. (Recall that we are treating *EFQ* as a structural rule.) Now consider Tennant's $\vee$-elimination rule for his systems **IR** and **CR**:

$$\vee_T\text{-E}, i\ \dfrac{A \vee B \qquad \begin{matrix}\Gamma_1, [A]^i \\ \vdots \\ C/\bot\end{matrix} \qquad \begin{matrix}\Gamma_2, [B]^i \\ \vdots \\ C/\bot\end{matrix}}{C/\bot}$$

Tennant explains:

> In the statement of $\vee E$ [$\vee_T$-E in our notation] the slash notation $C/\bot$ is to be understood as follows: we allow a subordinate conclusion of either one of the cases to be brought down as main conclusion if the other subordinate conclusion is $\bot$. Of course, if both subordinate conclusions are of the same form, the main conclusion has the same form (Tennant 1987, p. 258).

Tennant's rule, by design, allows us to derive disjunctive syllogism,

$$\vee_T\text{-E}, 1\ \dfrac{A \vee B \qquad \dfrac{[A]^1 \qquad \neg A}{\bot} \qquad [B]^1}{B}$$

$$\wedge\text{-GE}, i\ \dfrac{\begin{matrix}\Gamma \\ \vdots \\ A \wedge B\end{matrix} \qquad \begin{matrix}\Gamma', [A, B]^i \\ \vdots \\ C\end{matrix}}{C}$$

$\wedge$-GE is provably equivalent to our standard $\wedge$-elimination rule in the sense that replacing one for the other in any system does not affect that system's output. Yet the set of $\wedge$-GE's direct consequences is clearly more extensive because it licenses inferences not just to A and B but also to any of their consequences. All of $\wedge$-E's indirect consequences are direct consequences of $\wedge$-GE.

This is what distinguishes Tennant's systems from traditional Anderson and Belnap-type systems of relevant logic. However, in the absence of the *ex falso* rule, the principle of disjunctive syllogism is not derivable. Hence, under these structural assumptions $\vee_T$-E is stronger than the usual $\vee$-E rule.

The notion of the strength of an elimination rule having been thus clarified, we can now see that the principle of functionality requires only that an equivalence class of elimination rules is uniquely determined, rather than an individual rule. Viewed in this way, the functionality claim can be seen to be a direct consequence of the notion of general harmony we started off with. For suppose we did allow for cases in which two elimination rules, \$-E and \$-E$'$, say, of unequal strength (but comparable with respect to $\prec$) are associated with the same set of introduction rules. Necessarily one of the two rules will have a statement A as a consequence that is not a consequence of the other rule. Suppose $A \in \mathsf{C}(\text{\$-E})$ but $A \notin \mathsf{C}(\text{\$-E}')$ (hence $[\text{\$-E}']_R \prec [\text{\$-E}]_R$). If the inference to A is legitimate already on the basis of the grounds for introducing \$, then \$-E$'$ deprives us of information and so introduces P-weak disharmony. If, on the other hand, A is not already available at that point, \$-E has been shown to be too strong and hence leads to P-strong disharmony.

Therefore, if harmony is to be understood, as we have been assuming throughout, as a kind of balance or equilibrium between the grounds for introducing a constant and the consequences of having done so, we are inevitably wedded to the principle of functionality. Of course, since we wish to remain neutral with respect to the question of the meaning-theoretic priority of introduction versus elimination rules, we have to endorse an analogous story taking elimination rules as our point of departure. The principle of functionality then similarly guides our choice of sets of introduction rules. How does our argument from the previous paragraph carry over to this case? We can run an analogous argument to the one just presented. This time we suppose that, for a given set of elimination rules \$-E, there are several harmoniously matched introduction rules and we show how all but one of those sets of rules must in fact be in disharmony with \$-E. Again we are strictly speaking about equivalence classes of rules, where two sets of introduction rules \$-I and \$-I$'$ are in the same equivalence class if and only if the set of deductive consequences of their respective premises is the same. We can order these equivalence classes in accordance with respect to the subset relation between the sets of deductive consequences of the premises. Where there are several rules, as for instance in the case of or-introduction, we take the intersection of the sets of consequences of those premises. So in the case of disjunction with the familiar introductions rules $A \vdash A \vee B$ and $B \vdash A \vee B$, we have $\mathsf{C}(\vee\text{-I}) = \mathsf{C}(A) \cap \mathsf{C}(B)$. Where there are several premises per rule, as in the case of conjunction, we take the union of the sets of consequences of those premises. In the case of conjunction with the introduction rule $A, B \vdash A \wedge B$, we get $\mathsf{C}(\wedge\text{-E}) = \mathsf{C}(A) \cup \mathsf{C}(B)$. In this way we can, for any permissible set of elimination rules, uniquely determine an equivalence class of harmoniously balanced introduction rules. This guarantees that for every set of permissible introduction rules there will be a unique set of elimination rules that harmoniously matches it, and that for every set of permissible elimination rules, there will be a unique set of introduction rules that harmoniously matches it, exactly as required by the principle of functionality. Stability obtains if any set of rules is matched with its harmonious counterpart, and that, starting with the said counterpart, we again end up with the same initial set of rules.

So far, so good. We seem to have made some headway in articulating a more rigorous version of the pre-theoretic notion of general harmony we started off with. Before we continue on with our attempt to develop a suitable principle of stability, we must return to some issues surrounding the modal operators. Do such operators have a place in logical inferentialism? In particular can they be accommodated within our burgeoning account of harmony?

5.4 A note on modalities

As the last sections have shown, modal operators fit rather uneasily with the idea that the relationship between introduction rules and elimination rules ought to be a functional one. Standard formulations of natural deduction rules for modal operators rely on various restrictions—on introduction rules in the case of the $\Box$-operator and restrictions on elimination rules in the case of the $\Diamond$-operator—in order to achieve the fine-tuning required to express the distinct operators of differing strength. Consequently, as we noted above, the necessity operators of the standard systems **T**, **S4** and **S5** differ only in the restrictions they impose on the hypotheses on which the '$\Box$'-formula to be introduced may depend.[3] The following general form of introduc-

[3]The operators in **K**, **K4**, **KB** are even stranger from the point of view of harmony, as they have no $\Box$-elimination rule at all.

tion and elimination rules is shared by all systems:

$$\Box\text{-I}\ \frac{A}{\Box A}$$

$$\Box\text{-E}\ \frac{\Box A}{A}$$

We obtain **S4**-strength necessity if we require that all hypotheses on which $\Box A$ depends are modal (i.e. of the form $\Box B$, $\neg\Diamond B$, $\bot$). We can then weaken the system by adding further restrictions to get **T**, or strengthen the system to **S5** by relaxing the restrictions in certain ways, and so on. The same holds for $\Diamond$ and its elimination rule. But now we have a number of distinct $\Diamond$-operators that share the same introduction rule but have different elimination rules; and a number of $\Box$-operators that have different introduction rules! In the former cases we have violations of injectivity, in the latter cases the operators flout the requirement of functionality.[4]

As we mentioned earlier, the view that functionality and injectivity are suitable constraints for pinning down the notion of harmony for the logical constants is not as such incompatible with the position that modal operators ought to be accounted for in a theory of harmony. However, even if it were the case that one of the several available pairs of inference rules for necessity, say, satisfied the constraints of functionality and injectivity (and thereby necessarily satisfied these constraints uniquely), it may appear bizarre that the properties of the notion of necessity should be determined by logic alone. Why, for example, should the question of the axiom commonly referred to as (4)—that $\Box A$ entails $\Box\Box A$—holds be a question for logic to answer rather than metaphysics or conceptual analysis? If (4) holds, then any divine agency would be conscripted by the laws of logic; if (4) does not hold, then a divine agent would be free to choose the laws of logic. Are such apparently theological debates really just matters for logic to settle? Let us take another example. Suppose that **S4** turns out to be stable. It would follow that the logical positivist credo—that all necessity stems from linguistic convention—is false, since it is incompatible with (4). Many would of course agree that the logical positivist's story ought to be rejected, but it seems odd that logic should decree this.

One way to get around this difficulty is to distinguish clearly between *logical* necessity and other types of necessity. Our constraint of harmony could then be thought to single out a particular system of modal logic as purely logical, without

[4]Note that privileging elimination rules over introduction rules would not change much.

thereby discrediting other systems. Such systems could still be serviceable *qua* mathematical theories or *qua* applied logics fulfilling certain descriptive functions, as in the case of the ever-increasing numbers of systems of temporal, epistemic, doxastic and deontic logic. As such these systems will remain worthy of study (though the title 'logic' would be misplaced). We would find ourselves with a pair of purely logical notions of necessity and possibility whose meanings would be given entirely in terms of harmonious inference rules. They would constitute the one modal *logic*, which could easily co-exist with several *theories* of modality. The metaphysician would be under no obligation to devise theories that conform to the logical notion of necessity. He could happily claim that this or that other theory of modality better describes how things stand in the metaphysical realm. Conversely, an account of harmony that deals with modality would have nothing to do with physical or metaphysical notions of necessity. It would be concerned solely with the notion of logical necessity. Ideally, of course, the principle of harmony would pick out modal notions that mesh with our established ideas concerning the properties of logical necessity.

Although we are accustomed to the idea that several distinct notions of necessity and possibility coexist rather, well, harmoniously, the idea that harmony might pick out one of them as properly logical does not strike us as unreasonable. But the question is how, given the peculiar features of the inference rules of modal operators, can modal notions be incorporated in an account of harmony?

5.5 Read on modal operators

Our problem is that even if we can make sense of the idea that harmony should privilege particular modal operators in the way described in the previous section, it is still unclear what such an account would look like. Moreover, as we have noted, harmony, in the setting of the modal operators, comes at a price: we are forced to endorse the view that *both* introduction and elimination rules are at times individually determinative of meaning. This is a price many authors of both the introductions-first and eliminations-first bent would be unwilling to pay.

Prawitz has shown versions of the natural deduction systems for **S4** and **S5** to be normalizable, but this can be done only with considerable tweaking. Also, Prawitz's method does not provide us with local procedures that could be worked into a suitable principle of harmony. Another attempt at formulating a system of modal logic composed of harmonious rules of inference is again due to Read (2008). His proposal takes the form of a labelled deductive system for modal logic. Roughly, the idea is that every formula carries an index intuitively indicating the world at which it holds. Moreover, we introduce an explicit relation $<$ such that $i < j$ if world j is accessible from i. While the non-modal rules preserve labels in the obvious way, the 'labels are affected by, and determine the correctness of, the modal rules' (Read 2008, p. 15). Read gives the following rules for $\Box$, where $i \neq j$ and j does not occur in Γ.

$$\begin{array}{c}\Gamma, [i < j]^m, \\ \vdots \\ \Box\text{-I, } m \; \dfrac{A_j}{\Box A_i}\end{array}$$

$$\Box\text{-E} \; \dfrac{\begin{array}{c}\Gamma \\ \vdots \\ \Box A_i\end{array} \qquad \begin{array}{c}\Gamma', i < j \\ \vdots \\ A_j\end{array} \qquad \begin{array}{c}A_j \\ \vdots \\ C_k\end{array}}{C_k}$$

These rules are supplemented with 'equally straightforward and intuitive' (ibid., p. 17) rules for the possibility operator. Read claims that the rules for the modal operators so construed are indeed harmonious. He takes the existence of a levelling procedure for his rules to be indicative of this. Even leaving aside the details of Read's conception of harmony for now, this seems somewhat surprising.

It is true that the standard natural deduction systems for the modal operators do not admit of a levelling procedure. To see why consider the following proof of $\Box A \wedge \Box B$ from $\Box(A \wedge B)$:

$$\supset\text{-E} \; \dfrac{\supset\text{-E} \; \dfrac{\supset\text{-I, 2} \; \dfrac{\supset\text{-I, 1} \; \dfrac{\Box\text{-I} \; \dfrac{\wedge\text{-I} \; \dfrac{\Box\text{-E} \; \dfrac{[\Box A]^2}{A} \qquad \Box\text{-E} \; \dfrac{[\Box B]^1}{B}}{A \wedge B}}{\Box(A \wedge B)}}{\Box B \supset (\Box(A \wedge B))}}{\Box A \supset (\Box B \supset (\Box(A \wedge B)))} \qquad \wedge\text{-E} \; \dfrac{\Box A \wedge \Box B}{\Box A}}{\Box B \supset (\Box(A \wedge B))} \qquad \wedge\text{-E} \; \dfrac{\Box A \wedge \Box B}{\Box B}}{\Box(A \wedge B)}$$

The $\supset$-maximum $\Box A \supset (\Box B \supset (\Box(A \wedge B)))$ cannot be eliminated. A straightforward application of our levelling procedure does not result in a well-formed derivation.

$$\Box\text{-I}\,\cfrac{\wedge\text{-I}\,\cfrac{\Box\text{-E}\,\cfrac{\wedge\text{-E}\,\cfrac{\Box A \wedge \Box B}{\Box A}}{A} \qquad \Box\text{-E}\,\cfrac{\wedge\text{-E}\,\cfrac{\Box A \wedge \Box B}{\Box B}}{B}}{A \wedge B}}{\Box(A \wedge B)}$$

The last application of the $\Box$-introduction rule is not permissible, because the conclusion depends on $\Box A \wedge \Box B$ which is clearly not modal. That being said, the fact that Read's system admits of a levelling procedure is not of much help, at least not on its own. For we have seen that intrinsic harmony is not sufficient for harmony.

However, the real deep-seated flaw in Read's proposal is that the only way to make sense of the rules and the system in which they are couched is by appeal to possible worlds. Now, Read repeatedly stresses that the invocation of possible worlds is 'purely pedagogic' and so 'at best a useful metaphor' (ibid., p. 15). We incur no other-wordly ontological commitments, nor is the semantic detour essential to grasp the meanings of '$\Box$' and '$\Diamond$', according to Read. The account is proclaimed to be a purely inferentialist one: the meanings of the operators are given solely by the rules.

But how really does Read's system relate to inferential practice? Read offers no response to this question. He is surely right in saying that this 'is part of a wider question, how any part of the theoretical systematization of logic relates to practice' (ibid.). Nobody would deny that the question of what exactly makes a formal system an adequate representation of the informal practice it codifies is a tremendously difficult one. However, the absence of a satisfactory solution to this problem (indeed even the in principle insolubility of this problem) certainly does not license one to sever all ties between the formal system and the practice it represents. And this is where Read's system falters for there simply is no conceivable way of relating Read's system to our practice. This complete lack of consonance between formalism and practice is epitomized by the labels themselves and the auxiliary symbol '$<$', whose 'meaning, if any, is conferred by the rules' (ibid.). The question is not whether Read's system distorts our practice too much or not. This question does not even arise given that Read's system does not even purport to reflect our practice, or even any conceivable practice. Therefore the only way to make sense of Read's account is either via semantics (Kripke semantics to be exact), which would be tantamount to abandoning the inferentialist project altogether; or via a rather militant form of formalism, which would mean parting with the meaning-theoretically motivated inferentialist project we are engaged in (see section 2.2). Either way, the solution cannot lie here.

Of course it does not follow from the failure of Read's system that modalities have no place in an account of harmony. However, as our discussion of stability in the chapters to follow shows, the prospects for accommodating modal operators in such an account are bleak, to say the least. It would lead us too far astray to give an in-depth analysis of why modal operators cannot admit of harmonious inference rules. Nevertheless there is a clear intuitive sense why this is the case. Harmony requires there to be an equilibrium between V-principles and P-principles. But how can there be an equilibrium in the case of the modal operators? The case of the possibility operator illustrates the problem. For there to be equilibrium it would have to be the case that anything that follows from the grounds for asserting $\Diamond A$ must also follow from asserting $\Diamond A$ (on pain of P-weak disharmony). But $\Diamond A$ follows from A, but not everything that follows from A also follows from $\Diamond A$. In particular, of course, A itself follows from A but not from $\Diamond A$. Hence, if we really insist on maintaining equilibrium, the modal operators would collapse into the trivial operator t discussed in section 4.7. Consequently, on the inferentialist conception of logicality outlined in the beginning of our discussion of harmony (see section 3.3), neither the truth predicate nor the modal operators qualify as logical expressions.

Chapter 6

Stability: Tennant's principle of harmony

Two important conclusions were defended in the previous chapter. First, our notion of general harmony commits us to the principle of functionality and its cousin, the principle of injectivity. Given either the introduction rules or the elimination rules for a (putative) connective, the corresponding harmoniously matching set of rules will be determined uniquely. Second, we found that, while the principles of functionality and injectivity are not incompatible with modal operators, and while we can make sense of the notion of logical modalities picked out by the principle of harmony, there are good grounds for doubting that any account of harmony can be extended to encompass modal operators.

In the present chapter we return to the main aim of our inquiry, that of developing a workable principle of harmony. We shall begin by taking a closer look at existing accounts of stability. I begin by taking a brief look at Dummett's notion of stability, but then argue that it is Tennant who has provided us with the most promising account of our intuitive notion of harmony. Two putative arguments against Tennant are considered and rejected. We then present a counterexample of our own: prohibitively strong freak quantifiers are introduced and shown to be harmonious by Tennant's standards.

6.1 Dummett on stability

Let us begin with Dummett's own notion of *stability*. The intuitive idea is the following. Whereas intrinsic harmony constrains elimination rules in such a way that they do not outrun the introduction rules, stability also incorporates the inverse constraint: it requires also that the introduction rules do not overpower the elimination rules. Dummett puts the matter as follows:

> a little reflection shows that harmony is an excessively modest demand. If we adopt a verificationist view of meaning, the meaning of a statement is determined by what we acknowledge as grounds for asserting it. The fact that the consequences we conventionally draw from it are in harmony with these acknowledged grounds shows only that we draw no consequences its meaning does not entitle us to draw. It does not show that we fully exploit that meaning, that we are accustomed to draw all those consequences we should be entitled to draw (1991, p. 287).

The problem is that Dummett never follows through and formulates this more demanding constraint in a more rigorous manner. Moreover, what discussion he does provide of stability is obscured by the fact that it is closely tied up with his attempt at providing a justification procedure for the logical laws. The idea is that certain laws of logic are self-justifying and complete in the sense that all the remaining laws can be derived on their basis. Dummett considers first the verificationist approach according to which the class of self-justifying laws is just the set of introduction rules. Levelling procedures then provide immediate justification of the corresponding elimination rules. All other laws are then accounted for with the aid of what Dummett calls the *fundamental assumption*: the assumption, roughly, that any proof of A from hypotheses Γ can be converted into one which terminates with an application of one of the introduction rules corresponding to the main operator in A. As Dummett points out, one could, with the same intuitive plausibility, take a pragmatist approach and use the elimination rules as our point of departure, justifying appropriate introduction rules on *their* basis. The pragmatist justification procedure would be 'simply the mirror image' of the verificationist one. Stability in Dummett's sense obtains when both procedures come to the same result: if, in applying the verificationist procedure to a set $\mathcal{I}$ of introduction rules, we obtain a set $\mathcal{E}$ of elimination rules justified in terms of them; and if, conversely, starting with

$\mathcal{E}$, the pragmatist arrives at a justification of $\mathcal{I}$—in other words if the verificationist and the pragmatist agree on the meanings of the constants so determined—then $\mathcal{I}$ and $\mathcal{E}$ are stable.

The problem with this is twofold. First, it is far from clear how Dummett's rather sketchy account ought to be filled in. Starting with $\mathcal{I}$ we can rely on levelling procedures to guide us in choosing a suitable corresponding $\mathcal{E}$. It is hinted that there should be an analogous procedure available to the pragmatist. But Dummett leaves us in the dark, as to what this procedure should amount to. Second, Dummett's fundamental assumption is also beset with difficulties. And not even Dummett himself appears to be overly optimistic about the prospects of salvaging the fundamental assumption. For these reasons, I propose to turn now to Tennant's proposal.

6.2 Tennant's notion of harmony

Tennant has developed a way of dealing with the fundamental asymmetry that bedevils Dummett's notion of intrinsic harmony. The problem, let us recall, is that the existence of a procedure for levelling local peaks for pairs of sets of introduction and elimination rules fails to guard against P-weak disharmony, and in consequence can also lead to P-strong disharmony in certain contexts (as it did in the Q-example). Contrary to Dummett, Tennant offers us a fully worked out and rigorous solution to this problem, one that also captures our intuitive notion of harmony as a balance or equilibrium.[1]

We had said that for such a balance to obtain an elimination rule must exploit, by way of the inferential consequences it licenses, all and only the content conferred on its major premise by the corresponding introduction rule. In other words, we should be able to derive from a statement with the constant in question in a dominant position no more than what we are entitled to in virtue of being in a position correctly to assert it; nor should we be able to derive from it any less than the introduction rule allows. To the end of giving this intuitive idea a precise form, Tennant introduces the notions of the logically *strongest* and *weakest* propositions

[1] Tennant's 'harmony principle' has undergone a number of modifications over the years. To my knowledge it was first formulated in Tennant (1978). It was revised in Tennant (1987) and subsequently subject to minor tweaking. I am relying here on what I assume is the latest formulation in Tennant (forthcoming).

with a certain property.

- A is the *strongest* proposition with property P if, for any proposition B with the same property, A entails B.
- A is the *weakest* proposition with property P if A is entailed by any proposition B with the same property.

It should be noted that I am here following Tennant's non-standard use of 'proposition': a proposition A is the logical equivalence class of which A is a member, i.e. the class of statements logically equivalent to A.

With the aid of these tools, Tennant then proceeds to define the principle of harmony (the lower case 'h' is crucial here). An (arbitrary binary logical) operator $ is harmonious if the following two conditions are met:

> (S) $(A, B) is the strongest conclusion possible under the conditions described by $-I. Moreover, in order to show this,
>
> 1. one needs to exploit all the conditions described by $-I;
> 2. one needs to make full use [of] $-E;
> 3. but one may not make any use of $-I.
>
> (W) $(A, B) is the weakest major premise possible under the conditions described by $-E. Moreover, in order to show this,
>
> 1. one needs to exploit all the conditions described by $-E;
> 2. one needs to make full use [of] $-I;
> 3. but one may not make any use of $-E.[2]

When harmony obtains between $-I and $-E we may write $h(\$\text{-I}, \$\text{-E})$. Tennant states the motivating idea behind his formulation as follows.

> Introduction and elimination rules for a logical operator $ must be formulated so that a sentence with $ dominant expresses the strongest proposition which can be inferred from the stated premises when the conditions for $-introduction are satisfied; while it expresses the weakest proposition possible under the conditions described by $-elimination (ibid., p. 19, again with notational adjustments).

[2]See Tennant (forthcoming, p. 25, with some minor notational adjustments).

Let us briefly illustrate the way that Tennant's proposal works. Consider the rules for $\supset$. To demonstrate that they satisfy Tennant's principle of harmony, we need to show that $A \supset B$ is the strongest proposition that can be inferred by means of $\supset$-I and the weakest from which we can infer the conclusion of $\supset$-E. Moreover, our method for showing this has to conform to the constraints (S) 1–3 and (W) 1–3. Assume first that X is a conclusion obtained via $\supset$-I making 'full use' of it. By this assumption the following rule of inference ($\star$) may be used.

$$ {\scriptstyle \star,\, i}\ \dfrac{\begin{array}{c}\Gamma, [A]^i \\ \vdots \\ B\end{array}}{X} $$

Our task now is to show that in this case (S) obtains, i.e. that $A \supset B \vdash X$, making full use of the elimination rules for $\supset$, but not making any use of the introduction rules for $\supset$. The following derivation does the job

$$ {\scriptstyle \supset\text{-E}}\ \dfrac{[A]^i \qquad A \supset B}{{\scriptstyle \star,\, i}\ \dfrac{B}{X}} $$

We then have to show that if X satisfies the aforementioned requirements and features as major premise of an elimination as follows (call this rule of inference '$\bullet$')

$$ \bullet\ \dfrac{\begin{array}{c}\Gamma \\ \vdots \\ A\end{array} \qquad X}{B} $$

then it is the case that $X \vdash A \supset B$; i.e. we are required to show, in conformity with (W), that $A \supset B$ is the weakest proposition that can be inferred by $\supset$-E under the conditions specified. This can be done as follows

$$ {\scriptstyle \supset\text{-I},\, i}\ \dfrac{\bullet\ \dfrac{[A]^i \qquad X}{B}}{A \supset B} $$

6.3 The need for <u>H</u>armony

So much for exposition. We now turn to the critical evaluation of Tennant's proposal. One reason it might be faulted is because it too fails to accommodate modal operators. Such an argument would be alike in essentials to the one given in section 5.5. Consider, for instance, the elimination rule for $\Box$ (which, as we have seen, is shared by a number of systems, including **T**, **S4** and **S5**):

$$\Box\text{-E}\ \frac{\Box A}{A}$$

According to Tennant's account of <u>h</u>armony the proposition expressed by $\Box A$ ought to be the weakest proposition that entails A. But clearly the weakest proposition that entails A is A itself. Therefore, <u>h</u>armony requires that A and $\Box A$ be interderivable (and thus the same proposition in Tennant's sense). $\Box A$ would have to be derivable from A with no restrictions at all, in effect making it synonymous with the operator $t(A)$. Consequently we cannot, while retaining the restrictions on $\Box$-I, prove that $X \vdash \Box A$ and thus show that $\Box A$ is indeed the weakest such proposition.[3] However, as we have seen, incorporating modalities is a problem for proof-theoretic accounts in general. The failure of Tennant's theory to account adequately for such operators is therefore no reason to dismiss it.

However, Tennant's account, as we have presented it so far, encounters more serious problems. We now consider two objections that are decisive against the principle of <u>h</u>armony. This motivates a move from <u>h</u>armony to a stronger and more sophisticated notion which we denote <u>H</u>armony (note the upper case 'H') and which remains unscathed.

The first problem stems from the quantum-disjunction operator: on Tennant's account of <u>h</u>armony the rules for $\sqcup$ turn out to be <u>h</u>armonious. For consider a formula X that obeys the conditions described by $\sqcup$-I, i.e. the following inference rules (†) are legitimate

$$\begin{array}{c}\Gamma\\ \vdots\\ \dagger\ \dfrac{A}{X}\end{array}$$

[3]It has been brought to my attention that this point has also been made by Nick Denyer (Denyer 1989).

$$\dagger\ \dfrac{\begin{matrix}\Gamma\\ \vdots\\ B\end{matrix}}{X}$$

We now need to show that (S) obtains, i.e. that $A \sqcup B \vdash X$, using all the resources made available by ⊔-E but not using ⊔-I. This is achieved as follows:

$$\sqcup\text{-E},\ i\ \dfrac{A \sqcup B \qquad \dagger\,\dfrac{[A]^i}{X} \qquad \dagger\,\dfrac{[B]^i}{X}}{X}$$

Turning now to Tennant's condition (W), suppose this time that X satisfies the conditions stated by ⊔-elimination so that the following rule holds (‡)

$$\ddagger,\ i\ \dfrac{X \qquad \begin{matrix}[A]\\ \vdots\\ C\end{matrix} \qquad \begin{matrix}[B]\\ \vdots\\ C\end{matrix}}{C}$$

The following demonstrates that $X \vdash A \sqcup B$, in the requisite way.

$$\ddagger,\ i\ \dfrac{X \qquad \sqcup\text{-I}\,\dfrac{[A]^i}{A \sqcup B} \qquad \sqcup\text{-I}\,\dfrac{[B]^i}{A \sqcup B}}{A \sqcup B}$$

Thus one can demonstrate h(⊔-i, ⊔-e), i.e. the present account of harmony validates ⊔ in just the same way that it validates our ordinary disjunction operator. This is a damaging result for the notion of <u>h</u>armony, since there can be systems, which, though composed entirely of <u>h</u>armonious rules, are non-conservative and non-normalizable. (See section 4.4 and our discussion of the Q-example for details of this argument.) However, to ward off cases like these, Tennant imposes the further requirement of <u>H</u>armony:

> Given \$-E we determine \$-I as the *strongest* introduction rule \$-i such that h(\$-i, \$-E). Given \$-I we determine \$-E as the *strongest* elimination rule \$-e such that h(\$-I, \$-e) (Tennant forthcoming, p. 22, with slight notational adjustments).

By this method we arrive at H($-I, $-E). Applied to the problem of determining the correct elimination rule for disjunctions, we find that the second clause of Tennant's additional constraint forces upon us the choice of $\vee$ over $\sqcup$: we have $H(\vee\text{-I}, \vee\text{-E})$, but not $H(\sqcup\text{-I}, \sqcup\text{-E})$.

Another putative objection along similar lines advanced by Kurbis (unpublished) can also be neutralized by the additional constraint. Kurbis argues that Tennant's notion of harmony is unduly sensitive to the presence or absence of other connectives in the system. The example he proposes is Anderson and Belnap's system **E** of 'strict relevant implication'. Let us again follow the conventional practice and denote the conditional in relevant logics by '$\rightarrow$' to distinguish it from the horseshoe we have been using for the classical and intuitionistic conditionals we have been considering until now.[4] The introduction rule for $\rightarrow$ in **E** is

$$\dfrac{\begin{matrix}\Gamma, [A]^i \\ \vdots \\ B\end{matrix}}{A \rightarrow B}\ \mathbf{E}{\rightarrow}\text{-I},\ i$$

where all the formulas in Γ must be of the form $C \rightarrow D$. The elimination rule is the familiar *modus ponens*. Are the rules for $\rightarrow$ in harmony, i.e. do we have $h(\rightarrow\text{-i}, \rightarrow\text{-E})$? Well, certainly not in general, for when we try to show that $A \rightarrow B$ is the weakest formula that can be inferred in accordance with $\mathbf{E}{\rightarrow}$-E, we run into obvious difficulties. For suppose that we could make use of the rule $(\star')$

$$\dfrac{\begin{matrix}\Gamma \\ \vdots \\ A\end{matrix} \qquad X}{B}\ \star'$$

The usual procedure for showing that $X \vdash A \rightarrow B$ fails:

$$\dfrac{\dfrac{[A]^1 \qquad X}{B}\ \star'}{A \rightarrow B}\ \mathbf{E}{\rightarrow}\text{-I},\ 1$$

[4]We follow the established practice even though, as we have argued in section 2.6, this practice is misguided. There is no difference in meaning between the relevantist's $\rightarrow$ and the intuitionist's $\supset$. What separates the two systems are their respective structural assumptions, not the meanings of the logical operators involved in them.

It fails because we have no right to assume that X is of the required form. Therefore, in general, the rules for **E**→ fail to harmonize. However, Kurbis goes on to point out that they do harmonize in the context of the purely implicational fragment. Indeed it is not hard to see that in such a restricted context X can be neither atomic nor identical to B (because $B, A \vdash B$ violates relevance requirements) and so must be of the required form. The moral of the story according to Kurbis is that whether or not a constant is governed by harmonious inference rules depends on what other connectives are in the system. If this were so, Tennant's principle of harmony would indeed lack the desired property of locality. However, what Kurbis fails to realize is that the rules for → do not Harmonize, not in the full system **E**, nor in any of its subsystems. Once again, this is because the Harmony constraint forces us to select the strongest among harmonious introduction rules for → given the elimination rule. Our choice must therefore fall on the unrestricted introduction rule, thereby collapsing **E**→ into **R**→ (supposing that relevance is guaranteed by the appropriate structural restrictions against vacuous discharge and that *ex falso* is omitted). With the crucial additional requirement of Harmony in place, Kurbis's objection need not faze Tennant.

6.4 A counterexample to Tennant's account

So far, so good. But how exactly are we to understand the all important requirement of Harmony (as opposed to harmony)? *Prima facie* it seems tolerably clear what is meant when we are instructed to choose the *strongest* introduction/elimination rule harmonious with a given elimination/introduction rule. To establish the Harmony of a pair (possibly of sets) of rules \$-I and \$-E we hold a particular rule (or set thereof) \$-I fixed and then run through all of the \$-e rules such that $h(\$\text{-I}, \$\text{-e})$. Among these we choose the strongest, \$-E. We now repeat the same procedure, this time taking \$-E as our point of departure, and arrive at \$-I′. If \$-I′ and the original \$-I are identical we have established that $H(\$\text{-I}, \$\text{-E})$. If \$-I′ is stronger than \$-I we repeat the process until we do reach an equilibrium.

This presupposes that our grasp of what is to count as a legitimate candidate to be an inference rule (or set thereof) of the type required (introduction or elimination) is sufficiently clear for the talk of 'running through' a set of such rules to make sense. But even if we suppose this collection to be sufficiently determinate (or at least sufficiently determinable), there remains the problem of what exactly is meant by the 'strength' of a rule. The only hint that Tennant gives us is this.

> This captures the following thought concerning the roles of introduction and elimination. \$-E determines my commitments in asserting \$(A, B). To determine the corresponding \$-I I have to ask: what is the *least* I have to be able to do via \$-I in order to be able to meet my commitments? Likewise \$-I determines what I have to know in order to assert \$(A, B). To determine the corresponding \$-E I have to ask: how can I extract the *most* out of that via \$-E (Tennant 1987, p. 97, with notational adjustments).

The problem is that if we take this at face value it is unclear how we are to justify *any* kind of restriction at all. While the Harmony constraint was serviceable when it came to ruling out the unnecessarily restrictive quantum-logical elimination rule, are we not running the risk of overshooting and thereby obliterating the kind of restrictions on rules we *do* want? Consider the introduction rule for $\exists$:

$$\exists\text{-I}\ \frac{\begin{array}{c}\Gamma\\ \vdots\\ A[t/x]\end{array}}{\exists x A(x)}$$

As is standard, $A[t/x]$ is the result of uniformly substituting t for all free occurrences of x in A, and we require x to be freely substitutable for t in $A(t)$.[5] The corresponding elimination rule is this:

$$\exists\text{-E},\ i\ \frac{\begin{array}{c}\Gamma\\ \vdots\\ \exists x A(x)\end{array} \qquad \begin{array}{c}\Gamma', [A[a/x]]^i\\ \vdots\\ C\end{array}}{C}$$

where the parametric a may not occur in Γ', $\exists x A(x)$ or C.

First let us show that these rules are harmonious in Tennant's sense, i.e. h($\exists$-I, $\exists$-E). First we show that (S) is satisfied. To this end, suppose there is a proposition X which is, for every term t, entailed by $A[t/x]$. In other words, suppose the following rule ($\circ$) holds:

[5] Note that $A(x)$ need not have been obtained from $A[t/x]$ necessarily by substituting x for all occurrences of t. For example, it is the case that $R(t,t) \vdash \exists x R(x,t)$, but also $R(t,t) \vdash \exists x R(x,x)$, etc.

$$\begin{array}{c} \Gamma \\ \vdots \\ \circ\ \dfrac{A[t/x]}{X} \end{array}$$

We now want to show that $\exists x A(x) \vdash X$, making full use of the elimination rule and none of $\exists$-I. *Ex hypothesi*, X is implied by $A[t/x]$ *for every* t via $\circ$, therefore in particular for any t satisfying the constraints imposed by $\exists$-E. It follows that X follows from any arbitrary parametric a. We thus obtain the desired:

$$\exists\text{-E, 1}\ \dfrac{\exists x A(x) \qquad \circ\ \dfrac{[A[a/x]]^1}{X}}{X}$$

Turning now to (W), assume that X features as the major premise of $\exists$-E and thus entails every proposition that follows from $A[a/x]$, where a is again parametric and the usual conditions apply. Our assumption amounts to the following rule ($\circledcirc$):

$$\circledcirc, i\ \dfrac{X \qquad \begin{array}{c} \Gamma, [A[a/x]]^i \\ \vdots \\ C \end{array}}{C}$$

In particular, then, we obtain $X \vdash \exists x A(x)$:

$$\circledcirc, 1\ \dfrac{X \qquad \exists\text{-I}\ \dfrac{[A[a/x]]^1}{\exists x A(x)}}{\exists x A(x)}$$

This establishes that our standard rules for the existential quantifier are indeed $\underline{\text{h}}$armonious. But is $\exists$-E also the strongest elimination rule $\underline{\text{h}}$armonious with $\exists$-I? Herein lies the problem. For consider the elimination rule $\exists$-E′, identical to the regular elimination rule for the existential quantifier with the exception that it lacks some or all of the restrictions on the parameter a. A simple inspection of the above demonstration of $\underline{\text{h}}$armony reveals that the same demonstration goes through for the excessively permissive $\exists$-E′. Hence $h(\exists\text{-I}, \exists\text{-E}')$ holds also. But now let $\exists$-E′ be the elimination rule devoid of any restrictions. Surely *that* rule is the strongest such elimination rule as it enables us to 'extract' the most from sentences containing $\exists$ in a dominant position (in the sense discussed in section 5.3 where we defended the principle of functionality). In particular, of course, it is 'stronger' than the ordinary

∃-elimination rule. It follows that it is the wholly unrestricted elimination rule that is in Harmony with the introduction rule given above, i.e. we have $H(\exists\text{-I}, \exists\text{-E}')$.[6]

But do we really get $H(\exists\text{-I}, \exists\text{-E}')$? Or must we in turn revise our introduction rule? Only if we can also show that ∃-I is the strongest introduction rule harmonious with ∃-E′ will our conclusion pass muster. Again we might be tempted to try out a modified introduction rule devoid of any constraints. However, the restrictions in the standard elimination rule play a rather different role than the restrictions we impose on introductions of the existential quantifier. In the former case, the restrictions—the term a may not occur in any of the undischarged hypotheses, nor in the major premise, nor in the conclusion—are specifically tailored to capture the idea that a is arbitrary. This constraint is beholden ultimately only to the intended meaning of ∃. Or, to put it in a way more germane to the spirit of the inferentialist project, the restrictions are partially determinative of the meaning of the existential quantifier. By contrast, the restriction we impose on the standard ∃-introduction rule serves solely the purpose of ensuring that every application of the rule will result in a well-formed formula, i.e. that non-well-formed formulas like the following are prevented.

$$\exists\text{-I}\ \frac{\forall x R(x,a)}{\exists x \forall x R(x,x)}$$

Therefore, it seems that ∃-I really *is* the strongest introduction rule that matches ∃-E′. And we get $H(\exists\text{-I}, \exists\text{-E}')$ after all.

Something has gone very wrong here. Clearly the quantifier whose meaning we have so fixed—call it 'Ⅎ'—bears little resemblance to the existential quantifier that we have come to know and love. Tennant's principle of harmony failed to pick out the intended operator. But more importantly, the constant we end up with is incoherent. For it enables us to derive $F(b)$ for any b we like, provided only that $F(a)$ holds for some a.[7]

$$\text{Ⅎ-E, 1}\ \frac{\text{Ⅎ-I}\ \dfrac{F(a)}{\text{Ⅎ}xF(x)} \qquad [F(b)]^1}{F(b)}$$

[6] In correspondence Tennant conceded this point. In response he proposed a solution very similar to the one I (independently) consider and reject in section 8.1. I have not yet had the opportunity for an exchange with Tennant concerning the solution I propose in sections 8.2–8.7.

[7] The following proof violates the standard restriction imposed upon the ∃-E rule that the parameter a ought not to occur in the conclusion of the minor premise.

The introduction of Ⅎ not only violates total harmony, it is non-conservative in the even stronger sense that it violates the principle of innocence: it makes available new information about the world—indeed a lot of new information! By appeal to Ⅎ, we can justify the assertion of an atomic sentence $F(b)$ for any individual named by 'b' so long as the predicate F is correctly assertible of at least one individual. Clearly, therefore, the resulting system is not normalizable either: there cannot be any direct deductive path from $F(a)$ to $F(b)$ without introducing and then eliminating more complex formulas.

6.5 Expanding the counterexample

A similar story can be told with respect to the deviant quantifier ⅁. ⅁ employs the standard elimination rule for the universal quantifier:

$$\begin{array}{c}\Gamma\\ \vdots\\ \forall\text{-E}\ \dfrac{\forall x A(x)}{A[t/x]}\end{array}$$

However, instead of the customary introduction rule ∀-I

$$\begin{array}{c}\Gamma\\ \vdots\\ \forall\text{-I}\ \dfrac{A[a/x]}{\forall x A(x)}\end{array}$$

where a does not occur in Γ and the substitution of x for a must be uniform, we adopt an unrestricted rule ∀-I′ of the same form but devoid of the restrictions on a.

We again begin by showing that $h(\forall\text{-I}, \forall\text{-E})$. First we check that (S) is satisfied. Suppose there is a proposition X that is implied by Γ when Γ implies $A[a/x]$ parametrically so that the following rule of inference (⊛) holds:

$$\begin{array}{c}\Gamma\\ \vdots\\ \circledast\ \dfrac{A[a/x]}{X}\end{array}$$

Applying ∀-E we can show, as required, that $\forall x A(x) \vdash X$

$$\begin{array}{c} \Gamma \\ \vdots \end{array}$$
$$\circledast\,\dfrac{{\scriptstyle \forall\text{-E}}\,\dfrac{\forall x A(x)}{A[t/x]}}{X}$$

a's being parametric guarantees that X follows from $A[t/x]$ for any closed term t that we care to instantiate the universal quantifier with. As for (W), take X to be a proposition which, for any t, entails $A[t/x]$ in accordance with the following rule ($\ominus$)

$$\ominus\,\dfrac{X}{A[t/x]}$$

Given that this holds for any t whatsoever, we can pick a such that it does not occur in X or in any of the premises on which X depends. This licenses the inference to $\forall x A(x)$ as desired

$$\forall\text{-I}\,\dfrac{\ominus\,\dfrac{X}{A[a/x]}}{\forall x A(x)}$$

A moment's reflection shows that the same reasoning establishes $h(\forall\text{-I}', \forall\text{-E})$. Since $\forall$-I′ is stronger than $\forall$-I (though of course in a number of undesirable ways), we again come to the unfortunate conclusion that $\forall$-I′—the deviant rule—is H̲armonious with the standard rule $\forall$-E.

Again it can be shown that the freak quantifier $\backepsilon$, leads to unacceptable consequences:

$$\supset\text{-E}\,\dfrac{\supset\text{-I},1\,\dfrac{\backepsilon\text{-E}\,\dfrac{\backepsilon\text{-I}\,\dfrac{\exists x F(x) \qquad \backepsilon\text{-I}\,\dfrac{[F(a)]^1}{\backepsilon x F(x)}}{\backepsilon x F(x)}}{F(b)}}{F(a) \supset F(b)} \qquad F(a)}{F(b)}$$

As in the case of $\Finv$, we can prove any b whatsoever to be F provided that there is at least one F. The violation of the standard restrictions occurs in the $\backepsilon$-introduction rule on the right-hand side.

What can we say in defence of Tennant's principle? It may be said in response that such unrestricted elimination rules as $\exists$-E$'$, say, were never permissible in the first place; $\exists$-E$'$ was never a possible value for $\exists$-e. But what grounds do we have for ruling such rules out? Obviously, appealing to the meaning of the existential quantifier and claiming that only such and such restrictions do it justice is not an option: for why then could the quantum logician not with the same right lay claim to the primacy of *his* favoured meaning of $\exists$? The whole point is that harmony is to give us a well-motivated principle by which to arbitrate disputes between such rival interpretations. As Tennant himself puts it elsewhere,

> there is no independent access to the meaning of the word with respect to which one could then raise the question whether the rules governing it 'respect' that sense (Tennant 2005b, p. 629).

But how else can we justify the restrictions on our quantifier rules? Before we go on to propose our own solution to this puzzle, we shall consider another recent account, this one advanced by Read.

Chapter 7

Read's 'general elimination harmony'

7.1 General elimination harmony

Read's point of departure is once again Gentzen's remarks concerning the functionality of the relationship between introduction and elimination rules. According to Gentzen it ought to be possible, given a set of introduction rules, to 'read off' the harmoniously matched elimination rules. It is this intuitive idea that drives Read's characterization of harmony. Provided that Read's account manages to pin down our intuitive notion of harmony adequately, it will have one clear advantage over Tennant's account, for Read's account promises to furnish a method of determining the matching rule (or set of rules) for a given rule of inference—while Tennant's account had to rely on an assumed ability to grasp the totality of all matching rules for a given operator and to choose the strongest among them. Once we know under what circumstances the assertion of a statement containing the constant in question in a dominant position is licensed, we should be able to figure out the corresponding elimination rule because of the functionality inherent in our conception of harmony: for any introduction rule there should only be one matching elimination rule. Hence, an adequate account of harmony should allow us to identify a harmonious counterpart for any given rule.[1]

[1] Read believes that introduction rules have meaning-theoretic primacy. He is therefore only concerned with the problem of 'deriving' elimination rules on the basis of a given set of introduction rules. In conformity with the ecumenical line we have been taking so far, we would ideally want

Read presents matters as follows.

> In general, if $\{\Pi_i\}$ denotes the grounds for introducing some formula A (introducing an occurrence of a connective δ in A), then the elimination rule for δ should permit inference to an arbitrary formula C only if $\{\Pi_i\}$ themselves entail C (Read 2000, p. 130).

He thus proposes that, given the introduction rules (of which there may be several, one for each Π_i) for a particular connective, say the binary \$:

$$\text{\$-I}\ \frac{\Pi_i}{\$(A,B)}$$

harmony obtains if the corresponding elimination rule can be 'read off'. But how do we go about doing this?

Read suggests that readability is obscured by the standard format of elimination rules for some of the connectives. Instead of using the familiar natural deduction elimination rules, we would do better to adopt elimination rules modelled on those for $\vee$ and $\exists$. It turns out that such so-called *generalized elimination rules* can be formulated for all our usual constants, and that they are provably equivalent to our familiar elimination rules.[2] Read seems to go further, claiming that *any* permissible elimination rule can be formulated in this way. Rather than attempting to give a precise definition of generalized elimination rules at this point, it will be more illuminating to illustrate the notion by way of examples. We will consider the generalized elimination rules for $\wedge$ and $\supset$:

$$\wedge\text{-GE}, i\ \frac{\begin{matrix}\Gamma \\ \vdots \\ A \wedge B\end{matrix} \qquad \begin{matrix}\Gamma', [A,B]^i \\ \vdots \\ C\end{matrix}}{C}$$

$$\supset\text{-GE}, i\ \frac{\begin{matrix}\Gamma_0 \\ \vdots \\ A \supset B\end{matrix} \qquad \begin{matrix}\Gamma_1 \\ \vdots \\ A\end{matrix} \qquad \begin{matrix}\Gamma_2, [B]^i \\ \vdots \\ C\end{matrix}}{C}$$

to extend such an account into one that is functional in both directions. On such an extended account, any introduction/elimination rule would carry with it necessary information required to determine the corresponding elimination/introduction rule.

[2]See Schroeder-Heister (2004) for the history of generalized elimination rules.

A characteristic feature of generalized elimination rules is that the major premise containing the formula to be 'eliminated' is to be found in the left-most column. We can then distinguish two types of generalized elimination rules: those that match introduction rules involving subproofs and the possible discharge of assumptions (e.g. $\supset$-GE), and those that match introduction rules without subproofs (e.g. $\wedge$-GE).[3]

In the former case the rule will contain one or more columns representing subderivations of the assumptions. The far right-hand side will contain a column representing a subderivation of some formula C from the formula (in the case of $\supset$-GE, B) that constitutes the conclusion of the subproof of the introduction rule. The shift from the standard elimination rule, *modus ponens*, to the generalized elimination variant consists in the fact that we are simply building the subsequent derivation from B into the very elimination rule in the form of a subproof. What would usually be downstream from the elimination rule, i.e. the derivation of C from B,

$$\dfrac{\begin{matrix}\vdots \\ A \supset B\end{matrix} \qquad \begin{matrix}\vdots \\ A\end{matrix}}{\begin{matrix}B \\ \vdots \\ C\end{matrix}}\ \supset\text{-E}$$

is integrated into the generalized elimination rule as its right-most column.

The same general idea applies to the latter case of elimination rules corresponding to simple introduction rules not containing subproofs. In such rules, the grounds for asserting the major premise appear in the form of dischargeable assumptions in columns on the right (one or several). The underlying idea is, roughly, that whatever can be deduced from formulas that constitute sufficient grounds to introduce the major premise of the elimination rule can also be deduced from the major premise itself. Therefore, whenever it can be shown in a subdeduction nested within an application of the rule that some formula C can be derived from the grounds for introducing the major premise, then these grounds can be discharged, thereby showing that C follows from the major premise alone (or in the case of the former type of elimination rule, on the basis of the assumptions and the major premise alone).[4]

[3] We will elaborate on this distinction in section 8.4.

[4] Note that $\wedge$-GE involves assumption classes with more than one formula. Despite the unfamiliarity of such assumption classes there is nothing spooky about them. The tacit assumption underlying them is that ⌜supposing A and supposing B⌝ can be understood independently of understanding ⌜supposing $A \wedge B$⌝ (cf. section 12.4). This assumption is crucial if we want to avoid

The advantage of generalized elimination rules is that they offer a uniform format in which all elimination rules can be expressed. This uniformity promises to facilitate the task of characterizing the functional relationship between introductions and eliminations. If our hope is to uncover a general method for determining the matching elimination rules for any suitable set of introduction rules, this will be an important preliminary. We know the general shape of the rule. All that remains to be done is to establish a set of guidelines that tell us how to fine-tune the elimination rule in accordance with the demands of harmony and the distinctive features of the corresponding introduction rules.

According to Read, generalized elimination rules thus, as it were, carry the harmony constraint on their faces. They neatly encapsulate Read's informal characterization of harmony: whatever can be inferred via an elimination rule can be arrived at more directly on the basis solely of the premises required for establishing the corresponding major premise. To put it bluntly: we cannot get more out than we put in. On the other hand, the rules show that *all* the consequences of the grounds for introducing the major premise can be obtained from the elimination rules. So, it looks as if we should also be able to get everything we put in out again. On the face of it it seems that we have improved on the Dummettian notion of intrinsic harmony as the existence of a levelling procedure and so moved a good step closer to a rigorous formulation of the notion of general harmony. Moreover, such an account would be preferable to Tennant's account. Rather than invoking the notion of the strength of propositions, we are building the notion of deductive equilibrium into the very form of the elimination rule.

---

circularity. Yet it also seems reasonable. In fact, we already implicitly endorse it when we countenance instances of $\wedge$-I where both of the immediate premises are assumptions. It also underpins familiar sequent formulations of natural deduction. In any event, we could equally well reformulate $\wedge$-GE as two separate elimination rules, one for each conjunct:

$$\wedge\text{-GE}, i\ \frac{\begin{array}{c}\Gamma\\ \vdots\\ A \wedge B\end{array} \qquad \begin{array}{c}\Gamma', [A]^i\\ \vdots\\ C\end{array}}{C}$$

And likewise for B.

7.2 Criticism of Read's account

One problem with Read's account is that it is inaccurate to say that *everything* that follows from the grounds for asserting the major premise also follows from the major premise itself via the elimination rule. The case of disjunction makes the need for further refinement plain. Obviously it is not the case that whatever follows from *either* of the dischargeable premises of the disjunction elimination rule (along with collateral hypotheses) must also be inferable from the corresponding elimination rule, but rather only what follows from *both* premises. A simple illustration of this point is the proof

$$\vee\text{-E}, 1\ \dfrac{\begin{array}{c}\vdots\\ A \vee (A \wedge B)\end{array} \qquad [A]^1 \qquad \wedge\text{-E}\ \dfrac{[A \wedge B]^1}{A}}{A}$$

Clearly, B follows from $A \wedge B$ but not from $A \vee (A \wedge B)$. Only those formulas that are consequences of *both* of the canonical grounds for asserting the major premise need to be deducible from $A \vee B$. By contrast, the $\wedge$-elimination rule hinges on the fact that there is just one subproof the conclusion of which may be a consequence of *either* (or both) of the dischargeable formulas (jointly with the premises). But *these* are precisely the characteristic features of the rules that account for the individual senses of the constants. Read's slogan that everything that follows from the grounds for introducing the constant in question ought also to be derivable via the elimination rules seems to be too crude. It ignores the fact that what we need to know in order to know the meaning of the constants governed by the rules in question is exactly *which* consequences to take into account and *how* they follow from their premises. Short of providing such information, the rules of inference in question cannot be said to do inferentialist duty. Nor will we be able to determine matching elimination rules on their basis.

More importantly, there appears to be nothing in Read's formulation that prevents P-weak disharmony. Even if the *form* of an elimination rule is determined by an introduction rule, we do not yet have directives to hand that would guide us in a choice, say, between the standard and the quantum-logical or-elimination rule. After all, both elimination rules enable us to infer anything that follows from both of the possible grounds (i.e. the dischargeable premises) for inferring the major premise. The question as to whether it is permissible to wheel in collateral hypotheses is not resolved.

One might of course try a trick analogous to Tennant's move to his principle of Harmony to prevent such (from a logical perspective) unnecessary weakening of elimination rules. That is, one can add the further requirement that an elimination rule should enable us to extract the maximum of deductive consequences commensurate with the grounds under which the introduction of its major premise is licensed. Allowing additional hypotheses, for example, certainly maximizes the consequences we can infer from our elimination rule. Must we not therefore conclude that we ought to adopt the maxim that in selecting among several structurally adequate elimination rules, we ought always to choose the strongest one (i.e. the one without any restrictions on collateral hypotheses)? But now we of course run into the same difficulties concerning quantifier rules that we encountered on Tennant's account. How can we calibrate our choice of elimination rules in such a way as to prevent P-strong disharmony (as in the case of the unrestricted quantifier rules) *and* P-weak disharmony (as in the case of the quantum-logical elimination rule)?

In summary, then, Read's account requires shoring up on two fronts. First, it does not supply us with a general method for determining elimination rules from introduction rules. Second, it does not succeed in barring P-weak disharmony. In what follows I propose an account that overcomes these difficulties.

Chapter 8

Harmony: A new proposal

8.1 Reincorporating intrinsic harmony

Let us begin with the problem of P-strong disharmony. We have seen that Tennant's account ultimately fails to account for the necessary restrictions on the rules for quantifiers (and, perhaps inevitably, modal operators) and hence falls victim to the counterexample we presented in section 6.4. What can be done? A promising point to begin with is that although the rogue quantifier ⅃ is validated by both accounts, it is not validated by the principle of intrinsic harmony (i.e. the requirement that all peaks must be susceptible of being levelled).

Consider an ordinary ∃-peak

$$
\dfrac{\dfrac{\begin{matrix}\Gamma_0\\ \Pi_0\\ A[t/x]\end{matrix}}{\exists x A(x)}\;\scriptstyle{\exists\text{-I}} \qquad \begin{matrix}\Gamma_1, [A[a/x]]^i\\ \Pi_1\\ C\end{matrix}}{C}\;\scriptstyle{\exists\text{-E},\, i}
$$

Such peaks can be levelled as follows

$$
\begin{matrix}\Gamma_0\\ \Pi_0\\ A[t/x]\\ \Gamma_1, A[t/x]\\ \Pi_1[t/a]\\ C\end{matrix}
$$

A closer look reveals that the restrictions on parameter a are indispensable here. Only when the restrictions ensuring a's arbitrariness are in place is it in general possible to concatenate the proof Π_0 introducing the major premise with the proof Π_1 of the minor premise on the right. The restrictions ensure that Π_1 is a schematic proof that functions as a blueprint for other proofs with the same inferential structure. That is, Π_1 goes through in just the same way if we relabel it uniformly by replacing all occurrences of a by occurrences of any closed term t. That this is so is guaranteed by the

> **Substitution Lemma**: if Σ is a proof of B from hypotheses Δ and u is a closed term and b a parameter, then $\Sigma[u/b]$ is a proof of $B[u/b]$ from hypotheses $\Delta[u/b]$.[1]

Here we assume that no name u occurs in Σ. Then $\Sigma[u/b]$ is the result of substituting all occurrences of b in every formula in Σ by u. Similarly for Δ and the conclusion B. The levelling procedure for the existential quantifier is a special case of the lemma. What is of central importance for us here is that the lemma and hence the levelling procedure essentially depend on the restrictions imposed by $\exists$-E, and hence are inapplicable to Ⅎ-E. For only if a occurs in neither Γ_1 nor A nor C does the proof remain unperturbed; only then can we be sure that the same conclusion is deducible from the same premises via the same proof *modulo* relabelling.

The same is true for the deviant 'universal' quantifier ⅁. Any $\forall$-maximum

$$\cfrac{\cfrac{\begin{matrix}\Gamma\\ \Pi\\ A[a/x]\end{matrix}}{\forall x A(x)}\;\scriptstyle\forall\text{-I}}{A[t/x]}\;\scriptstyle\forall\text{-E}$$

can be levelled by the following procedure:

$$\begin{matrix}\Gamma\\ \Pi[t/a]\\ A[t/x]\end{matrix}$$

Again the levelling procedure cannot be extended to ⅁. Given that a may occur in Γ, its arbitrariness is not guaranteed.

[1] See e.g. Tennant 1978, p. 67 and van Dalen 1997, p. 96 for details.

We thus find that despite the problems faced by Dummett's notion of intrinsic harmony, it does succeed in ruling out pathological cases that both stymie Tennant's account. Can it be brought to bear also on modal operators? The answer is 'yes—for now'. A moment's reflection shows that no matter what restrictions are imposed on the introduction rule (in the case of $\Box$) or on the elimination rule (in the case of $\Diamond$), straightforward levelling procedures go through in the obvious way. However, the modal operators will likely prove ultimately to be incompatible with any of the ways of spelling out our intuitive notion of harmony for the reason sketched in section 5.5.

Intrinsic harmony thus proves to be an effective creativity constraint on elimination rules. By demanding that any local peak must be removable by judicious cutting and pasting of proofs, it provides us—in a well-motivated way—with exactly the justification of the restrictions on our quantifier rules we need. In other words, intrinsic harmony gives us a natural upper bound for the strength of elimination rules; it rules out P-strong disharmony.

8.2 The problem of P-weak disharmony

Reactivating the notion of intrinsic harmony is thus extremely useful. But of course this should not lead us to forget that there is a good reason why we must *re*introduce it in the first place. While it successfully weeds out elimination rules that provide too much deductive freedom, it offers no protection against the opposite vice, that of not fully exploiting the meaning conferred upon the constant by the introduction rule. The Q-example involving the quantum-logical elimination rule for disjunctions was but one example of this type of disharmony, i.e. P-weak disharmony. We have seen, however, that in this respect Tennant's account did us good service. The question, then, is: Can we formulate a principle of harmony that combines the best of both accounts? A principle, that is, which, when satisfied, guarantees that a true equilibrium between introductions and eliminations obtains?

The most obvious way to go about this is to form a hybrid account amalgamating Tennant's formulation and intrinsic harmony. We could take Tennant's theory of $\underline{\text{h}}$armony as basic, while slightly weakening the principle of $\underline{\text{H}}$armony. Rather than choosing the strongest elimination rule among all those for which $h(\text{\$-I}, \text{\$-e})$ obtains, the principle of $\underline{\text{H}}$armony ought to read

Weakened principle of H̲armony:

- Given the introduction rule \$-I we ought to choose the strongest elimination rule \$-E among the rules \$-e that satisfy $h(\$\text{-I}, \$\text{-e})$ *and such that (\$-I, \$-E) is intrinsically harmonious.*
- Given the elimination rule \$-E we ought to choose the strongest introduction rule \$-I among the rules \$-i that satisfy $h(\$\text{-i}, \$\text{-E})$ *and such that (\$-I, \$-E) is intrinsically harmonious.*

Insofar as Tennant's account is otherwise acceptable, this formulation patches things up satisfactorily. In order to determine whether a given pair of rules is harmonious in this sense, we first need to establish that they are both h̲armonious in Tennant's sense and, assuming that they are, that each is the strongest 'levellable' rule of the other.

The downside is that we end up with an unfortunately complicated principle of harmony. But that is not the only disadvantage. This account—like Tennant's original account—does not tell us how, when presented with one set of rules, to determine the other, harmoniously matching one. Moreover, the current formulation also makes it particularly difficult to establish whether harmony obtains, since one must be able to establish negative results. That is, to show that a given set of inference rules really is disharmonious, we have to *prove* that there can be no levelling procedure for the set of rules in question. But how are we to do that? The only solution, it seems, is to couple our weakened principle of harmony with further principles (e.g. conservativeness) that are necessary conditions for harmony and whose violation we can demonstrate. Obviously, such a solution makes the account all the more cumbersome.

The route to simplification leads us back to Read's general elimination harmony. Generalized elimination rules have the virtue of making it particularly clear to the eye how a local peak of a particular kind can be levelled. Levelling consists of reconstituting the relevant subproofs in such a way as to form a new derivation devoid of the maximum in question. Representing a maximum in the manner of Read's system has the advantage of explicitly laying out all of the requisite subproofs that are to be reconstituted in a systematic way—all the necessary raw materials are lined up and ready for our cutting and pasting procedure. Thus Read's account holds the promise of facilitating the formulation of an effective method by which we can determine the harmonious counterpart to a given rule by checking whether all the subproofs at our disposal can be rearranged so as to form a direct deductive pathway. The hope is that this method will be sufficiently reliable to enable us to establish negative results.

The task we are setting ourselves thus falls into two subtasks. In a first step—the *readability task*—we will devise a *readability procedure*. Given a set of permissible introduction or elimination rules R and certain parameters of R, the readability procedure enables us to 'read off' what we will call a *structural counterpart* of R.[2] This structural counterpart can be thought of in the first instance as a skeletal blueprint for matching elimination or introduction rules; it is the abstract form, so to speak, of the rule or set of rules that match R. Full-blooded rules that yet fit the structural mould are then generated by the imposition of various further restrictions (e.g. restrictions on collateral assumptions or on minor and/or major premises). The structural counterpart can thus be thought of as a set of rules that all share a common structural form but are each individuated by a different set of restrictions. E.g., the regular or-elimination rule and the quantum or-elimination rule are born of the same structural stock, but differentiated by the distinct sets of restrictions defining them.[3] Readability procedures can thus only ever produce a class of rules with appropriate structural properties; it does not take us all the way to harmony.

Our second subtask—the stability task—consists in providing a procedure that will enable us to determine the unique rule (or set of rules) harmonious with R. Call this rule (or set of rules) R's *harmonious counterpart*. The stability task thus consists of identifying a harmonious counterpart (if it exists) among the possible rules extractable from a structural counterpart. How do we accomplish this? Well, we know that the harmonious counterpart of R may introduce neither P-weak, nor P-strong disharmony with respect to R. As we have seen above, these constraints can be met by choosing, among the members of the class determined by our readability

[2]In order to isolate *operational* restrictions (as opposed to restrictions on discharge policies that induce different structural assumptions) it is preferable to operate in a natural deduction system in sequent format (see section 2.8).

[3]Our procedure for determining structural counterparts thus possesses at least the following advantages over Tennant's account: where Tennant offered no way of narrowing down the somewhat hazily defined class of all possible introduction/elimination rules for a given connective; our readability procedures allow us to determine a rule's structural counterpart and thus to narrow down the class of possible candidates significantly.

procedure, the strongest rule amenable to a levelling procedure. To this end we will devise a procedure—the *stability procedure*—that enables us to check whether, among the rules included in the class we called the structural counterpart, there is a rule which admits of a levelling procedure with respect to R. If there is no such rule, R has no harmonious counterpart; if R passes the test, all that is left to be done is to choose the strongest (i.e. the least restricted) among the levellable rules. Since we are choosing the strongest levellable rule, we can be assured that the rule so constructed will not be P-weakly disharmonious; because the strength of the rule is nonetheless bounded by the requirement that the pair of rules admit of a levelling procedure, we know the rule will not be P-strongly disharmonious.

The readability procedure and the stability procedure taken together thus enable us to determine, for any permissible rule (or set of rules) of inference presented to us, a harmonious counterpart, if it exists; if the rule(s) in question do not admit of an harmonious counterpart, our procedures enable us to establish this as well. We can summarize the procedure for determining harmonious counterparts for a given rule as follows. Given a rule of inference R we can determine its harmonious counterpart R' by the following steps.

1. Determine R's structural counterpart using our readability procedures.

2. Using the stability procedure determine whether R admits of a levelling procedure with respect to any of the rules in the structural counterpart.

3. If it does not, we have determined that R has no harmonious counterpart.

4. If it does, we choose the strongest rule that admits of levelling with respect to R—and this is R'.

We can think of the present approach as an algorithmic approach to harmony.

Let us begin with the readability task. We begin by analysing general elimination natural deduction systems with a view to identifying the relevant parameters that will allow us to establish structural correspondences between classes of structurally similar introduction rules and classes of structurally similar elimination rules. Our aim is to develop procedures for determining structural counterparts, on the basis of these correspondences.

8.3 The readability task I

To this end, let us return to Read's account of harmony. Read proposes a general schema for determining elimination rules on the basis of given introduction rules (Read 2000, p. 130).[4] Let us illustrate. Suppose we are given a set of introduction rules for the k-ary connective \$:

$$\text{\$-I}\ \frac{\Pi_i}{A}$$

where $1 \leq i \leq n$ and A is a formula containing \$ as its main connective.[5] Note also that the Π_i may be (non-empty) sets of sentences. Every introduction rule may thus contain multiple premises. The corresponding schematic elimination rule has the following form

$$\text{\$-GE, } i\ \frac{\begin{array}{c}\Gamma_0\\ \vdots\\ \$(A_1,\ldots,A_k)\end{array}\quad \begin{array}{c}\underbrace{\Gamma_1,[\Pi_1]^i},\ldots\\ \vdots\\ C\end{array}\quad \begin{array}{c}\ldots,\underbrace{\Gamma_n,[\Pi_n]^i}\\ \vdots\\ C\end{array}}{C}$$

As it stands, this will not do for our purposes since the schema does not leave room for elimination rules that match improper introduction rules. Improper rules, recall, are rules that contain subproofs from dischargeable assumptions as premises—rules like $\supset$-I, $\neg$-I and all generalized elimination rules. Improper rules are to be contrasted with *proper* ones, all of whose premises are sentences. We may call elimination rules matching improper introduction rules *tertiary* rules, and elimination rules for proper introduction rules *binary*.[6]

As we have already seen in the case of the generalized $\supset$-elimination rule, tertiary rules require extra columns containing derivations of the assumption involved in the

[4]Recall that even if Read's account were wholly unproblematic—which it is not—it would not be sufficient for our purposes, for our account ought to guarantee readability in both directions in conformity with our principle of functionality (see section 5.3).

[5]Read himself requires only that the \$-I rules introduce an 'occurrence of a connective' in A without demanding that this occurrence be principal, i.e. that \$ be the main connective in A (Read 2000, p. 130). Read does well not to impose the additional restriction since he later on considers multiple-conclusion systems, which do in fact countenance the introduction of logical constants into subordinate positions. However, as I argue in part three (see in particular chapters 12 and 13), there are good grounds for rejecting such systems. Unlike Read, we will therefore enforce this requirement.

[6]As we will soon see, tertiary rules comprise three types of columns, whereas binary rules require only two.

corresponding introduction rule. To accommodate elimination rules that match improper introduction rules we need to modify the above schema slightly:

$$\dfrac{\begin{matrix}\Gamma_0 \\ \vdots \\ \$(A_1,\ldots,A_k)\end{matrix} \qquad \begin{matrix}\ldots,\Gamma_i,\ldots \\ \vdots \\ \ldots B_i \ldots\end{matrix} \qquad \begin{matrix}\ldots,\underbrace{\Gamma_j,[\Pi_j]^m},\ldots \\ \vdots \\ C\end{matrix}}{C}\ \text{\$-GE}, m$$

The left-most column terminates in the major premise. For future reference, let us call it the *α-column*. We may assume for the time being that each elimination rule features only one major premise and hence contains only one α-column.

Let us call the middle columns the *β-columns*. The β-columns are required only in the case of tertiary elimination rules, β-columns contain subproofs concluding with those formulas that play the role of assumptions in the subproofs of the corresponding introduction rules. Since there may be introduction rules that involve more than one subproof, there may be several β-columns. In the binary case, there will be no β-columns.

Finally, we have the columns on the far right-hand side. Let us call them the *γ-columns*. Here it is again essential to distinguish between binary and tertiary rules. In the binary case, the γ-columns contain as assumptions the grounds necessary to introduce the major premise of the binary rule. That is, the γ-column contains all those formulas that can figure as immediate premises for the corresponding introduction rules the conclusion of which constitutes the major premise. Where an introduction rule contains more than one premise, there will be multiple assumptions in the γ-column. If there are several introduction rules there will be one γ-column for each, i.e. for each of the ways in which the major premise may be introduced. We may assume here that where there are several γ-columns the subproofs they contain all lead to the same conclusion.
We can summarize our findings for elimination rules corresponding to *proper* introduction rules by the table 8.1.[7]

Let us postpone further discussion of improper I-rules and their corresponding tertiary rules—in particular, a discussion of γ-rules in such a context—until the next

[7]Note that if we wish to eschew assumption classes containing more than one sentence (the possibility of which we signalled in chapter 7), this can be achieved by increasing the number of rules rather than the number of assumptions in accordance with the number of premises per rule.

Feature of introduction rule	**Feature of elimination rule**
number of rules	number of γ-columns
number of premises per rule	number of assumptions per γ-column

Table 8.1: Structural correspondences for proper rules

section. Let us briefly pause to remind ourselves of the motivations for the structural correspondences we have laid down so far. The number of introduction rules determines the number of γ-columns. So for instance, we have two $\vee$-introduction rules, and thus require two γ-columns, one for each disjunct. The introduction rule for $\wedge$ has two premises, but only one introduction rule; correspondingly the elimination rule will contain one γ-column with two assumptions (in the same assumption class). These structural correspondences neatly mirror our basic assumptions about harmony. In the former case, if each of $A_1, \ldots, A_m$ is a ground for asserting a statement $\$(A_1, \ldots, A_n)$ where $2 \leq n$ and $1 \leq m \leq n$, i.e. there is an introduction rule \$-I for each A_i, $1 \leq i \leq m$, such that

$$\text{\$-I}\ \frac{A_i}{\$(A_1, \ldots, A_n)}$$

then any formula C that follows from each A_i (separately) must also follow from $\$(A_1, \ldots, A_n)$ alone via \$-E. Of course in general the sets $\mathsf{C}(A_i)$ and $\mathsf{C}(A_j)$ of statements entailed by two distinct grounds A_i and A_j will also be distinct. Nevertheless, since any one of the A_i is sufficient for $\$(A_1, \ldots, A_n)$, \$-E may not permit us to infer any statements that are entailed by some but not all A_i. In other words, the elimination rule ought to license inference to those (and only those) statements contained in the intersection

$$\bigcap_{1 \leq i \leq m} \mathsf{C}(A_i)$$

It is for this reason that our requirement that all permissible elimination rules contain only γ-columns terminating with the same conclusion (which is then carried down as the overall conclusion of the elimination rule) is justified.

For consider a rule that violates our requirement:

$$\text{\$-E}\ \frac{\begin{array}{c}\Gamma_0\\ \vdots\\ \$(A, B, C, D)\end{array} \quad \begin{array}{c}\Gamma_1, [A]\\ \vdots\\ C\end{array} \quad \begin{array}{c}\Gamma_2, [B]\\ \vdots\\ D\end{array}}{D}$$

We would be unable in general to level local \$-peaks

$$\dfrac{\dfrac{\begin{array}{c}\Gamma_0\\ \vdots\\ A\end{array}}{\$(A,B,C,D)}\ \$\text{-I} \qquad \begin{array}{c}\Gamma_1,[A]\\ \vdots\\ C\end{array} \qquad \begin{array}{c}\Gamma_2,[B]\\ \vdots\\ D\end{array}}{D}\ \$\text{-E}$$

Having arrived at $\$(A, B, C, D)$ on the basis of A, we arrive at a conclusion D that may not be warranted by A. We thus find ourselves with a clear-cut case of P-strong disharmony. (Appending our proof of A to the proof of C from A and Γ_1 we only arrive at C, not D.)

Introduction rules involving more than one premise (e.g. the $\wedge$-I rule) can be given a symmetric explanation. If all of $A_1, \ldots, A_m$ are required for it to be legitimate to infer $\$(A_1, \ldots, A_n)$, \$-E should enable us to infer any statement inferable from *any one* of the A_i. In other words, the elimination rule should enable us to derive all and only those statements contained in the union

$$\bigcup_{1\leq i\leq m} \mathtt{C}(A_i)$$

from $\$(A_1, \ldots, A_n)$ alone.

Note that we can make sense of a wide range of conceivable composite logical constants by showing them to be aggregates of these two simple operations. Let \$ be an n-ary (at least binary) constant whose meaning is given by m introduction rules for $1 \leq m \leq n$. Each of these rules has the form

$$\dfrac{A_1, \ldots, A_k}{\$(A_1, \ldots, A_n)}\ \$\text{-I}$$

where $k \leq m$ and it is understood that k is a function f of the rule in question. The matching elimination rule should then license the inference to those and only those statements that are in the intersection

$$\bigcap_{1\leq i\leq m} B_i$$

where B_i is the union of the deductive consequences of all the $f(i) = k$ premises involved in the i-th introduction rule, i.e.

$$B_i = \bigcup_{1 \leq j \leq f(i)} \mathtt{C}(A_j)$$

Correspondingly, our elimination rule should contain m γ-columns, one for each introduction rule. A column corresponding to rule i ($1 \leq i \leq m$), say, will then contain $f(i) = k$ assumptions (in one assumption class).

Note that our quantifiers can readily be seen to be infinitary extensions of these elementary operations. The quantifier restrictions permit us to capture the notion of an arbitrary parameter, which makes it possible to represent these infinite extensions in a finitely graspable, and hence proof-theoretically palatable, way.[8]

8.4 The readability task II

Things become slightly more complicated in the case of tertiary elimination rules. The γ-column fulfils a somewhat different function in elimination rules of this type. While the β-columns contain proofs of the assumptions of the corresponding introduction rule, the γ-column takes as its initial assumption the statement that constitutes the conclusion of a subproof of the corresponding introduction rule. In the paradigmatic case of $\supset$

$$\dfrac{\begin{array}{c}\Gamma, [A]^i \\ \vdots \\ B\end{array}}{A \supset B}\ {\scriptstyle \supset\text{-I},\, i}$$

the provisional conclusion B figures as an assumption in the γ-column of the corresponding elimination rule, whereas a proof of A figures in the β-column.

In line with our previous explanations, we may think of the $\supset$-GE rule as enabling us to draw all the deductive consequences of B from $A \supset B$ alone, provided that we have shown B to be a consequence of A. What the introduction rule shows us is that we are entitled to assert $A \supset B$ once we established that $B \in \mathtt{C}(A)$. From this it follows that $\mathtt{C}(B) \subseteq \mathtt{C}(A)$: anything that follows from B can now be assumed

[8] See (Read 2000, p. 136–138) for a lucid discussion.

to follow from A alone (supposing we are taking our deducibility relation to be unrestrictedly transitive). And this aligns exactly with what $\supset$-GE 'says'.

Okay, but surely there are other improper introduction rules, which may be rather different from the paradigm case of the conditional. How do we deal with them? A special case worth mentioning is that of negation. Here instead of some statement B, the introduction rule states that $\bot \in \mathsf{C}(A)$.[9] Following the analogy from the case of the conditional we would expect $\neg$-GE to have the following form:

$$\neg\text{-GE}, i \; \dfrac{\begin{matrix} \Gamma_0 \\ \vdots \\ \neg A \end{matrix} \quad \begin{matrix} \Gamma_1 \\ \vdots \\ A \end{matrix} \quad \begin{matrix} [\bot]^i \\ \vdots \\ B \end{matrix}}{B}$$

However, we are here confounding the (operational) elimination rule for negation and the (structural) *ex falso* rule, and incorporating the structural rule into our elimination rule. We do well to keep these matters separate.[10] We can do this simply by re-presenting our standard $\neg$-elimination rule in the form of a generalized elimination rule:

$$\neg\text{-GE}, i \; \dfrac{\begin{matrix} \Gamma_0 \\ \vdots \\ \neg A \end{matrix} \quad \begin{matrix} \Gamma_1 \\ \vdots \\ A \end{matrix} \quad [\bot]^i}{\bot}$$

The general elimination rule for $\neg$ thus coincides by and large to the regular negation-elimination rule.

We said above that we would allow introduction rules containing more than one subproof. Indeed the introduction rule for the biconditional is naturally presented in this way:

$$\leftrightarrow\text{-I}, i \; \dfrac{\begin{matrix} \Gamma, [A]^i \\ \vdots \\ B \end{matrix} \quad \begin{matrix} \Gamma', [B]^i \\ \vdots \\ A \end{matrix}}{A \leftrightarrow B}$$

[9]As we mentioned in section 2.7, I do not take $\bot$ to be a 'propositional constant' standing proxy for any contradiction (expressible in the language in question) whatsoever. Rather, following Tennant, I consider $\bot$ to be a logical punctuation mark that signals the occurrence of a 'logical dead-end' (Tennant 2004a, p. 8). Tennant shows (in ibid.) how an appeal to $\bot$ can be eschewed altogether by giving a simultaneous inductive definition of the notions of proof and of disproof. Having shown that $\bot$ can be dispensed with in this way, we can happily return to appealing to $\bot$ to characterize $\neg$.

[10]See our discussion of *ex falso*'s status as a structural rule in section 2.7.

What form should the corresponding elimination rule take? Note that because we are only considering single-conclusion rules, there can only be one γ-column. How, then, can we deal with $\leftrightarrow$-I? The answer is that there must be two distinct generalized elimination rules: one with A in the γ-column, one with B in the γ-column. Given that the introduction rule contains two subproofs, each elimination rule must contain a proof of the appropriate assumption in its β-column. $\leftrightarrow$-GE must therefore have the form

$$\dfrac{\begin{matrix}\Gamma_0 \\ \vdots \\ A \leftrightarrow B\end{matrix} \qquad \begin{matrix}\Gamma_1 \\ \vdots \\ A\end{matrix} \qquad \begin{matrix}\Gamma_2, [B]^i \\ \vdots \\ C\end{matrix}}{C}\ {\scriptstyle \leftrightarrow\text{-GE},\ i}$$

The other elimination rule is identical except that we substitute B for A in the β-column and A for B in the γ-column. Thus we may have improper introduction rules containing several subproofs which need not lead to the same conclusion.

The last example raises the question of whether we can establish structural correspondences between improper introduction rules and their corresponding elimination rules, as we did above for the simpler rules. Let us, for purposes of illustration, examine two possible improper introduction rules. Consider first $\ominus$, governed by the following introduction rules:

$$\dfrac{\begin{matrix}\Gamma, [A]^i \\ \vdots \\ C\end{matrix}}{\ominus(A, B, C)}\ {\scriptstyle \ominus\text{-I},\ i}$$

$$\dfrac{\begin{matrix}\Gamma, [B]^i \\ \vdots \\ C\end{matrix}}{\ominus(A, B, C)}\ {\scriptstyle \ominus\text{-I},\ i}$$

Given that a statement $\ominus(A, B, C)$ may be inferred from a proof of C *either* from A or from B, the elimination rule must contain *both* a proof of A and of B to ensure that all $\ominus$-peaks can be levelled. Moreover, the γ-column must feature C as an assumption. We thus arrive at the following elimination rule:

$$\dfrac{\begin{matrix}\Gamma_0 \\ \vdots \\ \ominus(A,B,C)\end{matrix} \quad \begin{matrix}\Gamma_1 \\ \vdots \\ A\end{matrix} \quad \begin{matrix}\Gamma_2 \\ \vdots \\ B\end{matrix} \quad \begin{matrix}\Gamma_3, [C]^i \\ \vdots \\ D\end{matrix}}{D}\ \ominus\text{-GE}, i$$

Now consider the ternary connective $\circledcirc$, associated with the following introduction rule containing two subproofs:

$$\dfrac{\begin{matrix}\Gamma, [A]^i \\ \vdots \\ C\end{matrix} \quad \begin{matrix}\Gamma', [B]^i \\ \vdots \\ C\end{matrix}}{\circledcirc(A,B,C)}\ \circledcirc\text{-I}, i$$

The statement $\circledcirc(A,B,C)$ is assertible whenever we possess *both* a deduction of C from A and from B. In order to be able to eliminate $\circledcirc$-peaks, it is thus sufficient that we dispose of *either* a proof of A or a proof of B. By this reasoning we arrive at the following elimination rules,

$$\dfrac{\begin{matrix}\Gamma_0 \\ \vdots \\ \circledcirc(A,B,C)\end{matrix} \quad \begin{matrix}\Gamma_1 \\ \vdots \\ A\end{matrix} \quad \begin{matrix}\Gamma_2, [C]^i \\ \vdots \\ D\end{matrix}}{D}\ \circledcirc\text{-GE}, i$$

$$\dfrac{\begin{matrix}\Gamma_0 \\ \vdots \\ \circledcirc(A,B,C)\end{matrix} \quad \begin{matrix}\Gamma_1 \\ \vdots \\ B\end{matrix} \quad \begin{matrix}\Gamma_2, [C]^i \\ \vdots \\ D\end{matrix}}{D}\ \circledcirc\text{-GE}, i$$

Taking a closer look at what the introduction rules in question express we find that our choices are justified. If $\ominus$ 'expresses' the fact that all of C's consequences also are consequences of *either* A or B (i.e. $\mathsf{C}(C) \subseteq \mathsf{C}(A) \cup \mathsf{C}(B)$), $\circledcirc$ 'tells us' that anything that follows from C also follows from *both* A and B (i.e. $\mathsf{C}(C) \subseteq \mathsf{C}(A) \cap \mathsf{C}(B)$). In order to arrive at a formula D that follows from C directly without a detour through $\ominus(A,B,C)$, possessing a proof of just one of A or B will generally not be enough. In the case of $\circledcirc$, on the other hand, we are guaranteed that a proof of either one of A and B will give us access to D. Following this pattern, the most natural policy to adopt is that, since the proofs of A and B are both required in the case of $\ominus$, both should be grouped together in the elimination rule, as in $\ominus$-GE. By contrast, there should be two $\circledcirc$-GE rules, one corresponding to each of the subproofs in the introduction rule.

Is our policy viable? The crucial test will be whether it allows us to handle constants combining both of the features exemplified by our two previous case studies: namely constants governed by several introduction rules some of which contain more than one subproof. Let us consider a simple case to examine how such cases can be dealt with. Consider the five-place connective $\circledast$ governed by the rules

$$\circledast\text{-I}, i\ \dfrac{\begin{array}{cc}\Gamma, [A]^i & \Gamma', [B]^i \\ \vdots & \vdots \\ E & E\end{array}}{\circledast(A, B, C, D, E)}$$

$$\circledast\text{-I}, i\ \dfrac{\begin{array}{cc}\Gamma, [C]^i & \Gamma', [D]^i \\ \vdots & \vdots \\ E & E\end{array}}{\circledast(A, B, C, D, E)}$$

How do we avail ourselves of a $\circledast$-elimination rule? Can we apply our policy to the present case? In line with our explanations above, we find that the connective $\circledast$ expresses the fact that

$$\mathsf{C}(E) \subseteq (\mathsf{C}(A) \cap \mathsf{C}(B)) \cup (\mathsf{C}(C) \cap \mathsf{C}(D))$$

Applying our policy to $\circledast$ we find that a $\circledast$-elimination rule would have to contain a proof of one of A and B and a proof of one of C and D. It will thus contain two β-columns: one containing a member of the former pair, the other a member of the latter pair. This leaves us with four elimination rules

$$\circledast\text{-GE}, i\ \dfrac{\begin{array}{cccc}\Gamma_0 & \Gamma_1 & \Gamma_2 & \Gamma_3, [E]^i \\ \vdots & \vdots & \vdots & \vdots \\ \circledast(A, B, C, D, E) & A & C & F\end{array}}{F}$$

$$\circledast\text{-GE}, i\ \dfrac{\begin{array}{cccc}\Gamma_0 & \Gamma_1 & \Gamma_2 & \Gamma_3, [E]^i \\ \vdots & \vdots & \vdots & \vdots \\ \circledast(A, B, C, D, E) & A & D & F\end{array}}{F}$$

Feature of introduction rule	**Feature of elimination rule**
Proper rules:	
number of rules	number of γ-columns
number of premises per rule	number of assumptions per γ-column
Improper rules:	
number of subproofs per rules	number of elimination rules
number of rules	number of β-columns

Table 8.2: Structural correspondences for proper and improper rules

$$\circledast\text{-GE}, i\ \dfrac{\begin{matrix}\Gamma_0 \\ \vdots \\ \circledast(A,B,C,D,E)\end{matrix} \quad \begin{matrix}\Gamma_1 \\ \vdots \\ B\end{matrix} \quad \begin{matrix}\Gamma_2 \\ \vdots \\ C\end{matrix} \quad \begin{matrix}\Gamma_3, [E]^i \\ \vdots \\ F\end{matrix}}{F}$$

$$\circledast\text{-GE}, i\ \dfrac{\begin{matrix}\Gamma_0 \\ \vdots \\ \circledast(A,B,C,D,E)\end{matrix} \quad \begin{matrix}\Gamma_1 \\ \vdots \\ B\end{matrix} \quad \begin{matrix}\Gamma_2 \\ \vdots \\ D\end{matrix} \quad \begin{matrix}\Gamma_3, [E]^i \\ \vdots \\ F\end{matrix}}{F}$$

Generalizing our findings we may now complete our table of structural correspondences in the table 8.2

We may thus conjecture that with table 8.2 we possess all the information required to determine the harmonious counterpart to any given permissible introduction/elimination rule, thus validating our principle of functionality (see section 5.3). This completes the readability task we had set ourselves. However, before we turn to the remaining task—the stability task—it is worth pausing to point out an interesting feature of our analysis.

8.5 Composite and primitive rules

We have tried to ensure that our readability procedures are applicable to a broad range of inference rules. This is not to say that we have not made substantive assumptions about the permissible form inference rules may take. We assumed, in the form of the two-aspect model of meaning, that inference rules fall into exactly one of two disjoint camps, introductions and eliminations. Moreover, we assumed that all rules of inference feature only one logical constant. This assumption corresponds to the principle of separability, which we will discuss in chapter 13.4. Finally, we assumed that there are no idle wheels. That is, we took it for granted that all of the sentences mentioned in the conclusion of an introduction rule or in the major premise of an elimination rule have an active role to play. Our principle of functionality stands and falls with this assumption. Consider for example the following rules of inference for the tertiary connective \$:

$$\text{\$-I}\ \frac{A}{\$(A, B, C)}$$

$$\text{\$-I}\ \frac{B}{\$(A, B, C)}$$

Assume that these rules are taken to determine the meaning of \$. By the correspondences we set out above, the structural counterpart for \$ would be identical to that of our standard $\vee$-elimination rule. Assuming we had our stability procedure in place as well, the two procedures—readability and stability—would determine the same rule, $\vee$-elimination, thereby violating the principle of functionality. Moreover, going the other way, i.e. starting with the $\vee$-elimination rule, our procedures would be incapable of picking out a unique introduction rule. Rather, we would be faced with an indeterminate number of matching introduction rules: each introduction rule, though identical in form and sharing the same premises, would give rise to a different conclusion $\$(A, B)$, $\$(A, B, C)$, $\$(A, B, C, D), \ldots$. However, C (and D, etc.) are but idle wheels in the statement $\$(A, B, C)$. Their presence does not manifest itself in any way in the inferential behaviour of the constant in question. Since these inert components make no detectable difference to our use of \$, they cannot be semantically significant. Thus, we may discard any constants containing idle wheels.

If we are prepared to grant these assumptions (at least for the time being), we find that the correspondences we have established hold also between complex introduction and elimination rules. Yet, as closer examination reveals, most of these possible logical constants are unnecessary: they turn out to be governed by redundant *composite* introduction rules. All we really need to worry about are the constants that are given by *primitive* introduction rules. But what makes an introduction rule primitive? I want to suggest that a helpful notion of primitiveness can be defined in terms of the explanations offered in the previous two sections. The logical constants $\vee, \wedge$ and $\supset$ can be understood as the set-theoretic operations of, respectively, intersection, union and the subset relation on sets of consequences. Negation can be understood as a special case of the subset relation and the quantifiers may be thought of as generalizations of disjunction and conjunction. The primitive operations, on this picture, just amount to the most basic set-theoretic operations: they are the basic building blocks out of which complex connectives corresponding to complex set-theoretic operations can be constructed.

Our hypotheses are supported by the above examples. Our analyses of sentences like $\ominus(A, B, C)$ and $\circledcirc(A, B, C)$ as complex set-theoretic operations neatly corresponds to the observation that these seemingly simple logical constants in fact correspond to logically complex sentences: $(A \supset C) \vee (B \supset C)$ (or equivalently $(A \wedge B) \supset C$) and $(A \supset C) \wedge (B \supset C)$ (or equivalently $(A \vee B) \supset C$) respectively. Both rules are *composite* rules of inference in the sense that they are aggregate rules composed of elementary operations. The elementary or *primitive* rules of inference are just the ones that cannot be decomposed into simpler rules.

What is meant by 'decomposition' here? Take the operator $\circledcirc$. A sentence of the form $\circledcirc(A, B, C)$ may be asserted when it can be shown that C can be deduced from A and from B. But this means that an assertion of $\circledcirc(A, B, C)$ is warranted just in case both $A \supset C$ and $B \supset C$ are assertible, i.e. exactly when $(A \supset C) \wedge (B \supset C)$ is assertible. The $\circledcirc$-I rule thus in fact bundles two inferential steps into one. It masks inferential structure. Similarly, on the basis of the premises of $\circledcirc$-I, we could have arrived at $(A \vee B) \supset C$ by the following deduction

$$\vee\text{-E}, i\ \dfrac{[A \vee B]^j \qquad \begin{matrix}\Gamma_1, [A]^i \\ \vdots \\ C\end{matrix} \qquad \begin{matrix}\Gamma_2, [B]^i \\ \vdots \\ C\end{matrix}}{\supset\text{-I}, j\ \dfrac{C}{(A \vee B) \supset C}}$$

These cases show that the composite introduction rule for $\circledcirc(A, B, C)$ conceals potentially essential information concerning the inferential fine-structure of the proof in question. Our analysis in terms of set-theoretic operations brings this to light.

It may of course be retorted that the system and associated language in which we operate may simply lack the requisite resources to express primitive operations. We could certainly imagine a people who adopted logical constants governed by composite introduction rules and yet did not have the relevant elementary operators in their repertoire. Surely logic should not have any say in what expressions may or may not figure in our language. In particular, then, it may not mandate the presence of certain primitive logical operations. Does this mean that such a people deploys potentially disharmonious inference rules? Are the expressions whose meanings are so governed not really logical?

The response, surely, must be that, so long as the complex constants used by our imagined logical community are decomposable into harmonious ones governed by primitive introduction rules, they are indeed deploying genuine logical constants. But this does not change the fact that, from a logical point of view, our concern should be to identify the elementary inferential transitions. This point was famously put by Frege who argued that the logician's foremost task is to analyse complex arguments into

> a chain of deductions with no link missing, such that no step in it is taken which does not conform to some one of a small number of principles of inference recognized as purely logical. To this day scarcely one single proof has ever been conducted on these lines; the mathematician rests content if every transition to a fresh judgement is self-evidently correct, without inquiring into the nature of this self-evidence, whether it is logical or intuitive. A single such step is often really a whole compendium, equivalent to several simple inferences, and into it there can still creep along with these some element from intuition. In proofs as we know them, progress is by jumps (Frege 1884, §90).

What goes for Frege's logicism applies equally to our project of elucidating the concept of harmony. We have seen it in our discussion of immediate and mediate inferential transitions in section 2.5. Indeed we can do better than Frege. Contrary to Frege, who identified the 'purely logical' principles of inference with a somewhat random collection of axioms, our choice of primitive operators is well-motivated. Primitive operators enjoy meaning-theoretic priority, since unlike complex constants their meanings cannot be explained compositionally in terms of simpler operators. These primitive operators are also *self-justifying* in the sense that any logical constant properly so-called can be constructed out of them.

Note that the class of primitive constants contains most of the expressions traditionally treated as logical: conjunction, disjunction, the universal and existential quantifiers, negation and the conditional, although the primitive negation operator and conditional are intuitionistic. Moreover, the class of primitive constants does not include the truth predicate or the modal operators. On the other hand, certain operators may qualify as primitive that are not usually dignified with a place in our laundry list of logical constants. For instance, the t-operator, the truth predicate's banal cousin, which we encountered in section 4.7, is primitive according to our definition. It is an interesting question whether our notion of primitiveness could constitute a passable criterion of logicality. But we shall not try to explore it here. All that matters for the moment is that the distinction between composite and primitive constants gives us a principled reason to focus our attention on a restricted class of logical constants.[11]

8.6 The stability task I

Let us return to our main business. In sections 8.3 and 8.4 we accomplished the first part of our task, the readability task. We now face the task of devising a procedure for testing whether a given pair of rules is levellable or—and this is equally or more important—unlevellable. In devising our procedure it will again be useful to treat binary and tertiary inference rules separately.

Suppose the k-ary constant \$ has n introduction rules associated with it. We proceed by creating a maximum for each of the n rules. We do this by appending each introduction rule with the appropriate elimination rule (as determined by our readability procedure) and then considering each maximum in turn.[12] The maximum corresponding to rule i ($1 \leq i \leq n$) has the following form:

$$\text{\$-GE}, r\;\dfrac{\text{\$-I}\;\dfrac{\begin{matrix}\Gamma_0^1 \dots & \Gamma_0^m\\ \Pi_1 \dots & \Pi_m\\ A_1 \dots & A_m\end{matrix}}{\$(A_1,\dots,A_k)} \qquad \begin{matrix}\Gamma_1, \underbrace{[A_1^1,\dots,A_{f(1)}^1]^r}\\ \Sigma_1 \cdots \\ C \cdots\end{matrix} \qquad \begin{matrix}\Gamma_n, \underbrace{[A_1^n,\dots,A_{f(n)}^n]^r}\\ \cdots \Sigma_n\\ \cdots C\end{matrix}}{C}$$

Our elimination rule contains n β-columns, one for each introduction rule, among them the column associated with the ith rule. Note that the number of premises is a function f of the introduction rule in question. We thus have $f(i) = m$. Given

[11] Our distinction justifies some of the simplifications we allow ourselves in the next sections. In particular, because of it we need not worry about mixed introduction rules that contain subproofs and immediate premises.

[12] Recall that according to our structural correspondences binary rules only ever contain a single elimination rule.

our policy of not admitting idle wheels, the total number of premises occurring in all the β-columns must equal k, the arity of the connective \$; more pedantically,

$$\sum_{1 \leq i \leq n} f(i) = k$$

Our levelling procedure then simply consists in matching the proofs $\Pi_1, \ldots, \Pi_m$ of the premises $A_1, \ldots, A_m$ from the hypotheses $\Gamma_0^1, \ldots, \Gamma_0^m$ with the corresponding proof Σ_i of C from the same premises and possibly collateral hypotheses Γ_i. The resulting proof has the form

$$\begin{array}{c} \Gamma_0^1 \ldots \quad \Gamma_0^m \\ \Pi_1 \ldots \quad \Pi_m \\ \Gamma_i, \underbrace{A_1 \ldots \quad A_m} \\ \Sigma_i \\ C \end{array}$$

To summarize: binary rules of inference are harmonious in our sense if

- our levelling procedure can be carried out for each of the n introduction rules;
- neither introduction nor elimination rules impose any restrictions, other than those necessary in order to ensure levellability.

8.7 The stability task II

Once again things are slightly more complicated in the case of tertiary rules. Yet the point of departure is strictly analogous to the previous case: we begin by running the readability procedure to find the structural counterpart to the rule we are given. Suppose we are dealing with n introduction rules, each of which contains m subproofs.[13] For simplicity we may assume that all subproofs lead to the same

[13] One could also consider cases where the number of subproofs per introduction rules varies; i.e. cases in which the number of subproofs would be a function f of the rule in question. The total number of elimination rules would then amount to

$$\prod_{1 \leq i \leq n} f(i)$$

For simplicity, we shall ignore such cases. The considerations in the previous section assure us that there will be no loss of generality in doing so. The same goes for the other simplifying assumptions we are making here.

conclusion C. Moreover we shall suppose that all assumptions are distinct.[14] The operation \$ may thus be represented as

$$\mathtt{C}(C) \subseteq \bigcup_{1 \leq i \leq n} B_i,$$

where

$$B_i = \bigcap_{1 \leq j \leq m} \mathtt{C}(A_j).$$

Each elimination rule will contain n β-columns. Each column corresponds to one of the introduction rules, i say, and contains a proof of one of the m assumptions featuring in the ith introduction rule. How many elimination rules will there be in total? On assuming that each introduction rule has m subproofs per rule, there would be m^n elimination rules.

We then arrive at our levelling procedure as follows. As before, we consider all the possible \$-maxima of the form

$$\dfrac{\dfrac{\begin{matrix}\Gamma_0^1, [A_1] & \dots & \Gamma_0^m, [A_m] \\ \Pi_1 & \cdots & \Pi_m \\ C & \dots & C\end{matrix}}{\$(A_1, \dots, A_k, C)}\,\text{\$-I} \qquad \begin{matrix}\dots \Gamma_j \dots \\ \dots \Sigma_j \dots \\ \dots A_j \dots\end{matrix} \qquad \begin{matrix}\Gamma_{n+1}, [C]^r \\ \Omega \\ D\end{matrix}}{D}\,\text{\$-GE}, r$$

Here $k = m \times n$.

We survey all of the β-columns in every elimination rule until we find a proof Σ_j of A_j which we graft onto the proof with A_j from which (together with hypotheses Γ_0^j) we obtain a proof of C. Concatenating the resulting proof with the proof of D from C and hypotheses Γ_{n+1} then gives us the desired result

$$\begin{matrix} \Gamma_j \\ \Sigma_j \\ \underbrace{\Gamma_0^j, A_j} \\ \Pi_j \\ \underbrace{\Gamma_{n+1}, C} \\ \Omega \\ D \end{matrix}$$

[14] These assumptions are not essential. Generalizing is just an exercise in combinatorics that I happily leave as an exercise to more competent readers.

For the sets of \$-introductions and \$-eliminations to be in harmony, it must be the case that for every \$-introduction rule there is an elimination rule containing the appropriate β-column. If the search for such a β-column yields no results for at least one of the introduction rules, it follows that no levelling procedure is to be found. Unsuccessful systematic searches thus amount to proofs of the non-existence of a levelling procedure. The procedure we have described likens determining 'levellability' to a game of solitaire. The challenge is to arrange the columns in such an order as to create a complete deductive path from the uppermost hypotheses to the final conclusion. When this can be accomplished, the rules can be levelled.

As in the case of binary rules, we can summarize our findings by the following two points. Tertiary rules of inference are in harmony if

- our levelling procedure can be carried out for each of the n introduction rules;
- neither introduction nor elimination rules impose any restrictions, other than those necessary in order to ensure levellability.[15]

Conjoining our readability procedures and our stability procedures, we are now able to 'read off' the harmonious counterpart (if it exists) of any permissible rule of inference presented to us; in the absence of any harmonious counterpart for the rule in question our procedures enable us to establish this rigorously. Moreover, of course, our procedures jointly enable us to establish, for any given pair of inference rules (or sets thereof), whether or not they are harmonious.

8.8 Summary

In this chapter we have put forth a novel account of harmony. It solves the two major difficulties which have, in one way or another, bedevilled all the existing accounts: the problems of P-weak disharmony and of P-strong disharmony. P-strong disharmony is avoided by requiring that a pair of rules admit of a levelling procedure; P-weak disharmony is avoided by the proviso that we may not impose any restrictions on the rule in question apart from those necessitated by the requirement of levellability. What is more, our account of harmony—with its readability and stability procedures—enables us to read off algorithmically the harmonious counterparts of (permissible) inference rules. Lastly, our stability procedures enable us to

[15] As is the case of the rule for $\forall$.

establish conclusively negative results; i.e. to prove that no harmonious counterpart exists for a rule or that a pair of rules does not admit of a levelling procedure.

Part III

Proof-theoretic arguments

Chapter 9

Introduction

9.1 Proof-theoretic arguments

In the previous part, we offered a thorough analysis of the notion of harmony and proposed a new and improved version of a principle of harmony. With our revised principle on board, we will now turn to the fundamental question, 'Which logic is the correct one?' More precisely, we will be concerned with a particular approach to this question, an approach that has been articulated within the context of the realism/anti-realism debate.

The realist is someone who holds that the central ingredient in any explanation of meaning is that of classical mind-independent truth. The anti-realist, as we will understand him, is a revisionary who favours an epistemically constrained notion of truth and advocates broadly intuitionistic reforms of our logic.

The dispute between realists and anti-realists is closely linked to a dispute over which proof-theoretic framework be given priority, whether, e.g. natural deduction systems or sequent calculi should be preferred. Let me explain. The anti-realist assumes that the use we make of the logical constants, as it is manifested in our inferential practice, can be adequately represented in Gentzen-Prawitz-style natural deduction systems. Our meaning-giving inferential practice can be profitably summarized by the introduction and elimination rules associated with each operator. Natural deduction systems may thus be regarded as theories of meaning (in roughly Dummett's sense of the term) of our logical vocabulary: the inference rules governing a given operator state what a speaker needs to grasp in order to master that operator's meaning. Our task will largely be to determine to what extent the emphasis on natural deduction systems is defensible on the basis of the central inferentialist premises laid out in part one.

Authors like Dummett, Prawitz and Tennant have elaborated well-known arguments against classical logic on the basis of their inferentialism together with the said bias for natural deduction-style representations of our inferential practices. The unifying thought is this. Because natural deduction inference rules specify the meanings of the logical operators, they must be subject to the same constraints that regulate any viable theory of meaning. Such general meaning-theoretic considerations are brought to bear on natural deduction systems in the form of constraints on the form of acceptable inference rules. For example, languages must be learnable, hence the rules characterizing the meanings of the logical constants must be finitely stateable. Language is molecular, so the principle of separability must hold; i.e. it must be possible to isolate the meaning of a particular logical constant (or at least of a small group of constants).[1] Or again, and this point is of particular importance here, languages ought to be semantically well-behaved (or should at least strive towards such a state), therefore they must satisfy the constraints of harmony. A language is semantically well-behaved if the circumstances of correct application of each expression are in equilibrium with the consequences of having so employed that expression, which in turn, as we have seen, goes hand-in-hand with compositionality. If we accept this inferentialist framework and implement it within a natural deduction setting, we find that the meanings of the classical logical constants do not pass muster. The rules concerning negation, in particular, fail to satisfy the principle of harmony.

I call arguments promoting revision of our logical practice in this way *proof-theoretic arguments*. If sound, such arguments establish that classical logicians fail to attach coherent meanings to the logical expressions and indeed that no system of logic stronger than intuitionistic logic can receive a proof-theoretic justification.[2] Proof-theoretic arguments thus demonstrate the inadmissibility of certain forms of inference relative to our meaning-theoretic principles, i.e. relative to the fundamental principles underlying any coherent theory of meaning as they are represented in a natural deduction system (e.g. Dummett 1991, p. 303). It is in this way—by demonstrating which forms of inference can be justified by proof-theoretic means in

[1]The principle of separability will be defended in section 13.4.

[2]Tennant has argued for more thoroughgoing revisions of our logic; he endorses the adoption of the even weaker intuitionistic relevant logic.

accordance with meaning-theoretic strictures—that proof-theoretic arguments purport to provide a basis for the criticism of our existing logical practice and so bear on my opening question, 'Which is the correct logic?'

There are a number of ways realists might try to resist the conclusions of the proof-theoretic argument. Of course the realist may wish to register objections already at the level of the anti-realist's premises: he may find fault with the inferentialist commitments undertaken by the anti-realist; or he may cast doubt on the anti-realist's meaning-theoretic assumptions. But let us assume for present purposes that we are dealing with a defender of classical logic with inferentialist sympathies who buys into the general meaning-theoretic framework set out by the anti-realist. We assume, that is, that our realist believes that our ordinary logical practice is amenable to a systematic, explicit representation within a proof system and that, codified in this way, we ought to assess it in the light of the demands imposed by our meaning-theoretic principles. More precisely, our realist is a logical inferentialist. In particular he buys into our assumption of minimal molecularism (see section 2.4) and to the two-aspect doctrine of meaning (see section 2.1). Being a realist, however, he will maintain that classical logic will resurface unscathed from this examination.

How can a realist of this ilk accept the premises of the proof-theoretic argument while avoiding its revisionary conclusions? The key to the answer might reasonably be thought to lie in the anti-realist's choice of proof-theoretic framework. It is critical to the success of proof-theoretic arguments that they establish their results *without loss of generality*. That is, proof-theoretic arguments ought not to be sensitive to the particular structural properties of the proof-theoretic framework in which they are formulated, i.e. the specific structural features of natural deduction systems. Yet pretty much everything we have said about harmony so far has assumed the background of a natural deduction setting. Our account of harmony was specifically tailored to such systems. Consequently, a promising avenue for the realist might be to accept the premises of the proof-theoretic argument but simply opt for a different proof system; namely, a proof system that validates classical logic. After all, what reasons do we have for thinking, as the anti-realist would have it, that natural deduction systems enjoy a privileged status among all of the deductive systems on the market?

Indeed several authors adhering to classical logic and yet subscribing to inferentialism have expressed doubts concerning the privileged status attributed to the format of natural deduction by anti-realists.[3] Such scepticism turns into outright rejection when natural deduction-based accounts of the meanings of the logical constants culminate in a call for reforms of our logical practice. Surely, the realist insists, the fact that classical logic fails to satisfy proof-theoretically articulated meaning-theoretic constraints must tell against the format chosen by the anti-realist—not against our time-honoured classical logic!

The onus is thus on the anti-realist to explain why systems of natural deduction should be regarded as providing the adequate formal correlates to the constraints imposed by our theory of meaning in spite of the highly counterintuitive results they give rise to. Otherwise, the revisionary moral the anti-realist draws from the proof-theoretic argument must appear like an avoidable consequence of his choice to carry out the argument in a natural deduction setting. The realist, on the other hand, must present us with acceptable alternative proof-theoretic frameworks: his task is to find a proof system that fits the bill while validating classical logic. Given the preponderance of classical reasoning in ordinary as well as mathematical contexts, an inferentialist account that does away with the need for reform is *prima facie* to be preferred.

9.2 The project

The sequent calculus is the most obvious candidate for the realist here. The proof-theoretic argument crucially relies on the demonstration that each of the possible ways in which the intuitionistic natural deduction system (NJ) can be extended to yield its classical counterpart (NK) violates at least one of the meaning-theoretic constraints: NK is obtained from NJ by adjoining to the latter one of the characteristically classical rules of inference—the law of excluded middle, *reductio ad absurdum*, classical dilemma, double negation elimination, etc.—all of which, the anti-realist claims, fall foul of meaning-theoretic requirements. However, it is precisely the said sequent calculus, Gerhard Gentzen's other innovation, that casts doubt on the anti-realist's argument for reform.[4] Having introduced the sequent

[3] To name but a few: Došen (1994), Hacking (1979), Kneale (1956), Kremer (1988), Read (2000) and Restall (2005) all fall into this category (albeit in different ways).

[4] It should be noted that, although I make use of Gentzen's terminology, I shall slightly depart from his presentation of sequent systems (Gentzen 1969a, p. 81). I take the *relata* of the relation denoted by ':' to be sets of statements rather than sequences. We may thus dispense with the

calculus, Gentzen observes that in this system it is possible to move between the classical variant (LK) and the intuitionistic one (LJ) simply by requiring that in the intuitionistic case, succedents be restricted to at most one formula.[5] As Gentzen immediately recognizes,

> the distinction between intuitionistic and classical logic is, externally, of a quite different type in the calculi LJ and LK from that in the calculi NJ and NK. In the case of the latter, the distinction is based on the inclusion or the exclusion of the law of the excluded middle [or any of the other rules mentioned] whereas for the calculi LJ and LK the difference is characterized by the restriction on the succedent (Gentzen 1969a, p. 86).

For this reason sequent calculi are often said to lend themselves better to the formalization of classical logic than natural deduction systems.[6] From the point of view of the proof-theoretically-minded classicist, the sequent calculus offers a promising framework: if the formalization of classical logic afforded by the standard sequent calculi indeed turns out to be meaning-theoretically acceptable, the classicist could rightly claim to have neutralized the proof-theoretic argument and so to have successfully defended classical logic.[7] However, the meaning-theoretic legitimacy of the classical sequent system needs to be demonstrated; elegance alone is not enough. Any rival system must satisfy three conditions:

1. it must deliver an inferentialistically acceptable account of the classical meanings of the logical constants;

structural rules of interchange and of contraction—rules that are irrelevant for our purposes. Also, I will allow myself to speak somewhat loosely of, say, '*the* intuitionistic natural deduction system' in the singular, even though there is, strictly speaking, a multitude of systems, all of which are adequate for intuitionistic logic.

[5] I will speak of the *succedent* of a sequent to designate the set on the right-hand side of the sequent sign in order to distinguish it from the overall conclusion of the derivation, which is itself a sequent rather than a set. I shall use the slightly misleading adjective 'multiple-succedent' to refer to systems that allow for succedents of cardinality greater than one. Where the context makes this sufficiently clear, I will also at times speak of 'multiple-conclusions' especially in contexts where both multiple-succedent systems and other types of multiple-conclusion systems are at issue.

[6] E.g., Bostock 1997, Cook 2004, Hacking 1979 and Read 2000.

[7] Classical accounts of the meanings of the logical constants along inferentialist lines and based on sequent systems have been proposed by a number of authors including Došen (1994), Hacking (1979), Kremer (1988) and Read (2000).

2. it must allow for the expression of the meaning-theoretic constraints that bear on these meanings (in particular, it must be possible to formulate a principle of harmony for it);

3. and the classical meanings must be compatible with these constraints.

Only if the classical sequent calculus meets all three conditions will the realist have vindicated classical logic on the basis of the proof- and meaning-theoretic assumptions we have been considering.

The plan is the following. In the next chapter, I will be concerned with the question of how the notion of harmony might be applicable to sequent systems. I propose a way of formulating a principle of harmony within the sequent setting and demonstrate that, contrary to the received view, the standard classical sequent calculus *is* harmonious. Having established this, I turn to the central question of whether the result constitutes a vindication of classical logic. Chapter 11 addresses a possible worry as to whether sequent calculi pass muster as inferentialist frameworks (and by meaning-theoretic lights). I examine a number of possible objections against the notion that sequent calculi are suitable codifications of our inferential practice capable of representing our meaning-conferring principles of inference. Contrary to first impressions, these arguments do not carry much weight and can be parried by the realist. The upshot of our discussion is that *au fond* the issue turns on the legitimacy of multiple-conclusion systems generally, rather than of sequent systems in particular. The question of the legitimacy of multiple conclusions is thus the concern of chapter 12. I begin by addressing the question of whether multiple-conclusion systems have an in-built bias towards non-constructivity. An argument by Tennant to this effect is considered, but exposed as viciously circular. However, I show how, by expanding on an idea of Dummett's, Tennant's point can be transformed into a powerful argument against multiple-conclusion systems: such systems are intelligible to us only if we have a prior grasp of the meaning of disjunction. Therefore, they prove to be unsuitable for inferentialist purposes. A possible rejoinder for the classicist is then discussed. The idea is to devise an alternative interpretation of multiple-conclusion systems by invoking a speech act of denial alongside that of assertion. I argue that this attempt fails even if we countenance a notion of denial because it too is incompatible with inferentialist strictures. Finally, chapter 13 addresses the question of separability. We begin by considering Peter Milne's explanation of the 'magical fact', i.e. the rather surprising way in which sequent calculi encapsulate the transition from constructive to classical reasoning in the move from single to multiple succedents. Having thus located the exact source of non-constructivity in multiple-conclusion systems, we note that the real issue separating the classicist and the anti-realist is the question of whether or not to espouse the principle of separability. I argue that, relative to the initial meaning-theoretic assumptions laid out in part one, the principle of separability should be endorsed. Granted these assumptions, the realist's attempts of reconciling inferentialism with multiple-conclusion systems—or indeed with any system that does not satisfy separability must therefore ultimately fail. We thus arrive at the following conditional conclusion: *if* the proof-theoretic argument goes through in the case of natural deduction, it goes through *tout court*.

Chapter 10

Harmony in the sequent setting

I noted that the sequent calculus affords an elegant systematization of the principles of classical logic. That is to say, it is free of the displeasing asymmetries that plague its natural deduction analogue. For the anti-realist, of course, these flaws, far from relying on arbitrary aesthetic considerations, are indicative of substantial meaning-theoretic shortcomings. The realist, however, contends that classical logic does not fare well on natural deduction presentations because these systems have a certain built-in bias towards constructivist thought; that such N-systems have been specifically tailored to privilege constructive modes of reasoning. As William Kneale puts it,

> Gentzen's success in making intuitionist logic look like something simpler and more basic than classical logic depends, as he himself admits, on the special forms of the rules he uses, and in particular on the requirement that they should all be rules of inference (Kneale 1956, p. 253).

Unfortunately Kneale does not explain himself any further on this issue. Be that as it may, the realist has good reasons to put his money on the sequent calculus. Yet the realist who, *ex hypothesi*, abides by inferentialist strictures and is committed to the principle of harmony, will acknowledge that the 'naturalness' with which the standard sequent calculus represents classical logic is significant only insofar as it is compatible with meaning-theoretic constraints. That is, the classical sequent calculus, elegant though it may be, is a real option only if it can be shown to be harmonious.

Are there good grounds for believing that the neatness of the sequent calculus rendering of classical principles in general, and of the rules for negation in particular, is an indication that these principles are in harmony? And if it really turns out that the classical sequent calculus but not its natural deduction counterpart is harmonious, what would this tell us about the notion of harmony? Would harmony thereby be shown to be a relative notion—a notion relative to a particular proof system? In dealing with the realist we are dealing with someone who wishes to prove the proof-theoretic argument to be invalid (rather than unsound), and this the realist can achieve by showing that the classical meanings of the logical constants can receive a harmonious inferentialist treatment. In particular, then, a realist of this stripe will accept the principle of harmony as a *desideratum* for the semantic well-functioning of the logical fragment of language. But if harmony is a legitimate meaning-theoretic principle, it cannot be relative to a particular formalism. Both natural deduction and the sequent calculus purport to represent accurately the meaning-constitutive inferential relations in which the logical expressions partake. It cannot be the case that our principles governing the logical constants are harmonious according to one representation and not the other. Therefore, if such a situation should arise—and we will see in the following that it does—then one camp must either demonstrate the inadequacy of the proof-theoretic framework favoured by the other or will have to show that the other camp has misconstrued the notion of harmony.

Let us return to the question we began with at the beginning of the previous paragraph, 'Are the classical sequent rules harmonious?' Before we can hope to answer this question, we must ask how the question of harmony can even be raised within the framework of the sequent calculus. Recall that on our intuitive conception of harmony, we require that the grounds for introducing a constant be matched by the consequences of having introduced it. In the natural deduction setting this amounted to the demand that the introduction and elimination rules associated with each constant be in equilibrium—that the elimination rules allow us to infer *no more* and *no less* than what we could infer directly from the premises of the introduction rules.

The question now is whether it is possible to give the principle of harmony an interpretation within the sequent calculs. The problem is that harmony appears to be intimately linked to the architecture of natural deduction systems. How then can the realist appropriate the principle of harmony for the sequent setting? A promising point of departure, it turns out, is the notion of intrinsic harmony. Recall that we found intrinsic harmony (i.e. eliminability of local peaks) to be an essential part of harmony—a necessary condition, in fact. Our first task should be to uncover a corresponding procedure for preventing P-strong disharmony in the sequent setting.

A look at the evolution of Gentzen's project provides a hint here. The chief objective of his famous 'Investigations into logical deduction' was to prove the *Hauptsatz* within the newly introduced system of natural deduction. As it happens, Gentzen failed to prove the *Hauptsatz* for NK (although he did succeed in the case of the intuitionistic fragment)[1]. His failure led him to conclude that natural deduction systems 'proved unsuitable' for the purpose of establishing the *Hauptsatz* and that 'in order to prove the *Hauptsatz* in a convenient form' he would have to 'provide a logical calculus especially suited to the purpose' (Gentzen 1969a, p. 69)—the sequent calculus.[2] It is noteworthy that Gentzen uses the word *Hauptsatz* to refer to both results. In his mind the normalization theorem in natural deduction and the cut-elimination theorem in sequent calculi, as we have since become accustomed to distinguish them, constitute, in an important sense, one and the same result. The two theorems express the same underlying fundamental insight, namely that any proof admits of conversion into a certain 'direct' normal form. By showing that any proof can be converted into such a normal form, Gentzen in fact demonstrates the dispensability of CUT and hence provides a justification for the method of proofs by lemmas. Both in mathematics and in ordinary reasoning we tend to 'establish a result in stages, setting forth a series of lemmas, which are then composed to yield a further result' (Read 2000, p. 124). The advantage is a gain in clarity, understanding and surveyability. But this method is justified 'by the fact that the proofs of the lemmas can themselves be composed into single and direct proofs of the main theorem' (ibid.). The existence of a normal form guarantees that this process of disassembling proofs will again lead to a proof (albeit in general to a much much longer one) of the same result.

How does this idea find expression in the normalization theorem? The connection consists in the fact that when the method of proceeding by lemmas leads us to compose complex proofs, maxima will normally occur in precisely those places where lemmas are grafted together. Where one lemma joins another within the greater

[1]See von Plato (2008).

[2]Prawitz (1965) later showed how the *Hauptsatz* can be established directly in NK.

structure of a proof, the preliminary conclusion of the first lemma will generally be the conclusion of an introduction rule and simultaneously, in its function as a premise for the following lemma, the major premise of an elimination rule. Since any kind of 'indirectness' generates maxima (or plateaux), once we have shown that all such local peaks (as well as plateaux) can be eliminated, the proof we end up with will be a proof of the same conclusion but of the desired direct form. The normalization theorem thus furnishes a guarantee and indeed a method for dismantling proofs by lemmas producing proofs in which the same overall result can be established directly, from first principles, as it were.

Now, although we have emphasized the importance of distinguishing between the local procedure of levelling peaks and the global property of normalizability (which, in general, requires additional reduction procedures), it is also clear that the eliminability of maxima is the heart of (at least the standard) normalization theorems. Hence, it is precisely the notion of intrinsic harmony, as formally represented by the inversion principle, that ensures that proofs by lemmas can be reconstituted so as to form direct proofs of the same conclusion. But if intrinsic harmony is an essential component of the content of the *Hauptsatz*, and if the cut-elimination theorem really does, as Gentzen suggests, express the same content in a different guise, might we not expect the proof of the cut-elimination theorem also to incorporate a formal correlate to intrinsic harmony? If this is correct, the realist should be able to extract from the (proof-theoretic) proof of the cut-elimination theorem a procedure analogous to the levelling procedures involved in the inversion principle. Indeed if the constraints of harmony are to be more than a peculiarity of the natural deduction presentation, we should expect to find such an analogue in any proof-theoretic framework that affords a suitable representation of our inferential practice. Therefore, if the standard sequent calculus constitutes such a suitable representation, it should be possible to carve out a notion of harmony within the sequent setting.

10.1 Harmony and cut-elimination

We do not have to look far to find an analogue to the inversion principle. The crucial notion we are after is—unsurprisingly—the rule of *CUT* and the possibility of a procedure for its elimination. To see why, and to illustrate the considerations from the previous section further, consider the following two natural deduction proofs:

$$\begin{array}{c}\Gamma\\ \Pi\\ A\end{array}$$

and

$$\begin{array}{c}\Gamma', A\\ \Pi'\\ B\end{array}$$

As we said, our deductive practice, mathematical and other, depends to a great extent on the possibility of 'cumulative deductive progress' (Tennant 2005a, p. 8), i.e. the possibility of joining together the two results, a proof of A from Γ and a proof of B from Γ' and A, in order to obtain the overall result B from hypotheses $\Gamma \cup \Gamma'$. Schematically we may represent the case in question as follows in the natural deduction setting:

$$\begin{array}{c}\Gamma\\ \Pi\\ A\\ \underbrace{\Gamma', A}\\ \Pi'\\ B\end{array}$$

How does the same proof present itself in the sequent setting? Here too we have proofs of the two lemmas unified into a single proof.

$$\frac{\begin{array}{c}\Pi\\ \Gamma : A\end{array} \qquad \begin{array}{c}\Pi'\\ \Gamma', A : B\end{array}}{\Gamma, \Gamma' : B}$$

The unifying bit, encapsulated in the inferential step, is precisely an application of the *CUT* rule. This makes it clearer why Gentzen viewed normalization and cut-elimination as two manifestations of the same result. Where there are maxima in natural deduction, there are applications of *CUT* in the sequent calculus.[3] Whence the intimate connection between eliminating maxima and eliminating *CUT*. In the case of the normalization theorem in the natural deduction context, the effect of the global structural rule has to be simulated locally. It is for this reason that the normalization theorem relies on *operator-based* reduction procedures. It has

[3] Strictly speaking this is inaccurate. In reality only *some* instances of *CUT*, so-called *principal* occurrences, correspond to maxima. But this subtlety can safely be ignored for present purposes.

to examine all the different types of grafting points between lemmas depending on which rules of inference (and hence which logical constant) mediate them. By contrast, in the sequent setting, the assumption of transitivity that underlies the method of concatenating proofs (and hence of employing lemmas) is explicitly stated in the form of an inference rule. That is, the grafting is recognized as an inferential step in its own right. Accordingly, all grafting points can be dealt with using the rule of *CUT*. Hence, all we need to do in order to demonstrate that the method of lemmas is dispensable is to show that instances of *CUT* can be avoided.

But how, we may ask, does this relate to our intuitive notion of harmony? Let us start with the obvious. Within both natural deduction systems and sequent calculi, inference rules come in pairs. In the case of natural deduction, as we have seen, intrinsic harmony demands that the rules be balanced in such a way that the consequences warranted by the elimination rule do not outstrip the grounds for asserting its major premise via the corresponding introduction rule. How can this idea be transposed from the natural deduction context with its pairs of introduction and elimination rules to the sequent-calculus setting with it pairs of right-hand side (RHS) and left-hand side (LHS) introduction rules?

The correspondence of natural deduction introduction rules and RHS introduction rules in the sequent calculus is straightforward.[4] More interesting is the connection between elimination rules and LHS introduction rules. It is not obvious, at least on the face of it, that the two types of rules fulfil functionally equivalent roles in their respective systems. The elimination rule associated with a particular logical operator states in schematic form the inferences one is justified in making once one has asserted a sentence in which the operator in question has a dominant occurrence. In other words, an elimination rule specifies which inferences *from* the statement in question—taken as a major premise—are admissible (as opposed to introduction rules, which map out the (canonical) paths that lead *to* the statement in question). In what sense can LHS introduction rules be said to accomplish this task?

To see that LHS introduction rules and elimination rules amount, as it were, 'to the same thing', it is helpful to appeal to so-called 'translations' from natural deduction to sequent calculi. As we will see in greater detail in section 13.1, 'translations'

[4] As we will see below in our discussion of the 'seemingly magical fact' (see section 13.1), things are not as straightforward as they first appear in the case of sequent calculi which, like the standard classical system, allow several formulas in the succedent.

are really inductive clauses in proofs (by induction on the length of derivations) of the coextensiveness of the two systems concerned. Here, however, we are more concerned with understanding how both elimination rules and LHS introduction rules are specifications of the same aspect of an expression's meaning: its inferential consequences.

To this end consider the simplest case of the ∧-introduction rule on the left in a standard intuitionistic sequent calculus.

$$\dfrac{\begin{array}{c}\Pi\\ \Gamma, A : C\end{array}}{\Gamma, A \wedge B : C}\ {\scriptstyle \wedge\text{-LI}}$$

What the rule tells us is that if C is provable from the set of hypotheses Γ and A, then C will *a fortiori* be provable from the hypotheses Γ together with the proposition $A \wedge B$. The sequent rule can easily be justified from the perspective of natural deduction. For suppose we have a corresponding natural deduction proof Π_0 of C from Γ and A

$$\begin{array}{c}\Gamma, A\\ \Pi_0\\ C\end{array}$$

Armed with our ∧-elimination rule, we can convert the existing proof into a proof of C from Γ and $A \wedge B$ simply by appending an instance of an application of ∧-elimination, taking our new hypothesis $A \wedge B$ as its major premise.

$$\begin{array}{c}{\scriptstyle \wedge\text{-E}}\ \dfrac{\Gamma, A \wedge B}{A}\\ \Pi_0\\ C\end{array}$$

Viewed in this way, the LHS introduction rule for ∧ corresponds to an upward extension of the corresponding natural deduction proof beginning with an instance of ∧-elimination. Similar translations can be given for the remaining connectives, showing that elimination rules in natural deduction and LHS introduction rules in sequent systems really do play essentially the same role.

How does this bring us any closer to a notion of harmony for the sequent calculus? Well, consider an arbitrary binary connective $. Intrinsic harmony obtains between $-introductions and $-eliminations when no more can be inferred via $-elimination rules than could have been inferred on the basis of the grounds for applying $-introduction. Consider again the simplest case of an $\wedge$ elimination leading to the overall conclusion C,

$$\cfrac{A \wedge B}{\begin{array}{c} A \\ \Pi_2 \\ C \end{array}} \;\scriptstyle{\wedge\text{-E}}$$

If $\wedge$ is intrinsically harmonious, then it must be possible to infer C already from the premises that would warrant the assertion of $A \wedge B$. Filling in these details we obtain

$$\begin{array}{c} \begin{array}{cc} \Gamma_0 & \Gamma_1 \\ \Pi_0 & \Pi_1 \\ A & B \end{array} \\ \hline A \wedge B \quad (\wedge\text{-I}) \\ \hline A \quad (\wedge\text{-E}) \\ \Pi_2 \\ C \end{array}$$

Appending the proof Π_0 of A from Γ_0 with the proof Π_2 of C from A, we obtain the desired direct proof of C from A and Γ:

$$\begin{array}{c} \Gamma_0 \\ \Pi_0 \\ A \\ \Pi_2 \\ C \end{array}$$

(Obviously the same holds, *mutatis mutandis*, to arrive at B.)

Let us now consider how this carries over to the sequent setting. The sequent correlate of a $\wedge$-maximum has the following form:

$$\cfrac{\cfrac{\Gamma_0 : A \qquad \Gamma_1 : B}{\Gamma_0, \Gamma_1 : A \wedge B}\,\wedge\text{-RI} \qquad \cfrac{\Gamma_2, A : C}{\Gamma_2, A \wedge B : C}\,\wedge\text{-LI}}{\Gamma_0, \Gamma_1, \Gamma_2 : C}\,\mathit{CUT}$$

(Again, the same obviously holds if the premise for the LHS introduction rule is B.) Now our task is to show that the same result, a proof of C from the union of the hypotheses Γ_0, Γ_1 and Γ_2, can be obtained directly, without the need to first introduce $\wedge$ on the left and on the right. This would show that anything that follows from $A \wedge B$ (along with other hypotheses) is already deducible on the basis of $\Gamma_0 : A$ or $\Gamma_1 : B$. Since any C that follows from $A \wedge B$ follows from either A or B, we can simply apply CUT to the premises. That is, given

$$\wedge\text{-LI}\ \frac{\Gamma, A : C}{\Gamma, A \wedge B : C}$$

it must be possible to arrive at C solely on the basis of the grounds for inferring $A \wedge B$ (on the right), i.e. $\Delta_0 : A$ and $\Delta_1 : B$. That this is indeed the case can easily be shown with the help of the CUT rule:

$$CUT\ \frac{\Delta_0 : A \qquad \Gamma, A : C}{\Delta_0, \Gamma : C}$$

(The case where $A \wedge B$ is introduced on the left from B is similar.) This shows that if C follows from $A \wedge B$ we need no more than is required for an inference to $A \wedge B$ in order to arrive at C.

Consider now the case of disjunction. The LHS introduction rule mirrors exactly the natural deduction elimination rule:

$$\vee\text{-LI}\ \frac{\Gamma_0, A : C \qquad \Gamma_1, B : C}{\Gamma_0, \Gamma_1, A \vee B : C}$$

What we need to show is that the premises of the LHS $\vee$-introduction rule are matched by the corresponding RHS introduction rule. This is clearly the case since $A \vee B$ may be introduced either on the grounds that $\Delta_0 : A$ or on the grounds that $\Delta_1 : B$. Assuming—without loss of generality—the former, we obtain the desired result:

$$CUT\ \frac{\Delta_0 : A \qquad \Gamma_0, A : C}{\Delta_0, \Gamma_0 : C}$$

Turning now to $\supset$, we find ourselves with the following LHS introduction rule:

$$\supset\text{-LI}\ \frac{\Gamma_0 : A \qquad \Gamma_1, B : C}{\Gamma_0, \Gamma_1, A \supset B : C}$$

The RHS introduction rule has only the premise $\Delta, A : B$. Since both premises of the LHS introduction rule are relevant, we need to apply the CUT rule twice:

$$CUT\ \frac{CUT\ \dfrac{\Gamma_0 : A \qquad \Delta, A : B}{\Gamma_0, \Delta : B} \qquad \Gamma_1, B : C}{\Gamma_0, \Gamma_1, \Delta : C}$$

Finally, we deal with negation. Given the LHS introduction rule for negation

$$\neg\text{-LI}\ \frac{\Delta : A}{\Delta, \neg A :}$$

we arrive at the following reduction procedure:

$$\textit{CUT}\ \frac{\Delta : A \qquad \Gamma, A :}{\Delta, \Gamma :}$$

Applying the same method to the quantifiers, we obtain the following. The LHS introduction rule for the universal quantifier is (with the customary restrictions):

$$\forall\text{-LI}\ \frac{\Gamma, A[t/x] : C}{\Gamma, \forall x A(x) : C}$$

This harmonizes with the RHS rule in the following predictable way:

$$\textit{CUT}\ \frac{\Delta : A[a/x][t/a] \qquad \Gamma, A[t/x] : C}{\Delta, \Gamma : C}$$

Finally, let us consider the case of the existential quantifier:

$$\exists\text{-LI}\ \frac{\Gamma, A[a/x] : C}{\Gamma, \exists x A(x) : C}$$

This case can be dealt with as follows, completing our demonstration:

$$\textit{CUT}\ \frac{\Delta : A[t/x] \qquad \Gamma, A[a/x][t/a] : C}{\Delta, \Gamma : C}$$

We have thus shown how the notion of intrinsic harmony can be adapted to sequent calculi. However, so far—strictly for readability's sake, as we will now show—we have been considering only the intuitionistic sequent calculus. To demonstrate that our result carries over to the classical case—which is, after all, what we are interested in here—we need to show that the rules still display the desired harmony when we allow for sequents with multiple conclusions. But this is easily achieved at the cost of a slight loss of perspicuity. Let us illustrate this in the case of the conditional. The LHS introduction rule in the multiple conclusion calculus is now:

$$\supset\text{-LI}\ \frac{\Gamma_0 : A, \Theta_0 \qquad \Gamma_1, B : \Theta_1}{\Gamma_0, \Gamma_1, A \supset B : \Theta_0, \Theta_1}$$

Making the necessary emendations, we obtain the corresponding reduction procedure:

$$\textit{CUT}\ \frac{\textit{CUT}\ \dfrac{\Gamma_0 : A, \Theta_0 \qquad \Delta, A : B, \Sigma}{\Gamma_0, \Delta : B, \Theta_0, \Sigma} \qquad \Gamma_1, B : \Theta_1}{\Gamma_0, \Gamma_1, \Delta : \Sigma, \Theta_0, \Theta_1}$$

This completes our demonstration.

10.2 Intrinsic harmony in the sequent setting

Just as in the natural deduction setting, intrinsic harmony prevents cases of P-strong disharmony in the sequent setting. It ensures that *no more* can follow from taking a statement introduced on the left (along with other hypotheses) than already follows from the grounds for introducing the statement in question on the right.

To illustrate this, let us examine the sequent calculus analogue of Prior's connective `tonk`.[5] `Tonk` is introduced by one of the disjunction introduction rules

$$\text{tonk-RI}\ \frac{\Gamma : A}{\Gamma : A \text{ tonk } B}$$

Predictably, instead of the $\wedge$-elimination rule we model the `tonk`-elimination rules on the corresponding LHS $\wedge$-introduction rule

$$\text{tonk-LI}\ \frac{\Gamma, B : C}{\Gamma, A \text{ tonk } B : C}$$

Supposing the availability of the *CUT* rule, we now arrive at the familiar absurd consequences:

$$CUT\ \frac{\text{tonk-RI}\ \dfrac{A : A}{A : A \text{ tonk } B} \qquad \text{tonk-LI}\ \dfrac{B : B}{A \text{ tonk } B : B}}{A : B}$$

Clearly such a connective is ruled out by the principle of harmony as formulated above: B, the consequent of the premise of the LHS introduction rule, would have to be deducible from the grounds for inferring $\ulcorner A$ `tonk` $B\urcorner$ alone, namely A, which plainly is not the case (at least if we assume the system's antecedent consistency). The necessary appeal to *CUT* is no concession, since the natural deduction analogue depends no less on the transitivity of the deductive consequence relation.

We here see an important sense in which the focus on *CUT* is no surprise: the reduction procedures constitutive of the sequent calculus variant of intrinsic harmony turn out to be closely related to the inductive step in the proof of the *CUT*-elimination theorem. Each such procedure shows how an instance of *CUT* applied to a formula with a certain complexity (the *CUT* formula) can be reduced to an instance in which the *CUT* formula is of lower complexity; to each logical operator corresponds a particular reduction procedure. The common designation 'reduction

[5]See our discussion of `tonk` in section 3.2.

procedures' for both the procedures employed in the proof of the normalization theorem and in the proof of the cut-elimination theorem is thus no coincidence, since both procedures play functionally equivalent roles in the proof of the respective versions of the *Hauptsatz*. We have thus been able to verify our conjecture that a formal counterpart of the inversion principle can be extracted from the cut-elimination theorem within the sequent calculus.

Before moving on, let us briefly note that classical logic in its sequent calculus formulation also satisfies the stronger principle of total harmony. As we have seen in discussing Dummett's claim in section 4.6, classical natural deduction systems fail to display total harmony. Strengthening the NK fragment containing only $\supset$ by adding the classical rules for $\neg$ gives rise to a non-conservative extension; in the extended system, we can, for instance, in the extended system, prove Peirce's law $((A \supset B) \supset A) \supset A$, a formula that, despite making use only of the conditional, could not be proved in the restricted system. The same is true for the fragment $\{\vee, \supset\}$, which, when enriched by the same rules for negation, gives rise to new theorems like $(A \supset (B \vee C)) \supset ((A \supset B) \vee (A \supset C))$. This situation is rectified in the classical sequent calculus, where the addition of '$\neg$' to the language along with corresponding rules yields conservative extensions. To illustrate, Peirce's law can now be proven without having to invoke the rules for negation as follows:[6]

$$\dfrac{\dfrac{\dfrac{\dfrac{A : A}{A : B, A}\,\text{Right weakening}}{: A \supset B, A}\,\supset\text{-RI} \quad A : A}{(A \supset B) \supset A : A}\,\supset\text{-LI}}{: ((A \supset B) \supset A) \supset A}\,\supset\text{-RI}$$

Therefore, the classical sequent calculus meets the more stringent criterion of total harmony.

However, the attainment of intrinsic harmony in the sequent setting, as in the case of natural deduction setting, does not guarantee that the meanings of the logical constants have been fully exploited. That is, the presence of intrinsic harmony guarantees that P-strong disharmony is averted, but not that P-weak disharmony is averted. Again, the example of the quantum-logical disjunction operator serves to illustrate this point. The restricted LHS introduction rule distinguishing the

[6]It should be noted, however, that this result depends crucially on the structural rule of weakening.

quantum-logical disjunction operator from the standard disjunction operator takes the following form in the sequent setting:

$$\frac{A : C \qquad B : C}{A \vee B : C}$$

It is not hard to see that the sequent-style levelling procedure goes through for this case as well. How, then, can we formulate a constraint to ensure that we pick the strongest possible LHS elimination rule to match the RHS introduction rule? Well, as before, all we need to ensure is that the rules impose no restrictions on the contexts except for those necessary to ensure that the reduction procedures presented above can be carried out.

We have thus established that the sequent calculus provides far more than merely a proof-theoretic framework in which classical logic can be given a somewhat more 'elegant' representation. We have shown how the principle of harmony can be brought to bear on it and have established that the classical system is indeed harmonious. Therefore, insofar as the classical sequent calculus can be made out to be an inferentialistically legitimate proof-theoretic framework, the classical logician has resisted the anti-realist's challenge. The realist's victory hinges, however, on the question of whether the standard classical sequent calculus is legitimate by inferentialist standards. It is this question that occupies us in the next chapter.

Chapter 11

On sequent calculi

The classical logician's joy about the encouraging results of the previous section will be short-lived if there turn out to be independent grounds for rejecting the sequent calculus format. One indeed frequently encounters handwave-y claims to this effect. In the present section it will be our aim to spell out and to examine the grounds one might have for thinking that sequent calculi fail to provide an adequate framework for inferentialism.

11.1 The objections, in broad strokes

The accusation we will be examining runs as follows: the transitions that make up sequent proofs cannot really be thought of as inferences at all, nor are the proofs themselves readily interpretable as arguments. The underlying idea is simple: the expressions that constitute the premises and the conclusion of a proof in the sequent calculus are, as the name suggests, sequents rather than sentences. And it is sentences that we employ in ordinary practice. Natural deduction derivations, by way of contrast, are most naturally interpreted as arguments; the formulas occurring within them we interpret as sentences, and the transitions between formulas as inferential steps. Sequents, on the other hand, do not appear to be amenable to such an interpretation. Unlike statements made by means of sentences, sequents are not susceptible to being true or false: a sequent is valid or invalid (it being understood that invalid sequents can have no place in anything that merits the title 'proof'). Accordingly, correct inference rules in the framework of the sequent calculus, contrary to their natural deduction analogues, are more aptly described as 'validity-preserving' than truth-preserving. What is being claimed when a rule of inference in the sequent calculus is said to be correct is that any application of such a rule to one or more valid sequents of the form $\Gamma : \Delta$ (where Γ and Δ are sets of sentences) will again result in a valid sequent $\Gamma' : \Delta'$. An inference rule in the sequent calculus thus amounts to an instruction as to how a number of valid arguments represented by corresponding sequents may be manipulated so as to arrive at another valid argument represented by the conclusion of the inference rule.

According to this line of thought it would therefore seem that a proof conducted in a sequent system may not properly be regarded as an argument at all. Mathematical proofs seem to provide a decisive example. Mathematical proofs are the result of stringing together sentences in accordance with admissible principles of inference in such a way as to form a valid argument with the theorem to be proven as its conclusion. Crucially, in writing proofs, mathematicians reason *about sentences* (or the statements they express). Likewise, in our ordinary deductive reasoning, we wonder what follows from a number of statements that we take to be true. Rare, it seems, are the occasions when we inquire whether altering the structure of an argument in a certain way will again result in a valid argument. And even when this is indeed what we are doing, we consider ourselves to be engaged in a rather different kind of activity: we consciously adopt a meta-theoretical stance and reason about arguments rather than reasoning, as it were, from 'within' an argument. Yet this is precisely what proofs in the sequent calculus do: they 'talk about' arguments—in the case of multiple-conclusion systems, highly stylized (at best) arguments—rather than providing the medium in which to construct them. Sequent systems thus play the role of a meta-calculus for natural deduction. They accomplish this by incorporating into the object language the metalinguistic relation $\vdash$ (represented in our notation by ':' in the object language), which marks the presence of a valid argument from the premises on the left of $\vdash$ to the conclusion on its right. *Stricto sensu*, a sequent proof must be viewed as a proof of the existence of a 'proper' natural deduction derivation. It is nonetheless a constructive proof, since it generally allows us to recover (though not uniquely) a corresponding natural deduction proof. But if the sequent calculus plays the role of a meta-calculus, it must be answerable to the natural deduction system on which it is parasitic. And if this is true, then the realist, in privileging the sequent calculus over the more fundamental natural deduction system, is in a sense sawing off the branch on which he is perched.

Moreover, it may seem doubtful, in light of the present discussion, whether sequent inference rules can live up to inferentialist standards, given that they are transitions between sequents rather than propositions. Can such rules nonetheless adequately represent the inferential links that are constitutive of the meaning of the constant in question?

We have arrived here at a potentially more pressing objection still. Can the sequent calculus provide an adequate framework for inferentialism? The thought here is that if transitions between sequents cannot properly be interpreted as inferences, and if the meanings of the logical constants are to be given in terms of the licit deductive inferences involving them, then the sequent calculus is unable to make these semantically relevant inferential relations explicit. The opponent of the sequent calculus concludes that this system fails to qualify as a proof system suitable for inferentialist use.

Ian Rumfitt's argument against multiple-conclusion systems in his article '"Yes" and "No"' (Rumfitt 2000) can be seen as an extension of this train of thought. Sequents, as we have seen, are partial representations of the metalinguistic consequence relation. But, as Quine has tirelessly reminded us, the *relata* of the relation of entailment are not sentences or statements but names thereof (see e.g. Quine 1940, p. 28).[1] It follows that the sentences that figure in the sets that constitute the antecedents and succedents of sequents are merely mentioned, not used. For example the sequent '{'If the painting is a Vermeer, it cannot date from 1682', 'The painting is dated 1682'} entails {'The painting is not a Vermeer'}' relates sets of *names of sentences*. But as inferentialists we are 'exploring the idea that a connective's sense consists in the way in which it is correctly *used* in deductions' (Rumfitt 2000, p. 795). Therefore, Rumfitt concludes, the sequent calculus proves to be unsuitable for the inferentialist enterprise.

To summarize, we have identified two objections:

1. The sequent calculus mentions rather than uses sentences occurring in the arguments expressed by it, which makes it unsuitable for inferentialist purposes.

2. The sequent calculus is a meta-calculus for natural deduction and is therefore parasitic on it. We must therefore accord primacy to natural deduction systems.

[1] This holds as much for the proof-theoretic relation of deducibility $\vdash$ as it does for the semantic notion of consequence $\models$.

11.2 Replies

Must we therefore conclude that—the classicist's compelling case for harmony in the previous section notwithstanding—the sequent calculus fails to provide a refuge for the realist? Or can we mount a defence of the sequent calculus?

Let us begin with Rumfitt's objection. How forceful is the application of the Quine point here? Not very, I would argue. To be sure, inferentialism about the meanings of the logical constants is a form of use-theory of meaning (the relevant feature of use being that of the expression's employment in deductive inference).[2] Thus a proof-theoretic framework fit to serve the inferentialist's purposes should yield an adequate representation of the use, understood in terms of V- and P-conditions, that the logical expressions are put to in deductive inference. But must we conclude from this that the only acceptable framework is one in which sentences are only used, rather than mentioned? It is certainly at least conceivable that the use we make of our logical expressions could receive an adequate representation in a system in which the sentences containing these expressions are mentioned rather than used. We could understand such a system as a set of meta-theoretic instructions canonizing permissible inferences: from 'If the painting is a Vermeer, it cannot date from 1682' and 'The painting is dated 1682', infer 'The painting is not a Vermeer'. To be sure, *in following* the rules of the inferential game we will *use* rather than *mention* the sentences that are merely mentioned in the statement of the rules. So, from our acceptance of the premise that the meanings of the logical expressions are determined by the use they are put to in our deductive practices, together with the assumption that these meaning-conferring inferential links are capable of receiving a systematic representation, it does not follow that this systematic representation must be given at the level of the object-language. Indeed, not only does inferentialism not rule out such a meta-theoretical theory of meaning; some appeal to meta-theory seems inescapable in order even to render intelligible proofs containing only vocabulary of the expressively impoverished object-language. As Timothy Smiley puts it,

> Although logic is said to be the science of argument, the object-languages of modern logic are remarkable for containing no means for actually formulating an argument. 'Therefore' is not among the symbols of any

[2] Note that despite the slight equivocation on 'use' as it occurs in 'use-theory of meaning' and 'use vs. mention', Rumfitt's objection still stands.

> calculus I know of, and a bare sequence of sentences cannot be a deduction. In truth object-languages are inanimate objects. What is needed to animate them is provided in natural deduction systems by annotations like 'premise', 'assumption' or 'from lines 1 and 2', serving as signs of assertion, supposition, and inference respectively (Smiley 1996, p. 7).

Clearly, if a deductive system is to do meaning-theoretic duty, it has to be an interpreted system in Smiley's sense. And this in turn requires meta-theoretical commentary. The clear-cut object-language/metalanguage dichotomy on which Rumfitt's objection is premised thus begins to crumble when viewed from up close. As will emerge from our discussion below, natural deduction, insofar as it is to do inferentialist service, supports such a rigid separation between the two levels no more than the sequent calculus does; both require meta-linguistic commentary.

So much for Rumfitt's objection; let us now turn to the second type of objection. The thought is, let us recall, that the sequent calculus is a meta-calculus for natural deduction systems and that the former must therefore be answerable to the latter. The result would be the constitutive primacy of natural deduction systems over sequent calculi. We had seen that the argument underlying this objection is built on the broader charge that proofs carried out in sequent calculi do not really qualify as proper arguments. Is there a way the advocate of the sequent calculus can defend himself against *this* charge and so undercut the objection? The crux of the argument, recall, was that sequent systems deal in validity-preserving transitions between arguments rather than the truth-preserving inferences between statements, and that it is the latter rather than the former that we can engage in in ordinary practice.

The argument has some initial appeal when we look at simple cases, such as inferring $A \vee B$ from A in the course of an argument. Inferences of this kind strike us, in an intuitive sense, as *local*: the inferrer is required to appeal only to the premise (or, more precisely, to the logical form of the statement that constitutes the premise) and to no other statements or sub-arguments that may occur within the argument to justify his inference. On the face of it, matters present themselves quite differently in the case of the analogous rule of inference in the sequent calculus: 'there is a valid argument from Γ to A; therefore there is a valid argument from Γ to $A \vee B$' (where Γ is the set of assumptions on which A rests) does not appear to be local in the same way. But is there really a philosophically important difference underpinning this intuitive distinction? Might we not simply be dealing with a difference in notation here? The sequent notation simply represents the logical dependence between the hypotheses (by which we mean to include both assumptions and asserted premises) and the conclusion(s) horizontally with the aid of an explicit marker ':' interpretable as the relation of deductive consequence. But natural deduction systems capture the same dependency relations. The only difference is that in the natural deduction format deductive dependence is represented in a vertical array rather than in a horizontal one: any formula occurring within a deduction depends on the undischarged topmost formulas that lie on the same deductive path as the formula in question. Just as we do in ordinary reasoning, where we keep a mental record (or so one would hope) of the assumptions on which our intermediate conclusions depend, we help ourselves to this representational shorthand in order to avoid an explicit restatement of the assumptions at each and every node in the course of the deduction. Therefore, although natural deduction systems in tree (and linear) format can dispense with an explicit sequent symbol, they nonetheless express the same meta-linguistic relation. Sequent systems *say* it; natural deduction systems *show* it.

Where does this leave us with respect to the further point according to which sequent calculi fail to qualify as an adequate proof-theoretic framework because they deal with sequents, i.e. relational metalinguistic statements whose correctness is assessed in terms of validity, rather than with statements susceptible of being true or false? Has this objection too been neutralized? It may seem so: any occurrence of the inferential move from A to $A \vee B$ in the course of a deduction will result in an extension of the deduction preceding that occurrence such that the conclusion $A \vee B$ of this sub-deduction depends on the same (possibly empty) set of hypotheses as its immediate premise. Viewed in this way, there is, as we have pointed out, no more than a notational difference: once we supply the context, the inference from A to $A \vee B$ amounts to the transition from

$$\begin{array}{c}\Gamma\\ \vdots\\ A\end{array}$$

to

$$\cfrac{\begin{array}{c}\Gamma\\ \vdots\\ A\end{array}}{A \vee B}$$

and this is no more than an alternative presentation of the logical dependency represented in sequent notation as

$$\frac{\Gamma : A}{\Gamma : A \vee B}$$

Does this show that there is no basis to the criticisms against the sequent calculus we have been investigating? Not quite. What we have shown is that when we consider an application of the $\vee$-introduction rule *embedded within a deduction*, the natural deduction incarnation of the rule is no more local than its sequent calculus rival. And this is not surprising, because both modes of presentation express the holding of the *global* meta-linguistic relation $\vdash$ between the initial hypotheses and the conclusion.

But have we thereby demonstrated that inference rules in natural deduction operate on sequents rather than statements, in the same way as they do in sequent systems proper? Here, it seems, there remains a noteworthy difference. But, as I will go on to show, this difference only appears if we pick our examples carefully. Start with the example of the (RHS) $\vee$-introduction rule: to see that natural deduction rules operate on sentences, rather than on sequents, it suffices to note that any instance of an application of the $\vee$-introduction rule itself expresses the fact that $A \vdash A \vee B$. That is, any instance of the rule is a deduction in its own right that can, in virtue of the assumed transitivity of the relation of deductive consequence, be appended to deductions with matching premises or conclusions, i.e. in this case deductions that terminate with the conclusion A or contain $A \vee B$ as an hypothesis. As we have remarked earlier, the assumption of transitivity is built into natural deduction systems.[3] We could thus represent the above natural deduction proof terminating with the inference from A to $A \vee B$ meta-theoretically as follows:

$$\textsc{cut}\ \frac{\Gamma \vdash A \qquad A \vdash A \vee B}{\Gamma \vdash A \vee B}$$

If we now replace '$\vdash$' by its object-language analogue ':', we find ourselves with a fragment of a sequent calculus proof. The transitivity of $\vdash$, tacitly assumed in the natural deduction framework, has been made explicit in the form of an instance of the application of the CUT rule in the sequent setting. In the presence of CUT we could

[3]At least this holds true if we countenance non-normal proofs. This assumption is not shared by everyone. The systems **IR** and **CR**, for instance, presuppose that all proofs are presented in normal form.

thus achieve the same effect by replacing the standard RHS introduction rule for disjunction by the axiom $A : A \vee B$. The fact that the sequent calculus incorporates an explicit rule to mark the transitivity of the relation of logical consequence may again be taken to suggest that the sequent calculus has its place further removed from actual practice in the meta-theory.

11.3 Improper inference rules

But this just brings us back to our earlier point that the meta-theoretical nature of sequent calculi does not deprive them of their inferentialist credentials. The question, rather, is whether the fact that sequent systems deal in sequents rather than in sentences counts decisively against this type of system. So far we have seen that there is nothing about the format of sequents that should tell against such systems from an inferentialist point of view. But what of the practice of applying inference rules to the hypotheses of deductions as well as to conclusions? In answering *this* question it is important to remember that not all rules in natural deduction follow the pattern of the disjunction introduction rule. ∨-introduction is an example of what we have following Prawitz, called an 'proper inference rule' (Prawitz 1965, p. 23), i.e. a rule that marks inferential transitions, all of whose premises are statements and so, as we have just seen, can be incorporated into any suitably structured deduction. In standard natural deduction formulations, as we noted, the ∧-introduction, ∧-elimination, ∨-introduction, ⊃-elimination, ∀-elimination and ∃-introduction rules all fall into this category.

The category of 'improper inference rules' houses the remaining inference rules. Improper inference rules are characterized by the fact that their statement essentially involves the mention of subproofs as well as the possible discharge of hypotheses. It is because such improper inference rules introduce forms of inference that call for subproofs from unasserted assumptions that they are of particular interest in the current context. Consider yet again our ∨-elimination rule:

$$\dfrac{\begin{array}{c}\Gamma_0 \\ \vdots \\ A \vee B\end{array} \qquad \begin{array}{c}\Gamma_1, [A]^i \\ \vdots \\ C\end{array} \qquad \begin{array}{c}\Gamma_2, [B]^i \\ \vdots \\ C\end{array}}{C}\;\vee\text{-E}, i$$

Here the conditions for the correct application of the inference rule do not only take the form of statements of a certain form; some of the premises can also take the form of subproofs. Rules of this form cannot be stated by saying 'From $A_1, \ldots, A_n$ infer B'; rather, we must, to use the same example, resort to locutions like 'Given $A \vee B$, a derivation of C from A and a derivation of C from B, we may infer C from $A \vee B$ alone'. Although the inferential transition being described is itself primitive in the sense that it is irreducible to simpler ones, the grounds for applying the rule in question mention potentially highly complex chains of inference. The rule tells us that whenever we are in a position to assert $A \vee B$ and we can show that $A \vdash C$ and $B \vdash C$, we can infer from this that $A \vee B \vdash C$. As Dummett aptly puts it,

> 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, p. 186).

In admitting improper rules we therefore already essentially commit ourselves to sequents. Moreover, the statement of indirect inference rules crucially depends on our ability to talk *about* deductions; we have to be able to mention them as well as use them.

Smiley, who, as we have seen, appreciates the necessity for natural deduction systems to have some recourse to the metalanguage, nevertheless attempts to maintain a distinction between more or less thoroughgoing appeals to the metalanguage. The first level of meta-linguistic involvement is, according to Smiley, best 'described as supplementing the deficiencies of the object-language'. This is all that is required in order to 'animate' the expressively impoverished object-language. The further step, or second level of meta-linguistic involvement, then draws on resources sufficient to 'talk about the object-language' (Smiley 1996, p. 7) using, object-language propositions and complex expressions composed of them. However, if our previous considerations concerning improper inference rules are correct, Smiley's distinction seems to collapse when applied to natural deduction systems. On the assumption that improper inference rules are essential to their formulation—which they are so long as we wish to dispense with axioms—natural deduction systems partake in the second level of meta-linguistic involvement. If this is the case, natural deduction systems have no edge over sequent systems in this respect. If the fact that a proof system at times mentions rather than uses object-language expressions in the statement of inference rules is problematic for an inferentialist, this will count against natural deduction systems as much as against the sequent calculus.

Above we have argued, against Rumfitt, that there is no reason why an inferentialist account should not operate at the level of the metalanguage. Now we see that Rumfitt was mistaken in holding that standard natural deduction systems can avoid any such meta-linguistic excursions. Not only are meta-linguistic accounts legitimate; they are the only option available (at least if we restrict our attention to standard natural deduction and sequent systems).

Finally, having got this supposed 'difference in kind' between natural deduction and sequent systems out of the way, the claim that natural deduction is somehow the more fundamental of the two systems also becomes groundless: the fact that the sequent calculus is in some sense answerable to natural deduction, of which it would be the meta-calculus, can be based on no more than the contingent historical order in which Gentzen constructed them.

11.4 Getting to the real issue

Should there be any hard-headed advocate of natural deduction systems who remains unpersuaded by the defence of the sequent calculus that we have mounted so far, the realist still has a card up his sleeve. The defender of classical logic can even grant the anti-realist that the sequent calculus is unsuitable for inferentialist purposes and yet reap all the benefits of our above demonstration of harmony for classical logic. How? By simply adopting a multiple-conclusion natural deduction system. Such systems are obtained by allowing inference rules leading to more than one conclusion and by introducing structural rules of weakening and contraction familiar from sequent systems.

To give a taste of such systems without going into unnecessary detail, let us mimic the proof of Peirce's law above:

$$\supset\text{-I, 2}\ \dfrac{\text{contraction}\ \dfrac{\supset\text{-E}\ \dfrac{\supset\text{-I, 1}\ \dfrac{\text{weakening}\ \dfrac{[A]^1}{A, B}}{A, A \supset B} \qquad [(A \supset B) \supset A]^2}{A, A}}{A}}{((A \supset B) \supset A) \supset A}$$

We note that, as in the sequent proof we gave above of Peirce's law, this multiple-conclusion proof requires no more than the rules governing $\supset$ and the two structural rules. Indeed it turns out that this system, like the corresponding sequent system, satisfies the conservativeness constraint. As we will see below (see section 13.1), strengthening NJ by replacing the standard $\supset$-introduction rule by its multiple-conclusion correlate is sufficient to obtain a system adequate for classical logic:

$$\supset\text{-I, 1}\ \dfrac{\begin{matrix}[A]^1 \\ \vdots \\ B, C\end{matrix}}{A \supset B, C}$$

Let us call the system so obtained NKM. We shall be examining the rule that gives rise to NKM and others like it in more detail below. For present purposes all we need to note is that, as Stephen Read has shown, the system's rules of inference are harmoniously balanced (Read 2000, p. 150). Since NKM differs from the common natural deduction systems only in that it accommodates multiple conclusions, even if the critique of sequent calculi considered above were correct (though we have provided grounds for believing that this is not the case), we see that it is not enough, from the anti-realist's point of view, to rule out sequent calculi.

Summarizing, we may say that both objections against sequent calculi considered above fail. Our discussion has made it plain that the reliance on meta-linguistic notions is no more avoidable in natural deduction systems than it is in the sequent calculus. Moreover, even if the critic of sequent systems remains unimpressed by our discussion, the realist still has the option of appealing to NKM. Indeed, NKM makes it clear is that the crux of the issue is not the choice whether sequent calculi or natural deduction systems are to be accorded priority, but rather whether proof systems that admit multiple conclusions are legitimate. Arguably, it is precisely the admission of arguments of this form that lends a system its properly classical character. In the next section we will try to understand the mechanisms underlying multiple-conclusion systems.

Chapter 12

Multiple conclusions

So far we have seen not only how the all-important principle of harmony can be applied to the sequent calculus, but that it is indeed satisfied by the classical sequent system. Moreover, in the previous section we found two familiar objections against sequent calculi wanting. What can the anti-realist say in defence of the proof-theoretic argument in the face of such mounting pressure? Our conclusions at the end of the previous section suggest that the appropriate target of the anti-realist's riposte must be multiple-conclusion systems in general rather than sequent systems. Are there any remaining avenues open to the anti-realist hoping to salvage the proof-theoretic argument? Are there any meaning-theoretic grounds the anti-realist can enlist that count against the acceptability of multiple conclusions; grounds a realist would have to accept? Both Tennant and Dummett have put forward arguments to this effect and we turn to them now.

12.1 Multiple conclusions and constructivity

It is clear that the admission of multiple conclusions somehow introduces non-constructive content in the case of the systems we have considered.[1] But the situation is more complicated than it may at first appear. For one thing, it is not true, at least without qualification, that multiple conclusions are intrinsically non-constructive. Admitting more than one member in the conclusion or succedent is not a sufficient condition for non-constructivity. We thus need not necessarily part company with the anti-realist in admitting multiple conclusions. Indeed, in sequent

[1] We will try to gain a better understanding of the underlying causes in section 13.1.

systems constructivity can be restored by slightly modifying the right-hand side introduction rules for conditional, negation and for the universal quantifier.[2]

To illustrate this for the propositional case, consider a sequent system obtained by taking the standard classical sequent rules for the propositional connectives (along with the usual structural rules) with the exception of the rules governing the material conditional:

$$\supset\text{-RI}\ \frac{\Gamma, A : B, \Delta}{\Gamma : A \supset B, \Delta}$$

$$\supset\text{-LI}\ \frac{\Gamma_0 : A, \Delta_0 \qquad \Gamma_1, B : \Delta_1}{\Gamma_0, \Gamma_1, A \supset B : \Delta_0, \Delta_1}$$

Leaving the LHS rule untouched, we replace the right introduction rule by the following rule:

$$\supset\text{-RI}^*\ \frac{\Gamma, A : B}{\Gamma : A \supset B, \Delta}$$

The difference consists solely in the delayed introduction of the context Δ on the far RHS of the succedent. In the standard classical rule the context is already present in the premise. The system obtained is a multiple-succedent system for propositional intuitionistic logic.[3]

Anti-realists like Tennant, who advocate even more extensive logical reform and favour intuitionistic *relevant* logic, will find fault with the implicit appeal to weakening on the right in introducing the context Δ on the far RHS of the conclusion. It may be objected that the proposed system, by its very constitution, violates relevantist principles, and so precludes the possibility of further revision from the outset. However, the instance of the rule of weakening in the statement of $\supset$-RI* turns out to be inessential. Indeed, we obtain an equivalent system if we replace the standard classical RHS introduction rule for $\supset$ with the usual single-succedent rule.[4] The resulting system recognizably remains a multiple-conclusion system for

[2] For brevity's sake we may make the local stipulation that $\neg A =_{\text{def.}} A \supset \bot$. This allows us to dispense with extra rules for negation. In the light of our earlier remarks concerning $\bot$ as a structural rule, however, it is important to keep in mind that this shortcut is avoidable.

[3] The system can be extended to full intuitionistic first-order logic by amending the RHS introduction rule for the universal quantifier in an analogous way. See Troelsta and Schwichtenberg for details (2000, p. 68).

[4] Accordingly, in the case of the corresponding first-order systems, we replace the RHS rule for the universal quantifier with its standard single-succedent counterpart.

intuitionistic logic. Moreover, it contains no 'built-in' violations of relevantist principles. Allowing sequents with multiple succedents therefore does not *per se* violate constructivist principles, not even in the face of the additional constraint imposed by intuitionistic relevant logicians like Tennant.

Can it now be shown that the admission of multiple conclusions is not a necessary condition for non-constructivity either? At first sight it seems that this too can be done. For example, we can devise a single-succedent sequent system for (propositional) classical logic. This can be achieved by adding a version of the law of the excluded middle (*LEM*) restricted to atomic formulas to the intuitionistic propositional single-succedent calculus. The rule in question is this:[5]

$$\frac{A, \Gamma_0 : B \qquad \neg A, \Gamma_1 : B}{\Gamma_0, \Gamma_1 : B}\ {\scriptstyle LEM}$$

We immediately obtain the law of excluded middle as a theorem:

$$\frac{\dfrac{A : A}{A : A \vee \neg A}\ {\scriptstyle \vee\text{-RI}} \qquad \dfrac{\neg A : \neg A}{\neg A : A \vee \neg A}\ {\scriptstyle \vee\text{-RI}}}{: A \vee \neg A}\ {\scriptstyle LEM}$$

However, there is undeniably something fishy about this version of the law of the excluded middle. Unlike standard operational rules in the sequent calculus, this rule does not introduce a constant, but rather eliminates two formulas, and therefore behaves more like the *CUT* rule. Indeed, *LEM* violates the principle of harmony in that it has no associated rule with respect to which it can be shown to be harmonious. Moreover, the system, when extended to first-order logic, fails to satisfy the subformula property. Such an extension is therefore not a viable option.

But what examples like these do show is that such counterexamples are to be had rather cheaply. This should teach us to refrain from formulating general statements about the admissibility of proof systems based on superficial structural features rather than on an examination of the underlying mechanisms. Are we therefore justified in concluding with Greg Restall that the multiple conclusion system 'makes non-constructive proof *available*, if the vocabulary allows for it, but it doesn't *mandate* it' (Restall 2004, p. 12)?

It would be hasty to go along with Restall's verdict. All we did so far was present a multiple-conclusion *formalism* that outputs all and only the theorems of intuitionistic logic. Before we can free multiple-conclusion systems from the charge

[5] See (Negri and von Plato 2001, p. 114) for details.

of tilting the balance towards classical logic at the very outset, we would have to show that a suitable interpretation can be supplied that links the system in question with our ordinary practice. A proof-theoretic framework can only be said to confer meanings on the logical operators if the rules of inference that constitute it may be understood as pertaining to our ordinary ways of inferring. Without an interpretation that establishes such a connection with our inferential practice we are left with a purely formal construct devoid of any meaning-theoretic interest. An uninterpreted formalism advanced as a putative counterexample against the alleged intrinsic non-constructivity of multiple-conclusion systems has no force.

12.2 Tennant's argument

It is precisely the issue of interpretation that constitutes the point of entry for the anti-realist's criticism. This is what Tennant has to say about multiple-conclusion sequent calculi,

> the classical logician has to treat of sequents of the form $X : Y$ where the succedent Y may in general contain more than one sentence. In general, this smuggles in non-constructivity through the back door. For provable sequents are supposed to represent acceptable arguments. In normal practice, arguments take one from premisses to a single conclusion. There is no acceptable interpretation of the 'validity' of a sequent $X : Q_1, \ldots, Q_n$ in terms of preservation of warrant to assert when X contains only sentences involving no disjunctions. If one is told that $X : Q_1, \ldots, Q_n$ is 'valid' in the extended sense for multiple-conclusion arguments just in case $X : Q_1 \vee \ldots \vee Q_n$ is valid in the usual sense for single-conclusion arguments, the intuitionist can demand to know precisely which disjunct Q_i, then, proves to be derivable from X. No answer to such a question can be provided in general with the multiple-conclusion sequent calculus of the classical logician. It behoves us, then, to stay with a natural deduction system, and to present it in sequent form only if we observe the requirement that sequents should not have multiple conclusions (Tennant 1997, p. 320).

Tennant's objection involves two steps. The first step concerns the interpretation of multiple-succedent sequents: a sequent represents an acceptable argument adequately only if the sequent as a whole is interpreted disjunctively as a single-succedent sequent. But—second step—when interpreted in this way (i.e. disjunctively), it is not in general the case that it can be determined which of the disjuncts in the succedent of the end-sequent holds. In other words, multiple succedent calculi fail to satisfy the disjunction property, which requires that for every proof of a disjunction we can (at least in principle) produce a proof of at least one of the disjuncts. Therefore, multiple-succedent sequents fail to conform to constructivist strictures.

Let us suppose for the moment that Tennant is correct about the interpretability of multiple-conclusion proofs. Nevertheless, as he presents the matter, Tennant's argument strikes one as being vulnerable to the charge of circularity. This is also Read's verdict,

> [Anti-realists] exclude multiple conclusions from consideration because they allow the assertion of disjunctions neither of whose disjuncts is assertible. But that is to beg the question. The question is whether intuitionistic logic is superior proof-theoretically to classical logic. To exclude forms of proof which are intuitionistically unacceptable is to introduce a circle in the reasoning (Read 2000, p. 145).

Read seems to have a point here. Recall that the declared aim of the anti-realist is to translate meaning-theoretic principles into constraints on inference rules, thereby furnishing *independent* meaning-theoretic grounds[6] against classical forms of inference and in favour of intuitionistic ones—grounds other than the standard arguments for the rejection of the classicist's assumption of bivalence. But if this is our aim, we surely cannot appeal to the very principles we set out to justify (e.g. the disjunction property) when explaining to the classical logician why appealing to the sequent calculus cannot be a way of dodging the anti-realist's proof-theoretic argument. This, however, is precisely what Tennant appears to be doing: according to him, the classicist's appeal to the sequent calculus is illegitimate because, in order to obtain classical systems, we have to allow for sequents involving multiple members in their succedent; this the classicist cannot do, he claims, on pain of violating

[6]'Meaning-theoretic' is a more apt designation than Read's 'proof-theoretic', in my opinion. The significance of proof-theoretic criteria within the context of this discussion resides solely in the fact that they reflect certain meaning-theoretically relevant notions; e.g. we are interested in levelling procedures only insofar as they enable us in a particularly perspicuous way to explicate the notion of harmony.

constructivist principles. But these constructivist principles are the very principles the proof-theoretic argument aims to establish in the first place. Tennant's appeal to the disjunction property is thus patently circular.

That being said, must we not admit that Tennant is right about the following: if the dispute between realists and anti-realists is genuinely a dispute about the validity of certain logical principles, then the framework in which it is to be carried out should not favour either of the disputing parties by its intrinsic structural properties? If, for instance, the validity of the law of excluded middle is already guaranteed by the morphology of the system, we shall not expect to arrive at a satisfactory answer to the question, Which is the correct logic? Tennant's point would thus be perfectly justified, *if* it were understood that the conflict is to be construed as one between *equals*. That is, if we begin with the notion that intuitionistic logic (or some subsystem of it) is a contender *inter pares*, we shall have to require of the proof-theoretical framework in which we operate that it maintain strict neutrality with respect to the outcome of the dispute. If *these* were the circumstances, we would indeed have to reject multiple-conclusion systems for the reasons stated by Tennant.

However, this does not seem to capture the dialectical situation we are faced with: the situation is not that of a choice between several logics with the same *prima facie* right to the claim of being *the* correct logic. The point of departure for the anti-realist and our realist is our actual inferential practice. The Dummettian project consists in examining, criticizing and, where the need arises, amending our inferential practice and hence the use we make of our logical vocabulary on the basis of meaning-theoretic considerations. Given that we do in fact frequently appeal to characteristically classical principles of logic in our ordinary deductive reasoning, we would be misrepresenting matters were we to ascribe the same initial status to classical and intuitionistic logic. The anti-realist's position takes the form of a critique of an established order; he presses reforms against the conservative forces of our undisputedly classical logical practice. Consequently, the burden is on him to produce convincing evidence that this practice really does stand in need of revision. This he can do by demonstrating that our practice (when suitably regimented in a proof-theoretic framework) fails to obey fundamental meaning-theoretic principles—in particular, the principle of harmony. If this is an accurate representation of the dialectical situation, it must be clear that the anti-realist cannot lay claim to an automatic right to equal treatment. Inasmuch as sequent calculi are acceptable codifications of our practice, it does not tell against them that they do not lend themselves to the representation of constructive reasoning.

12.3 Dummett's argument

Given the dialectical situation, are there other arguments anti-realists can avail themselves of? There are. Indeed all that is required is to push Tennant's thought just a little further. The point can be put as follows: sequent systems, like any other type of proof system, are of interest to us only inasmuch as they are interpretable in terms of our ordinary deductive practice. If a multiple-conclusion system is to be more than just a formal game, it ought to be possible to relate sequent proofs to ordinary arguments. However, as Tennant puts it, 'in normal practice, arguments take one from premisses to a *single* conclusion' (Tennant 1997, p. 320, my emphasis). Therefore an acceptable interpretation will transform a multiple-conclusion system into a single-conclusion system. One way—the standard way—to achieve this is to read the commas occurring to the right of the sequent sign (or in the conclusion in the case of multiple-conclusion natural deduction systems) as disjunctions.[7] Thus the sequent $A_1, A_2, \ldots, A_m : B_1, B_2, \ldots, B_n$ should be understood as saying 'Whenever $A_1, A_2, \ldots, A_m$ are assertible, so is either B_1 or B_2 or... or B_n'. This reading seems to be legitimized by the interderivability of A, B and $A \vee B$: the succedent of any sequent of the form $\Gamma : A, B$ can be transformed into a disjunction by a simple application of the $\vee$-introduction rule on the right. Conversely, we can transform any sequent of the form $\Gamma : A \vee B$ into a multiple-succedent sequent:

$$\text{CUT}\ \dfrac{\Gamma : A \vee B \qquad \vee\text{-LI}\ \dfrac{A : A \qquad B : B}{A \vee B : A, B}}{\Gamma : A, B}$$

It is not hard to see that this procedure is readily generalizable to sequents whose succedents contain any finite number of formulas.

But there are severe problems with this interpretation. While we can grasp the way in which the premises function jointly in the antecedent of a sequent without

[7]Strictly speaking we must first interpret any set of formulas on the RHS as disjunctions of their members. Only then does it make sense to interpret the commas occurring on the right of the sequent sign as disjunctions.

having any prior understanding of the meaning of conjunction, no such understanding of the conclusions is possible without already understanding the meaning of 'or'. Therefore the very format of the proof system requires us to have a prior grasp of the meanings of some logical constants. Dummett (1991, p. 187) makes this very point, but again casts it as an argument based on the requirement of neutrality. The argument can be paraphrased as follows:

> The dispute between realists and anti-realists is recast as a dispute over the validity of certain fundamental logical principles. But disagreements about these matters must turn on questions of meaning; the meaning of the logical constants. Therefore a characterization of the meanings of the logical constants in question will be an indispensable preliminary for any such discussion. Moreover, by our inferentialist hypothesis, such a characterization is to be given within the confines of an interpreted proof system that codifies all meaning-theoretically relevant inferential relations. However, if the only possible (informal) interpretation of our proof-theoretic framework necessitates a prior understanding of certain logical operators, it will not be a suitable medium within which to settle questions of legitimacy of any of the principles containing the logical constants in question.

Clearly, put in this way, the argument runs into the same problems we discussed at the end of the previous section. But the anti-realist need not frame the argument in this way. There is no need to invoke neutrality. It turns out that multiple-conclusion systems already fail at a more fundamental level; they are incompatible with the very idea of inferentialism. If, as we said, the only plausible interpretation of such multiple-conclusion systems draws essentially on an understanding of the meanings of at least some logical constants, then such systems cannot play the role required of a proof-theoretic framework. After all, it is the very purpose of such a framework to provide an adequate means for specifying the meanings of the logical constants. On our reading of inferentialism, a system qualifies only if it yields a way of representing what it is a speaker has to grasp in order to be a semantically competent user of the expression in question. On this understanding of what, for the purposes of inferentialism, a proof system has to accomplish, multiple-conclusion calculi constitute a blatant failure, at least under the standard interpretation.[8]

12.4 An objection and its rebuttal

Could it be, however, that the anti-realist's previous argument proves too much? Why is it the case that an understanding of the premises does not likewise require an antecedent grasp of the meaning of conjunction? And so why should single-conclusion systems not be in a similar predicament? After all, it is obviously not the case that 'A' entails 'B' and '$A \supset B$' entails 'B'. Only when the premises A and $A \supset B$ are taken to be *conjunctively connected* can they be said to jointly entail B. We thus have $A_1, A_2, \ldots, A_m \vdash B$ just in case we have A_1 and A_2... and $A_m \vdash B$. Moreover we can—as we did above for the case of disjunctively connected conclusions—establish the formal interderivability of $A, B : \Delta$ and $A \wedge B : \Delta$. Any sequent of the form $A, B : \Delta$ can be transformed into the sequent $A \wedge B : \Delta$ by an application of the left-hand side $\wedge$-introduction rule. Conversely, given $A \wedge B : \Delta$ we get:

$$\dfrac{\dfrac{A : A \qquad B : B}{A, B : A \wedge B}\;{\scriptstyle \wedge\text{-RI}} \qquad A \wedge B : \Delta}{A, B : \Delta}\;{\scriptstyle CUT}$$

Again it is easy to see how this result may be generalized to any number of premises. Should we not then, by parity of reasoning, conclude that single-conclusion calculi too necessitate a prior understanding of the meaning of conjunction? The result would be not so much a disproof of inferentialism as a wholesale disqualification of any proof system with multiple premises (so, in practice, any proof system whatsoever) from playing the role of a proof-theoretic framework.

Fortunately the inferentialist need not despair. For there is an important disanalogy between the conjunctive connection of premises and the disjunctive connection of conclusions. The difference is this. Asserting A and asserting B is in a sense 'tantamount' (Dummett 1991, p. 187) to asserting ⌜A and B⌝. We are not obliged to understand the meaning of 'and', as long as we know how to assert both A and B. This is not to say that there is no distinction to be drawn at all between asserting

[8]Dummett eventually comes around to this view (Dummett 1991, p. 192). I therefore take issue with Restall, who appears to assimilate Dummett's and Tennant's arguments (Restall 2004, p. 16, fn. 6).

A and asserting B, on the one hand, and asserting $\ulcorner A$ and $B \urcorner$, on the other hand. The transition from one to the other requires, in both directions, the effecting of an inference, and it is a mastery of inferences following these patterns that constitutes knowledge of the (logically relevant) meaning of conjunction. This logical distinction notwithstanding, it makes no difference whether another person's claims are reported to me in the form of two separate assertions

> Henry said that aardvarks are nocturnal and he said that they are indigenous to South America.

or as affirming a single conjunctive proposition.

The same does not hold true in the case of disjunction. Here the distinction between

> Henry said that aardvarks are nocturnal or he said that they are indigenous to South America.

and

> Henry said that aardvarks are nocturnal or indigenous to South America.

is crucial. In the second case Henry speaks truly, since aardvarks are indeed nocturnal. In the first case, whether Henry speaks truly or not depends on which of the sentences Henry in fact asserted; he might be wrong. Therefore we cannot, in this case, replace an understanding of the assertion of a disjunction by an understanding of the disjunction of assertions. Indeed, even if we could it would not be of much help. We simply cannot understand the disjunctive nature of the connection between conclusions save by invoking the notion of disjunction itself. Consider an argument of the form $A \vdash B, C$. Surely we cannot read such an argument as issuing an inference ticket to either B or C, whichever we choose, provided only that A is assertible. For if we are warranted, upon asserting A, in asserting either B or C at will, we are in effect warranted in asserting $\ulcorner B$ and $C \urcorner$—obviously this is a much stronger claim. Nor does the following proposal work: 'If you have asserted A, you may either assert B or you may assert C or you may assert both'. For this interpretation still misses the essential point that B or C or both follow from A independently of our wishes and decisions. The conclusion that we need to presuppose a notion of disjunction thus seems inescapable.

If we now ask which notion of disjunction we must presuppose, we do indeed find—Tennant is absolutely correct on this point—that multiple conclusions are intrinsically classical in that we do not in general know which of the disjuncts within the conclusion obtains. Conjoining Tennant's and Dummett's arguments, we thus find that multiple-conclusion systems indeed not only necessitate a prior grasp of the meaning of 'or', but that this meaning must be that of classical disjunction. As such, they do not constitute a suitable inferentialist framework. Inferentialists therefore cannot avail themselves of such systems. May the anti-realist therefore rest his case? Not quite yet. Our argument here relied on the assumption that there is only one acceptable interpretation of multiple-conclusion arguments: the disjunctive one. Therefore, if the realist could offer an alternative interpretation that avoided the same problems, there may yet be a way out for him.

12.5 Bilateralism—an escape route?

Both Tennant's and Dummett's arguments assumed that the only inferentialistically acceptable interpretation of multiple-conclusion systems is the disjunctive reading. Our informal deductive reasoning proceeds by the construction of arguments, and arguments lead from premises to a *single* conclusion. Therefore the only way in which a multiple-conclusion system can be matched with our ordinary practice is by interpreting it as a single-conclusion system with a disjunctively connected conclusion. This was Tennant's point.

We have said, on the other hand, that the only way for the realist to escape the conclusions of the previous section while still being able to reap the benefits of multiple-conclusion systems is to devise an alternative interpretation. The realist's task is thus that of rendering arguments (or sequents) in such systems intelligible without presupposing an antecedent grasp of the notion of disjunction or any other logical constants. How is this to be achieved? Is there any room here for the realist to manoeuvre? If so, what shape might such an interpretation take?

One way of approaching the problem is by invoking the notions of rejection and denial.[9] The central idea is to introduce denial as a symmetric counterpart to the

[9] These notions have been the object of recent work by a number of authors, e.g. Restall 2005, Rumfitt 2000, Smiley 1996.

speech act of assertion. Similarly the notion of rejection may be understood as a negative mental attitude alongside the positive attitude of acceptance. I accept a statement if I judge it to be true; I reject a statement if I judge it to be *untrue*. Corresponding to the internal states of acceptance and rejection, we have the outward manifestations in the form of the speech acts of assertion and of denial. The crucial point is that the speech act of denial and its associated mental state are taken to be conceptually prior to the negation operator. It's one thing to deny a statement; it's quite another to assert the negation of that statement. Even if it turns out that the notion of denial should be construed so as to assimilate the assertion of the statement ⌜not-A⌝ and the denial of the proposition A (effectively identifying a proposition's being untrue with its falsity), as is the case with classical negation, the two are not 'the very same thing' (Smiley 1996, p. 1). In the first case we are dealing with a sign of pragmatic force; in the second case with a logical operator. The denial-theorist, who is all too aware of the difference, does not intend to replace the latter notion with the former. Even in the case of classical negation, there remains a residual difference between expressing assent to the negation of A and expressing dissent from the statement A, just as there is a difference between asserting ⌜A and B⌝ and asserting A and asserting B. The aim, rather, is to give a more complete account of our inferential practice and/or to give an inferentialistically satisfactory account of classical negation.

In the present context, however, the question is whether the availability of this new pragmatic tool also enables the classicist to give an alternative reading of multiple-conclusions without presupposing any familiarity with the meanings of the logical constants. And it seems as if such an interpretation is indeed available. We may interpret a sequent of the form: $A_1, \ldots, A_m : B_1, \ldots, B_n$ as follows:

> It is incoherent to assert each of $A_1, \ldots, A_m$ while simultaneously denying each of $B_1, \ldots, B_n$.[10]

Reading sequents in this way—call it the *denial-interpretation*—allows us to do away with the disjunctive reading and thus appears to eliminate the problems associated with the conventional interpretation. However, before we may hope to have resolved the realist's difficulties, we must ask whether our new interpretation can be accepted by the inferentialist. This, it would appear, is doubtful.

[10]Cf. Restall 2005 and Smiley 1998.

Let us grant, for the sake of the present argument, that denial has a place as a speech act alongside assertion and that it has a central (perhaps symmetric) role to play in the determination of the meanings of the logical constants. We put to one side for the time being the question of the exact relation between the mental act of rejection and the linguistic act of denial and their respective functions in an account of meaning.[11] We also need not worry about what exactly is meant by denial (whether, for instance, it is appropriate to deny any statement one is not in a position to assert, or whether there may be statements that are neither assertible or deniable).

The crucial point is rather that the denial interpretation constitutes a marked departure from the bilateralism of Smiley and Rumfitt. Smiley and Rumfitt seek to raise the notion of denial to the status of a speech act 'on all fours' with assertion (Smiley 1996, p. 1). The initially plausible idea is that both types of act have an equally important meaning-theoretic role to play. In particular, both are equally instrumental in fixing the meanings of the logical operators. What is needed, therefore, is a proof system that lays down a complete set of inference rules on the basis of both types of speech acts. And this is precisely what Smiley and Rumfitt deliver. The standard assertion-based rules regulating the usage of a given logical constant are supplemented with rules stating the inferences to which we are entitled by virtue of having denied statements involving the constant in question. Such systems lay down introduction and elimination rules for each logical constant specifying when a statement of that form may be denied as well as asserted. Importantly, however, both Smiley and Rumfitt are concerned exclusively with *single-conclusion* calculi; neither author seeks to employ bilateralism for the purposes of vindicating multiple-conclusion calculi.[12] Indeed Rumfitt explicitly repudiates such systems for reasons discussed in section 11.1 (see Rumfitt 2000, p. 794–796). Rather, he regards the bilateral framework as an *alternative* defence of classical logic from the proof-theoretic argument. As such, it merits careful consideration—more careful consideration than we are able to give it here. At present, however, our sole focus is the question of the legitimacy (from an inferentialist point of view) of multiple-conclusion calculi. Since

[11] All of these questions are disputed and may be settled in such a way as to provide grounds for ruling out the denial interpretation *ab initio*. See for example Dummett (2002) and Rumfitt's reply (2002).

[12] True, the aforementioned (Smiley 1998) does allude to the denial interpretation. But it is his (1996) paper that carries weight for our current discussion and no mention of multiple conclusions is made there.

Restall's use of the denial interpretation is, as far I am aware, the only sustained attempt at justifying multiple-conclusion systems by way of an bilateralist interpretation, we may focus our attention on it (leaving a discussion of the significance of Smiley- and Rumfitt-type single-conclusion systems for another occasion).

In Restall's denial interpretation the notion of denial is deployed in a rather different way than in the systems of Smiley and Rumfitt. For a start, no statements occurring to the left of the sequent sign can be denied on the denial interpretation. All the statements in the antecedent of the sequent carry assertoric force; all the statements in the consequent are denied. Plainly, we cannot have one sign of force within the scope of another. In particular, it makes no sense to assert a statement that is in the scope of the force marker for denial (or vice versa). So there can be no overlap between assertion conditions and denial conditions on this picture. But this means that Restall's explanation of the meanings of the logical constants in terms of how they 'constrain assertion and denial' fails to meet our fundamental inferentialist principle that the meaning of an expression should be statable within the framework of its V- and P-principles. Let me explain. The denial-theorist holds that an account of meaning based solely on the notion of assertion is insufficient; a complete account also takes denial conditions into account (i.e. the conditions under which a statement made by means of a sentence containing the constant in question in a dominant position is denied) and the consequences of denial (i.e. the consequences of denying such a statement). Therefore, the denial-theorist is committed to delivering a complete account of the assertibility conditions and the consequences of asserting *and*, on top of such an account, he promises to provide a similar account for denial. But clearly, Restall provides no such thing. All we learn is when it is inappropriate to deny a statement containing the said operator in a dominant position (relative to the statements that are simultaneously endorsed), and we learn—since the LHS introduction rules simulate elimination rules, as we have seen—what follows from assertions of statements with the constant in a principal position. We are thus in effect offered introduction rules for denial and elimination rules for assertion. But even these rules do not offer us an explanation of the conditions under which the assertion or the denial of a statement is legitimate. Rather, all we are told is that such and such a statement cannot be coherently asserted or denied given that we have incoherently asserted and denied certain other statements. The denial interpretation thus provides us with no means of extracting specifications of the two aspects of a constant's meaning for either of the two semantically relevant types of force.

Now it could perhaps be argued that Restall's interpretation does somehow convey the meanings of the constants; that this is all that is required to suitably 'constrain assertion and denial'. If so, however, Restall owes us a meaning-theoretic story to replace the inferentialist one based on the two-aspect model of meaning. In the absence of such an account we must conclude that the denial interpretation fails to establish a suitable link between our theory of the meanings of the logical constants and the use we make of them. This, of course, is not to dismiss the possibility that Restall's account might offer a potentially useful way of understanding the normative impact of logic as codified in multiple-conclusion systems.[13] It does mean, however, that multiple-conclusion systems cannot be rehabilitated to meet inferentialist standards with the help of the denial interpretation. As we noted two paragraphs up, we have not ruled out that the notion of denial as it figures in single-conclusion systems of the kind proposed by Smiley and Rumfitt might still open the door to an effective classicist defence against proof-theoretic arguments. All we have argued here is that an appeal to the speech act of denial is of no help when it comes to the question of the inferentialistic acceptability of multiple-conclusion systems.

[13]Significantly, Restall himself views this as the sole aim of his account: 'once one reads this turnstile as a form of *consequence* from X to Y, one must read X and Y differently—it is the *conjunction* of all X that entails the *disjunction* of all Y' (2005, p. 8, fn. 3).

Chapter 13

The seemingly magical fact

Is this enough to put multiple-conclusion systems in their place at last? Not yet. For it seems there is yet another way for the advocate of multiple-conclusion systems to evade such a conclusion. He will have to concede that the denial interpretation is not inferentialistically tenable. Yet perhaps there is a rather straightforward way in which he could appeal to our common-sense understanding to salvage the disjunctive reading after all. The thought is this. According to the inferentialist, a system that adequately represents our inferential practice should provide all that is required for somebody completely innocent of any understanding of the meanings of the logical constants to be able to acquire a working grasp of all of them. It is precisely on this count that the multiple-conclusion system fails. But can we not conceive of a sequential acquisition process? We might begin by familiarizing the speaker with the concept of disjunction within the context of a single-conclusion framework. Then we would proceed by introducing him to a multiple-conclusion setting and, within it, to the remaining logical constants. What would be objectionable about such a step-by-step process that uses the single-conclusion system as a spring board (to avoid the overworked Tractarian ladder-metaphor)?

A quick answer is that this still does not give us *one* unified account that codifies the use we make of the logical constants. Now it is unlikely that the classicist will be overly impressed with this. But there is a much more important reason why this last stance from the defender of multiple-conclusion systems will be to no avail. To see this, the anti-realist must dig deeper and identify the source of non-constructivity in such systems. He must show *what* the features are that account for non-constructive content, *how* they function, and finally *that* they violate meaning-theoretic principles. Of course, in so doing, the anti-realist must be cautious to appeal only to such meaning-theoretic principles as the realist will accept. In other words, what is required is an understanding of what Hacking calls the 'seemingly magical fact': the fact that the difference between intuitionistic and classical logic amounts—in standard proof systems—to the apparently innocuous question of whether or not to allow for more than one formula on the right-hand side of the sequent sign, the question, that is, of whether or not to admit multiple conclusions.[1]

To understand the transition from LJ to LK systems, it is convenient to examine the relation between natural deduction and sequent systems. As we did earlier in our discussion of harmony in a sequent setting (see section 10.1), we will reveal the commonalities between both systems by exploiting the notion of a 'translation' between N- and L-systems. Translating sequent inference rules into natural deduction rules will make explicit the classical commitments that remain implicit in the sequent formulation.

13.1 Milne's explanation

Peter Milne (2002) offers an explanation of the magical fact along similar lines, which is worth spelling out in some detail. It should be noticed before we start, however, that Milne's use of 'translation' can be somewhat misleading. The term standardly designates a proof of coextensiveness of two formalisms, for instance a proof that for every sequent $\Gamma : A$ provable in the system LJ there is a derivation in NJ of A from Γ, i.e. $\Gamma \vdash_{NJ} A$. Translations in this standard sense, are also important in Milne's work. They proceed by induction on the length of derivations, where the length of a derivation is the number of inference rules occurring between the topmost formula or sequent and the conclusion or end-sequent. To designate the length n of a derivation Π we write $|\Pi| = n$. If our aim is to establish the coextensiveness of the systems LJ and NJ, for example, we show that for all derivations of length one, i.e. instances of axioms of the form $A : A$, there exist equivalent natural deduction derivations—namely, natural deduction derivations of the form $[A]$, where A figures simultaneously as assumption and as conclusion of a zero-step inference. Our induction hypothesis, then, is that for every derivation of length n of a sequent

[1] For simplicity we will focus on the multiple-conclusion sequent calculus. Our discussion carries over straightforwardly to multiple-conclusion natural deduction systems.

$\Gamma : A$, there exists a derivation of the formula A such that all the assumptions on which A depends are contained in Γ. The induction step then consists in showing that, on this assumption, the same holds for any derivation of length $n+1$. With our induction hypothesis in place, it thus suffices to check that, for any of the possible LJ-inference rules introducing the end-sequent (corresponding to step $n+1$), we can find a sub-proof of the corresponding formula in NJ.

To illustrate this, consider the case where $\wedge$-introduction on the left corresponds to the step $n+1$ in the following derivation:[2]

$$\wedge\text{-LI}\ \frac{\begin{matrix}\Pi \\ \Gamma, A : C\end{matrix}}{\Gamma, A \wedge B : C}$$

Our induction hypothesis tells us that there exists in NJ a derivation Π_0 of C from the hypotheses Γ and A for Π, $|\Pi| = n$:

$$\begin{matrix}\Gamma, A \\ \Pi_0 \\ C\end{matrix}$$

But then, *a fortiori*, we obtain the derivation Π_1 of A from the assumptions Γ and $A \wedge B$ by a simple application of the $\wedge$-elimination rule:

$$\begin{matrix}\wedge\text{-E}\ \dfrac{\Gamma, A \wedge B}{A} \\ \Pi_0 \\ C\end{matrix}$$

We proceed in a similar fashion for the remaining inference rules.

Although Milne's discussion exploits equivalence proofs between N- and L-systems, his aim is to *explain*, not to *prove*. The aim is to motivate natural deduction inference rules on the basis of the corresponding sequent calculus rules: 'we—ahistorically, for this was not Gentzen's approach—read off our natural deduction rules for classical logic from its sequent calculus formulation' (Milne 2002, p. 515). That is, for the purpose of exorcising the magic we engage in a piece of counterfactual history by imagining that Gentzen took the sequent calculus as the

[2]The right-hand side introduction rules being structurally identical in both systems, the interesting cases concern introduction rules on the left.

starting point on the basis of which he then devised the system of natural deduction.[3] Milne first shows how the rules for NJ can be derived from the corresponding rules in the system LJ. He then examines the case of the corresponding systems of classical logic, identifying two RHS introduction rules in LK that, when translated, give rise to characteristically classical rules of inference. It is worth briefly illustrating this procedure. In the intuitionistic case, the N rule for $\supset$-elimination, for example, is obtained from the LHS L rule for $\supset$-introduction in the following way. Starting with the LHS introduction rule in LJ

$$\supset\text{-LI}\ \frac{\Gamma_0 : A \qquad B, \Gamma_1 : C}{\Gamma_0, \Gamma_1, A \supset B : C}$$

we deduce the matching natural deduction elimination rule, $\supset$-elimination or *modus ponens*, using the CUT rule as the only other rule:[4]

$$CUT\ \frac{\Gamma_0 : A \supset B \qquad \supset\text{-I}\ \dfrac{\Gamma_1 : A \qquad B : B}{A \supset B, \Gamma_1 : B}}{\Gamma_0, \Gamma_1 : B}$$

The remaining rules can be derived in a similar way.[5]

The magic starts (and ends) when we look at the 'translation' from LK into NK. In the classical sequent calculus, to remind ourselves, we allow for several members to the right of the sequent symbol. Moreover, we read the succedent of the sequent disjunctively. What happens if we try to 'translate' from LK to NK using this generalized conception of sequents? As we will see, what is really happening here is that the multiple-conclusion sequent rules are giving rise to natural deduction rules that introduce or eliminate formulas in which the constant governed by that rule is not in a dominant position. Take for instance the previously trivial case of the $\wedge$-introduction rule. It now has the form

[3]In reality, of course, the opposite is true: Gentzen first devised the system of natural deduction and then, purely as a matter of mathematical expediency, introduced the sequent calculus.

[4]Note the aforementioned difference between an equivalence proof properly so called and the type of explanation offered by Milne. We are here 'translating' from LJ into NJ, but the deduction of the elimination rule in the latter system via the application of the CUT rule draws on methods that are used in the inductive step of the proof establishing that

$$\text{if } \Gamma \vdash_{NJ} A, \text{then} \vdash_{LJ} \Gamma : A$$

That is, we are in fact working in the opposite direction, showing that we can simulate NJ elimination rules with the matching LHS introduction rule and the CUT rule.

[5]See (Milne 2002, p. 503–509).

$$\wedge\text{-IR}\ \frac{\Gamma_0 : A\{\vee C\} \qquad \Gamma_1 : B\{\vee D\}}{\Gamma_0, \Gamma_1 : (A \wedge B)\{\vee(C \vee D)\}}$$

where the strings within braces combine with the formula immediately to their left to form disjunctions whenever there are formulas to take the place of C and D.[6] (In the limit case where both slots are unoccupied, we of course find ourselves again within the intuitionistic system.) Thus whenever at least one of the two premises contains more than one member on the right-hand side of ':', the $\wedge$-introduction rule introduces a conjunction into a *subordinate position* with respect to $\vee$, which is the principal operator of the conclusion. That is, the conjunction is in effect being introduced into a subformula of the formula that constitutes the succedent of the conclusion whose principal connective is $\vee$. This generalized form of the $\wedge$-introduction rule on the right can be shown to be derivable in NJ, but this is not trivial (see ibid., p. 512).

Certain inference rules like this $\wedge$-introduction rule remain essentially unperturbed by the admission of multiple conclusions. The meaning of $\wedge$ has not been altered by the modified rule, since any conjunctions (whether embedded in disjunctions or not) could have been derived already on the basis of the standard rules. The same is true for the remaining propositional constants with two crucial exceptions: the RHS $\neg$-introduction and $\supset$-introduction rules. In the first case, the NK-rule that follows from the multiple-succedent version of the right $\neg$-introduction rule is this:

$$\begin{array}{c} \Gamma, [A] \\ \vdots \\ \underline{\quad C \quad} \\ \neg A \vee C \end{array}$$

In the case of the $\supset$-introduction rule, we end up with the following NK-rule:

$$\begin{array}{c} \Gamma, [A] \\ \vdots \\ \underline{\quad B\{\vee C\} \quad} \\ (A \supset B)\{\vee C\} \end{array}$$

[6] In general C and D could of course be sets of formulas. However, on the disjunctive reading adopted here we can equally well treat them as the disjunctions of their member formulas, since C and D are finite. Thus this justifies our mode of presentation.

From the first rule we immediately obtain the law of excluded middle. From the second we can derive the classical formula $(A \supset B) \vee A$ without any further assumptions. We thus obtain the full system of classical logic simply by replacing either the ordinary $\supset$-introduction rule or the ordinary $\neg$-introduction rule in a standard natural deduction system for intuitionistic logic by their non-separable counterparts. In short, our translation, in the cases of the RHS introduction rules for $\neg$ and $\supset$, delivers stronger, non-constructive natural deduction rules.[7]

The key to understanding the 'magical fact' resides in the insight that the acceptance of several members in the succedent leads to natural deduction introduction rules in which the constant in question is introduced into a subordinate position with respect to the disjunction that figures dominantly in the conclusion. Our disjunctive reading of multiple-conclusion sequents thus enables us to repackage the sequent rules as *single*-conclusion natural deduction rules. It is precisely this mode of presentation that makes Milne's analysis so illuminating.

Now, while for some rules the introduction into disjunctive contexts has no impact on the meaning of the connective, nor therefore on the overall strength of the system (cf. Milne, op. cit., p. 514), in the case of the RHS rules for negation and the conditional

> the extra logical strength imparted to the rules is not compensated for elsewhere in the system; the natural deduction translation of either one gives us a rule inadmissible by intuitionistic standards, a rule that suffices, when added to otherwise standard rules for $\wedge, \vee, \supset$ and $\neg$ to yield a formulation of classical logic (ibid., p. 515).

Milne therefore concludes that

> the difference between classical logic and intuitionistic logic is that the former, but not the latter, sanctions the introduction of $\neg$ and $\supset$ to a position subordinate to an occurrence of $\vee$. This is the explanation of Hacking's seemingly magical fact (ibid.).

Whether this really is, as Milne says, *the* difference between classical and intuitionistic logic (as opposed to *a* difference) need not concern us here. Suffice it to say that it is one of the guises in which the difference between these two systems can

[7]It is for this reason that the rules for negation and conditional specifically require modification in the intuitionistic multiple-succedent system described above, cf. section 12.1.

manifest itself. And it does explain the miraculously elegant way in which the step from constructive to classical reasoning is taken in the sequent setting. Having analysed the workings of the multiple-conclusion systems, the question is whether the anti-realist can exploit such insights for his ends, and if so, how.

13.2 The principle of separability

To see how the anti-realist might go about doing this, let us step back to see the relevance of the considerations of the previous section for the anti-realist's case. Recall that we are investigating possible ways to respond to the classicist who dodges the anti-realist's proof-theoretic argument by appealing to the sequent calculus. The anti-realist's response to such a manoeuvre cannot simply consist in rejecting multiple-conclusion sequent calculi on the grounds that such systems are somehow intrinsically classical. As we saw in our discussion of Tennant's argument against multiple conclusions (see section 12.2), such a response must rely on considerations of precisely the kind the proof-theoretic argument attempts to supplant, and thus leads to vicious circularity. Rather, the anti-realist's attacks on multiple-conclusion calculi must be launched from the neutral soil of widely acceptable proof- and meaning-theoretic motivations. This is where Milne's explanation enters the picture. Milne's analysis of how the step from single to multiple conclusions in sequent systems encapsulates the transition from constructive to classical logic may provide the anti-realist with just the tools he is after: having gained an understanding of the interrelations between natural deduction and sequent systems, we can see where exactly, from an anti-realist point of view, things go awry. The anti-realist's task must now be to identify the meaning-theoretic implications of this move. It is here that an argument against the legitimacy of multiple-conclusion systems is to be sought. If there are meaning-theoretic grounds for objecting to multiple conclusions, this is where we would expect them to manifest themselves.

But we have already seen where the source of non-constructivity lies. The question now is whether the potentially deviant feature of the strengthened rules for negation and the conditional—namely, that they introduce the connective (partially) governed by them into a *subordinate* position—constitutes a genuine violation of meaning-theoretic principles, a violation that the classicist would also have to recognize. Assuming that the meaning of a constant is determined by the rules of inference it obeys (its introduction and/or elimination rule or some combination thereof), it is significant that the strengthened rules for negation and conditional invoke logical operators other than the one being introduced. This means that the meaning of one constant is thereby linked to that of at least one other. The question is whether rules that involve more than one operator in this way contravene meaning-theoretic strictures. Are there any reasons for believing that the meanings of the logical constants should be separable and expressible in terms of rules treating but one constant at a time? Or may we allow for cases where the meanings of certain constants are tied to those of other constants?

To tackle these questions, it will be useful to introduce some vocabulary. We may—following Dummett—call a rule of inference *pure* if only one logical constant occurs in it. Moreover, a rule is said to be *simple* if every logical constant occurring in it figures as the principal connective of the sentence in which it occurs (see Dummett 1991, p. 256). Dummett's definitions are closely related with Tennant's *principle of separability*, which in turn goes hand in hand with the *principle of analytic systematicity*.[8] Analytic systematicity requires there to be 'basic rules for each logical operator' and that these be finite in number. The principle of separability requires that each rule deal with 'one dominant occurrence of a logical operator at a time' and be 'purely schematic' elsewhere (cf. Tennant 1997, p. 316–317). In other words, according to Tennant, each logical constant ought to be governed by a specific set of rules (a set of introduction rules and a set of elimination rules), the formulation of which should not explicitly mention any other logical constants. Let us use '$\neg'$' and '$\supset'$', respectively, to denote the constants governed by the strengthened (single-conclusion) natural deduction rules corresponding to the RHS introduction rule for $\neg$ and for $\supset$. Since the strengthened rules for $\neg'$ and $\supset'$ appear also to modify the meaning of the disjunction operator, we may likewise introduce '$\vee'$' to designate the strengthened version of disjunction. Clearly the introduction rules for $\neg'$ and for $\supset'$ violate the requirements of simplicity and purity and thus also violate Tennant's principle of separability. In both cases, the operator introduced by the rule is inseparably linked to disjunction, which is the principal connective in the conclusion.

With these new lexical resources to hand, we may now rephrase our central ques-

[8]Up until now, we have assumed the principle of separability. We must now provide grounds for our assumption.

tion: Must all logical constants be separable (and hence governed by simple and pure introduction rules)? To answer this question, we first note that although our assumption of minimal molecularism (see section 2.4) tells us that the logical constants enjoy a kind of semantic autonomy, it tells us nothing about the internal semantic relations the logical constants entertain with one another. Are the meanings of the logical constants wholly independent of one another? Or do they form a cluster in which all or at least some of the constants are semantically interdependent? We can conceive of three types of semantic structure:

1. **Logical holism**: The logical constants form a semantic cluster, and understanding of the meaning of any one constant presupposes a grasp of all the others. They must therefore be learned *en bloc*.

2. **Logical molecularism**: The relation of dependence between the meanings of the logical constants is not circular, as in the case of holism, but rather asymmetric. The meaning of a given constant may depend on that of another, but the logical constants can be acquired progressively by following the (partial) order of dependence.

3. **Logical atomism**: The meaning of each logical constant can be specified and learned separately. No appeal to the meanings of any other expressions is necessary.

We thus find that the *internal* semantic structure of the logical fragment reflects the familiar approaches to accounts of meaning in general. The microcosm recapitulates, as it were, the macrocosm. Gentzen himself was a proponent of logical atomism (in the sense just defined, not Russell's). And there can be no doubt that atomism makes for a particularly elegant theory. But, attractive symmetries aside, what meaning-theoretic basis is there for holding it?

13.3 Logical holism and logical molecularism

Surprisingly, one searches in vain for arguments in favour of atomism in the writings of its proponents. Despite his heavy reliance on atomism, Tennant, for example, limits himself to the following comment in his *Anti-realism and logic*:

> The mastery of any one colour-word requires mastery of all (or at least some of) the others in the language. [...] But by contrast it is difficult to point to expressions whose mastery is required for, or is sufficient for, mastery of the logical operator 'and'. For mastery of 'and', it would appear, rather, that one needed mastery (albeit implicitly) of the general concept of assertion, acquired through the use of whole sentences, and of warrants of such assertions. [...] It shares these features with all the logical operators: they offer outstandingly secure points of entry to the molecularist (Tennant 1987, p. 63).[9]

Tennant is of course right that it is 'difficult to point to expressions' upon which the meanings of the logical constants depend *but only insofar as* we take the standard separable, pure Gentzen-type formulations of natural deduction for granted. But the question is what we should make of rules that are not of this ilk. Clearly, in the present dialectical situation, Tennant's comment amounts to no more than a restatement of the position of logical atomism. Tennant simply affirms precisely what is at issue: that the logical constants be individually graspable.

In a similar vein, Dummett assures us that

> the logical constants form a uniquely simple case, since they *do not satisfy the generality constraint*: to understand ⌜A or B⌝, one need not understand ⌜A and B⌝ or ⌜if A, then B⌝ (Dummett 1991, p. 223, my emphasis).

Again, no argument is offered. Interestingly, in a later passage of the same work Dummett appears to allow at least for molecularism. Commenting on the idea that all introduction rules should be like those presented by Gentzen—pure, simple and single-ended (i.e. either an introduction or an elimination rule, but not both at once, see section 2.5)—Dummett remarks:

> Reflection shows that this demand is exorbitant. An impure **c**-introduction rule [where **c** is a logical operator] will make the understanding of **c** depend on the prior understanding of the other logical constants figuring in the rule. Certainly we do not want such a relation of dependence to be cyclic; but there would be nothing in principle objectionable if

[9]The last sentence is taken slightly out of context. Nevertheless, I believe that the quote, as presented here, is faithful to the spirit of the passage.

> we could so order the logical constants that the understanding of each depended only on the the understanding of those preceding it in the ordering (Dummett 1991, p. 257).

Why ought the dependence relation not be cyclic? In other words, why does Dummett declare logical holism 'certainly' untenable?[10]

Well, Dummett is certainly right to dismiss cyclic dependence relations if they commit us to saying that the conditions under which a statement containing a given constant, say \$, as a principal operator is assertible will appeal, directly or indirectly, to *all* the remaining constants. For an understanding of \$ will consist in grasping all of the associated \$-I rules, that understanding assumes prior knowledge of the meanings of all the remaining constants. However, according to our stipulations, each of the other constants will in turn presuppose an understanding of \$. This type of circularity, surely, is intolerable because it would render a compositional account of meaning impossible.

Imagine, for instance, that we are faced with a sentence of the form $A \wedge B$ for complex A and B. An understanding of this sentence consists in an understanding of A, B and $\wedge$. However, understanding $\wedge$ consists in a grasp of the conditions under which it is admissible to assert conjunctions, and on the picture we are considering, these conditions will involve other logical operators, as in $\neg(\neg A \vee \neg B)$, say.[11] The inference from $\neg(\neg A \vee \neg B)$ to $A \wedge B$ is treated as primitive as discussed in section 2.5. I.e. there cannot be more *direct* ways of deriving $A \wedge B$, e.g. from its canonical grounds A and B. It follows that the meaning of $\wedge$ is dependent on that of $\vee$ and $\neg$. But according to our stipulations some of the introduction rules for *these* connectives will in turn appeal to other connectives (which are partially determinative of *their* meanings). This will eventually bring us back full circle to $\wedge$. On this kind of rampant logical holism, there is no guarantee that our meaning analysis will ever reduce the complexity of the sentences whose meaning we are required to understand,

[10]Note that we are making the plausible assumption that for any constant either its introduction rules or its elimination rules are determinative of its meaning. It is never the case that some of the introduction rules *and* some of the elimination rules are jointly determinative of meaning, nor that meaning is determined by both sets (taken in their entirety) together. (Given the principle of functionality this would be redundant, see section 5.3.) Let us further assume, for the sake of simplicity (but without loss of generality), that introduction rules enjoy meaning-theoretic priority. The following applies *mutatis mutandis* to the opposite view that elimination rules (in some or all cases) fix the meanings of the logical operators.

[11]A rule of this kind flaunts the complexity condition formulated in 2.5. The present argument is a further point in its favour.

let alone that the process will bottom out at the atomic level.[12]

Such reflections, however, hardly constitute a conclusive argument against holism. Far from it. Indeed, for all we have said, it remains entirely reasonable for the realist to argue that, if the proof-theoretic argument really entails that any of the possible classical systems (or at least any of the known formulations of it) violate the principle of separability, then we should not callously jettison our treasured classical logic, but rather subject the assumptions underlying the proof-theoretic argument to very careful scrutiny. And the principle of separability might be just the vulnerable link in the anti-realist's argument that the realist needs. Surely, the realist can insist, in a face-off between giving up classical logic and dropping the principle of separability, classical logic must have the upper hand. That is, it should be granted priority unless a thoroughly convincing case can be made that the principle of separability is deeply anchored in the most fundamental principles that underlie our conceptions of meaning and logic.

Prima facie we have good reasons to doubt that such a case can be made. Is it not in fact rather plausible that there should be some degree of interdependence between the constants? After all, why should it not be the case that the meanings of, say, 'or' and 'not' are somehow interwoven? Could it not be said—relevantists, plug your ears!—that the law of disjunctive syllogism is partly constitutive of the meaning of $\vee$?

Milne's proposal is inspired by just such intuitions. On his picture, we get classical logic without having to introduce multiple conclusions simply by adopting either of the aforementioned liberalized introduction rules for negation or conditional. The price to pay is precisely that we are violating separability: either the meaning of the conditional or that of negation will be tied up with that of disjunction. Consider again the rules for $\supset'$. What is peculiar about them is that, given a derivation of the form

$$\begin{array}{c} \Gamma, [A] \\ \vdots \\ B \vee' C \end{array}$$

[12]Note that the problem consists in the circularity of the relation of *meaning-theoretic* dependence, not necessarily in the fact that the order of acquisition is circular. For even if the meanings of the logical constants were interdependent in the way described, there might still be ways of breaking into the circle. There is no reason why it should not be possible to learn the meanings progressively by adjusting the meaning of each constant at each stage when a new constant is added to the repertoire.

we have a choice of either inferring

$$\frac{\begin{array}{c}\Gamma, [A] \\ \vdots \\ B \vee' C\end{array}}{A \supset' (B \vee' C)}$$

in which case the inference is simply a special case of our usual ⊃-introduction rule; or inferring

$$\frac{\begin{array}{c}\Gamma, [A] \\ \vdots \\ B \vee' C\end{array}}{(A \supset' B) \vee' C}$$

in conformity with the liberalized '$\supset'$-introduction rule'. As the scare quotes indicate, the status of the rule is not entirely clear. It may introduce either a disjunction or a conditional. It is important to note that, contrary to Milne's claims, the admission of either rule is unjustifiable on the basis of molecularist principles. For the relation of meaning-theoretic dependence obtaining between the operators concerned is not in fact asymmetric: it is true that we must have mastered the meaning of $\vee'$ in order to understand the meaning of $\supset'$ in terms of which it is defined, but conversely we cannot be said to have full knowledge of the meaning of $\vee'$ until we have mastered the liberalized rule above. Yet we cannot grasp the rule without having a prior understanding of the meaning of $\supset'$. Neither constant can be grasped individually, but only as part of the double pack. Under closer inspection, the '$\supset'$-introduction rule' turns out to be simultaneously an introduction rule for $\vee'$ and for $\supset'$. The same is true for the liberalized negation rule: here, too, the meaning of $\neg'$ is inextricably bound up with that of $\vee'$ and vice versa.

It follows that if there are two speakers, who have mastered all the propositional connectives, and one operates with the standard Gentzenian inference rules while the other abides by one of the liberalized rules, then the two will use the conditional and disjunction (or negation and disjunction) in different ways: although both speakers may employ the same symbols, they will attach different meanings to them.[13]

[13]This has repercussions also for the realism/anti-realism dispute. Read asserts that on this formulation (involving the '$\supset'$-introduction rule') it is no longer classical negation that marks the difference between classicists and intuitionists.

13.4 An argument for separability

What are we to make of Milne's systems obtained by adopting either of the strengthened rules of inference? Clearly, even though both of the two cases of the operators must be learned simultaneously, the rules for $\supset'$ and $\neg'$ do not run into the same problems faced by the viciously circular dependence relations discussed in the previous section. That is, Milne's rules do satisfy the complexity condition. Should we still be concerned about the fact that these pairs can only be understood in aggregate?

Are there any arguments the atomist can avail himself of? One intuition the atomist might exploit is that logical constants ought to be graspable individually and hence that they should be learnable in a piecemeal fashion, in any order. Can this intuition serve as a ground for an argument? Well, let us ask what is involved in rejecting it. The holist is committed to the claim that someone who mastered only a proper subset of the complete set of the logical constants does not have a mastery of the *same* constants as someone who masters the entire set. Detached from their normal habitat, the symbols no longer have the same meaning. Therefore, a community whose repertoire of logical constants comprises only $\wedge$ and $\vee$, say, will attach different meanings to '$\wedge$' and '$\vee$', since he will not be in possession of the conditional and negation. This holds true even for partial holists who allow for clusters among the logical constants all of which are semantically interdependent, but where none of the clusters exhausts the entire set of logical constants. Milne's systems are of this type. Suppose we added the liberalized rule for the conditional to an intuitionistic system $\{\vee, \wedge, \neg\}$. Call the system obtained S. Someone operating in the system $S' = \{\vee, \wedge, \neg\}$ or even with $S'' = \{\vee, \wedge, \neg, \supset\}$ (where $\supset$ obeys the familiar introduction and elimination rule) will not mean the same thing by '$\vee$' or

> NK [...] is deeply misleading about the difference between classical and intuitionistic logic. Its formulation suggests that the two systems agree on disjunction, the conditional and so on, differing only in their treatment of negation. But NC [Read's multiple-conclusion natural deduction system incorporating the liberalized rules in question, NKM in our terminology (see section 11.4)] shows that formally, the calculi can be made to agree on disjunction and negation and disagree on the conditional (Read 2000, p. 151).

We can now see that he is wrong on this last point. The point of divergence, given either of the two non-separable systems we considered, must reside in the difference of meaning between either of the two interdependent pairs of connectives ($\supset'$ and $\vee'$ or $\neg'$ and $\vee'$), and their separable counterparts ($\supset$, $\vee$ and $\neg$) rather than in the conditional alone.

by '$\vee$' and '$\supset$', respectively. (Thus we have the need to introduce $\vee'$ and $\supset'$.) We could imagine similar scenarios on the molecularist conception. The molecularist denies that someone who does not acquire the constants progressively in accordance with the order of the asymmetric dependence relation can attach the same meaning to the symbols that he and the molecularist apparently share.[14]

For such a holist (or partial holist) it is thus inconceivable that somebody could master $\supset'$ while having no conception of disjunctions. Assuming that some of our constants really are semantically intertwined in something like the way described in Milne's examples, it would follow that we cannot so much as imagine a people that operates with some but not all of these co-dependent constants and yet ascribe to the 'shared' constants the same meaning as we do. On the one hand, it seems bizarre that a community innocent of some of our logical expressions should be unable to attach the same logically relevant meaning to 'or', say, as we do. Similarly, why should it not be possible to acquire the same logical constants in various orders?

To this it might be retorted that even Gentzen's formulations of natural deduction impose a certain order in which the meanings of the constants have to be learned. As Kneale points out, certain rules that are part and parcel of any natural deduction system—namely the improper ones involving the discharge of hypotheses—appeal no less to other constants; they just make this appeal tacitly. This, according to Kneale introduces 'something awkward into the system' because these rules 'depend in a certain way on other rules. For what follows from a supposition must follow in accordance with some principle, and obviously this cannot be the principle formulated by means of the supposition' (Kneale 1956, p. 244). The upshot seems to be that one would have to master at least some of the constants governed by proper inference rules before one could grasp any of those whose meaning is given by improper ones. Improper rules involve subproofs, and subproofs require rules of inference as their deductive raw material. Hence, since these rules must ultimately come from proper inference rules, we need proper inference rules to put the improper ones to work or so the argument goes.

It is undeniable that principles of inference employing premises that take the form of deductions do indeed presuppose *some* means of constructing deductions. However, it is not clear that a grasp of such derivability relations must, in the

[14]In the following paragraph we will, for simplicity, address only the holist (or partial holist). All that is being said applies *mutatis mutandis* to the molecularist as well.

first instance, be purely logical in nature. We could plausibly acquire a grasp of the introduction rule for 'if. . . , then. . . ', say, simply by acquiring competence at drawing chains of non-logical (material) inferences from certain assumptions. Arguments constructed by these means would not, to be sure, be formally valid. Nevertheless, such non-logical inferential chains are all that is required to acquire at least an initial grasp of the meanings of connectives governed by improper rules of inference. We can thus do without an antecedent understanding of logical vocabulary.

So the atomist's intuition, insofar as it is based on natural deduction frameworks, is not self-defeating—natural deduction systems do not in and of themselves violate separability. But of course, this does not show that the atomist's intuition should be preserved at all cost, especially if the cost is giving up classical logic. After all, it is highly plausible that there are semantic clusters in other regions of language—we mentioned colour predicates earlier. We generally take it that a grasp of the meaning of any one colour predicate necessitates the mastery of at least some other ones. If this is right, we should not be able to conceive of linguistic communities that use 'red' in exactly the same way we do even though their lexicon contains no other colour predicates. There does not seem to be anything shocking about this conclusion, on the contrary, it seems to be rather profoundly true. What, then, is so aberrant about the thought that similar interdependencies should obtain within the realm of the logical constants? We are thus led to the conclusion that, despite the undeniable intuitive appeal of the atomist's position, the mere intuition that the constants should be learnable individually is not sturdy enough a foundation to mount a conclusive argument.

Does this mean that the atomist—and with him the advocate of proof-theoretic arguments—has to concede defeat? Not quite yet. There remains an important worry in connection with the rejection of the principle of separability: Do constants defined by non-separable inference rules satisfy our requirement of harmony? The worry is both very real and very pressing. For example, consider again the introduction rule for $\supset'$.[15] Using $\supset'$-I we might arrive at a sentence of the form $(A \supset' B) \vee' C$. Now, according to our notion of the first layer of logical form, we are dealing with a disjunction. Since, according to our intuitive conception of harmony, general harmony, it is a necessary condition that there be a levelling procedure for the rules governing a logical constant, we should expect that any maximum of the

[15]The argument to follow applies equally to the '$\neg'$-introduction rule'.

form $(A \supset' B) \vee' C$ should be eliminable using the or-elimination rule. But this clearly will not work, as a mere inspection of a $\vee'$-peak of this kind makes plain:

$$\dfrac{\dfrac{\begin{matrix}\Gamma_0, [A]^i \\ \vdots \\ B \vee' C\end{matrix}}{(A \supset' B) \vee' C}\,{\scriptstyle \supset'\text{-I},\, i} \qquad \begin{matrix}\Gamma_1, [A \supset' B]^j \\ \vdots \\ D\end{matrix} \qquad \begin{matrix}\Gamma_2, [C]^j \\ \vdots \\ D\end{matrix}}{D}\,{\scriptstyle \vee\text{-E},\, j}$$

Since both the $\supset'$-introduction rule and the $\vee$-elimination rule are improper rules, there is no way of concatenating the minor premises with the deduction leading to the major premise so as to create a direct proof of the conclusion that avoids the introduction and immediate subsequent elimination of the major premise. Perhaps, then, the mistake is that we should have appealed to our $\supset$-elimination rule, rather than to the $\vee$-elimination rule? After all, we are dealing with an introduction rule for $\supset'$, which is just a liberalized form of the ordinary introduction rule for conditionals. However, the $\supset$-elimination rule turns out to be equally ill-suited for the purposes of eliminating our $(A \supset' B) \vee' C$-maximum. And for obvious reasons: the $\supset$-E rule takes only major premises of which $\supset$ is the principal operator. We thus see that there is no way to level peaks of this hybrid kind with our usual elimination rules.[16]

We are inclined to ask what consequences we may draw from a statement expressed by means of a sentence of this form. The introduction rules do not provide us with the faintest clue. They enable us to make inferences that we should not have been able to make under the same circumstances on the basis only of our separable introduction rules. Our intuitive conception of harmony therefore demands that the corresponding elimination rules be strengthened commensurately. But how? Are the individual elimination rules for $\vee'$ and $\supset'$ to be somehow liberalized? It is far from clear how this should work, given that the same rule can introduce either a conditional or a disjunction, as the case may be. Maybe, then, this calls for the introduction of a specifically tailored 'complex' elimination rule to reflect the interdependence of the two constants. But it is far from clear what form such a rule could take.

How persuasive is this argument for separability from harmony? The trouble with it is that our intuitive conception of harmony as a balance between V- and P-principles appears to be biased. The very notion of V- and P-principles (as they

[16] Clearly our readability procedures can be of no help here (see section 8.3).

apply to the logical constants) is informed by standard formulations of natural deduction systems and thus has a built-in prejudice in favour of separability. But in the present case we are precisely confronted with constants that share some of the very same V-principles but have different P-principles. Therefore, it should not come as a surprise to us if non-separable rules like Milne's rules do not satisfy harmony so understood. Our appeal to the notion of harmony is thus obviously circular. Does this mean that we have to abandon the idea of intrinsic harmony as a necessary condition for harmony? Must we completely rethink our conception of harmony in the light of these considerations? Or ought we perhaps to give up the notion of harmony altogether in the absence of an argument for separability?

Not so, I would argue, so long, at least, as we wish to retain the fundamental principles of innocence and of autonomy (see section 3.3). These principles in effect simply spell out what already follows from our assumption of minimal molecularism (see section 2.3), which states that logic is to form a semantically autonomous linguistic realm. The principles of innocence and of autonomy simply lay down the two conditions that need to obtain for logic to enjoy this independent standing. Logical expressions are innocent in that they do not impinge on the meanings of non-logical expressions; conversely, the non-logical regions of language may not distort the meanings of the logical constants. On our conception it is the principle of harmony that guarantees that these two principles obtain. More precisely, we have, drawing on the assumption of separability, argued for a local principle of harmony; a local principle that would ensure the autonomy of logic. If we are to part ways with the principle of separability and thus the notion of a local principle of harmony *while preserving a commitment to the principles of autonomy and innocence*, then we must find a constraint or principle that can fill the void. What we need, in other words, is a new principle that operates at the global level of the logical fragment and functions as a safeguard for the principles of innocence and autonomy. Clearly, such a global version of harmony would be weaker than our proposed local principles. For in addition to permitting the systems that satisfy local harmony, it would also countenance systems in which local violations of harmony are compensated by global and structural properties in such a way that the overall configuration of the system nevertheless satisfies innocence and autonomy.

The only candidates we have encountered that might serve as globally-acting proxies for our locally-acting principle of harmony are the notions of full harmony and normalization.[17] Does either of these notions fit the job description?

Let us start with the former. It does not take much to see that full harmony is particularly ill-suited for the purposes. The classical logician is looking for a principle that allows for semantic interdependence among the constants, and this is precisely what full harmony will not do. Assume we have a system S that contains two constants $\star_1$ and $\star_2$ such that former depends on the latter; i.e. we cannot learn $\star_1$ without first mastering $\star_2$. Even on this minimal form of molecularism, S will not satisfy full harmony, since adding $\star_2$ to $S - \{\star_2\}$ is bound to result in a non-conservative extension. Yet S might easily be well-behaved. Full harmony is thus clearly too strong a notion for the purpose at hand. For recall that we are after a notion that will guarantee that the logical fragment is well-behaved in the sense of satisfying the principles of innocence and of autonomy. To take another example, the classical fragment $\{\neg, \supset\}$ is well-behaved, yet it is not fully harmonious (as we have seen in section 4.6). So full harmony is of no help here.

What about normalization? If the very point of harmony is to ensure the good working order of language, and if the only thing we ask of the logical fragment is that it be well-behaved, then normalization might seem to be the perfect match. For normalizability ensures that any logical deduction can be cleansed of local peaks and plateaux. Since innocence is threatened in just those cases where local peaks are creative of content—i.e. when the elimination of a constant fails to 'undo' its prior introduction and the constant therefore affects the V-principles or P-principles of non-logical expressions—normalization may seem like just the notion we are after. There is an obvious problem, however: normalization requires levelling procedures.[18] We cannot, it seems, have normalization without intrinsic harmony. And, as we have seen, we do not get intrinsic harmony on any of the non-separable rules we have considered. Of course this does not prove that no other non-separable principles will be found that fit the bill. Nor does it exclude the possibility that other types of global constraints can be articulated that guarantee the principles of autonomy and innocence. But I will rest my case here; I believe the burden of proof has now

[17]Recall that a system is fully harmonious if any of the operators it contains is a systematic conservative extension of all the remaining operators. Note that total harmony, as a relational property—the relation holding between a constant and a base system—is not of the right kind.

[18]Prawitz's systems of modal logic are an exception. As we have seen in section 5.5, levelling procedures are generally unavailable for certain systems of modal logic that nevertheless can be shown to normalize by Prawitz's ingenious technique. However, these techniques do not appear to be generalizable to other types of systems.

been passed to the classicist.

Where does this leave us? Let us briefly recapitulate the argument to remind us of the dialectical situation. The classicist, in an attempt to parry the proof-theoretic arguments, challenged the anti-realist's reliance on the principle of separability. Give up separability and you can help yourself to Milne's rules—$\supset'$-I and $\neg'$-I—either of which is sufficient, in combination with intuitionistic logic, to get full classical logic. We then considered possible ways in which the anti-realist could defend the principle of separability. Suppose we do allow non-separable principles, as the classicist wants. Insofar as the classicist remains committed to the principles of innocence and of autonomy, he must find a safeguard that ensures the holding of these principles. I.e. he must find a global analogue to our local notion of harmony. We considered two potential candidates: full harmony and normalization. Full harmony fails on account of being too restrictive. The second, more promising proposal was the notion of normalization. The problem here was that the classicist would have to avail himself of suitable reduction procedures to deal with Milne's separability-defying rules. It seemed that no such reduction procedures were forthcoming for such rules. Of course it would serve the classicist's aims just as well were he to discover alternative non-separable principles that would be adequate for classical logic while still giving rise to a normalizable system. But it is not clear what such principles might look like.[19] Alternatively, the classicist could of course propose a global version of harmony other than normalization. But again it is unclear what this alternative constraint might amount to. Therefore, although the argument we gave is not conclusive, the burden is shifted to the classicist. He must either concoct suitable non-separable principles, or he must come up with an alternative to normalization as a global variant of harmony.

13.5 Summary

This brings our discussion of proof-theoretic arguments to an end. Having assumed that such arguments are successful in the natural deduction framework, we considered possible ways of parrying the anti-realist's conclusions. In particular, we considered the option of displacing the debate from the natural deduction into the

[19] The difficulties encountered by Weir in his discussion of a number of seemingly promising non-separable principles demonstrates that the task of devising such alternative principles, if possible at all, is no simple task (Weir 1984, especially p. 469 and p. 478).

sequent setting. We were forced to acknowledge, however, that the classicist's efforts bore no fruit. His appeal to multiple-conclusion systems turned out to contravene shared inferentialist principles. Attempts at salvaging systems of this form by reinterpreting them with the help of the speech act of denial equally failed to generate an account acceptable by meaning-theoretic principles. Finally, the classicist's attempt to devise a proof- and meaning-theoretically palatable formulation of classical principles through the introduction of non-separable principles has been shown to be a daunting, if not impossible task.

Chapter 14

Conclusion

Let us recapitulate the main conclusions of our foregoing discussions. Having laid down the cornerstones of logical inferentialism—the use-theoretic approach, the two-aspect model of meaning, inferentialism and molecularism—we examined the role structural assumptions play in determining the meanings of the logical constants. We found that natural deduction systems in sequent format most accurately capture the distinction between the structural rules expressing the structural properties of the system as a whole and operator-specific operational rules. In particular, I showed that, contrary to conventional wisdom, the rule of *ex falso* ought to be understood as a *structural rule* closely related to the rule of weakening on the right. We then considered Paoli's distinction between operational meaning, structural meaning and the overall meaning of a logical expression given by the entire set of inferential relations in which it participates. I argued that the notion of a constant's overall meaning and hence that of its structural meaning are groundless. The meanings of the constants are given by the rules of inference they obey and by these rules alone. From this it follows that disputes over logical principles can stem from two distinct sources. The debating parties can differ over the meanings of the logical constants, in which case their quarrel turns ultimately on meaning-theoretic considerations (as in the case of the debate between classicists and intuitionists). Or, as is the case in the debate between relevantists and classical logicians, the opposing factions can differ over the structural assumptions they countenance and hence over their conceptions of validity and acceptable argument structure (as in the case between classicists and relevantists).

Part two developed a novel account of harmony. Following the introduction of the notion of harmony and a survey of the ends for which it has been advanced, we laid down the basic intuitive notion that served as the touchstone for our formal accounts: the notion of general harmony. The fundamental idea was that equilibrium must prevail between the conditions under which the assertion of a sentence containing a constant as its main connective is legitimate (the V-principles) and the consequences of having asserted such a sentence (the P-principles). So long as these two aspects are balanced, we can rest assured that the introduction of a logical constant will enable us to deduce neither more nor less than we should have been in a position to do otherwise. I.e. we can rest assured that the use of the logical constant neither creates nor loses information. With our notion of general harmony in place, we proceeded to assess existing accounts of harmony, beginning with Dummett's work. According to Dummett, harmony as it applies to the logical constants falls into two kinds: a local version, intrinsic harmony, which is identified with the existence of levelling procedures, and a context-dependent version, total harmony. We had argued that systematic conservativeness, while a necessary condition for harmony, is not a suitable candidate for harmony as such, since it fails to guard against P-weak disharmony and illicitly relies on contextual factors. Intrinsic harmony must also be dismissed as a formal correlate of harmony on account of the Q-example. The solution, we concluded, must lie in a revised local principle that incorporates a defence against P-weak disharmony: the principle of stability. Unfortunately, since Dummett does not himself develop such a principle, it was incumbent on us to do so. Before pursuing this project, however, we defended Dummett's conjecture to the effect that stability implies conservativeness. In response to Prawitz and Read we pointed out that the non-conservativeness of the system resulting from supplementing **PA** with a Tarskian truth theory does not tell against stability. The source of non-conservativeness does not reside in the truth predicate as such, but in the ancillary Tarskian apparatus; the adjunction of the T-schema on its own results in a *conservative* extension of **PA** even if we allow the truth predicate to occur in instances of the induction schema. Moreover, the non-conservativeness of the truth predicate over first-order logic does not disprove Dummett's conjecture. The cause for non-conservativeness again does not lie with the rules of inference governing the truth predicate, but rather with the name-forming operator required to put it to use. It follows that the truth predicate is not a logical constant and that Dummett's conjecture stands.

In the next chapter, we argued for the principle of functionality and for its symmetric counterpart, the principle of injectivity. Taken together these principles guarantee that two introduction rules are distinct if and only if the corresponding elimination rules are distinct. On the background of this principle we then addressed the issue of the modal operators: Can modal operators have a place in an account of harmony? And, since harmony implies functionality, are they compatible with the principle of functionality? Although it turned out that the modal operators are not in principle incompatible with functionality, it remained wholly unclear how the standard rules for the modal operators could be incorporated into an account of harmony. Read's attempt to lay down inference rules for the modal operators in the context of a labelled deductive system was also considered and found wanting. For even if it were possible to formulate principles that satisfied harmony in the context of such a system, it would be at the cost of severing all ties with our ordinary inferential practice. There is no way of making sense of the rules of inference save by either appealing to possible world semantics or by resigning oneself to an extreme kind of logical formalism. Neither of these options is open to the inferentialist (or to Read). Our conclusion is thus a pessimistic one; the very notion of general harmony is at odds with our modal notions.

After this interlude, we returned to our main task of devising a principle of harmony in line with the notion of general harmony. We considered Tennant's account of harmony as a possible solution to the problem set out by Dummett. Having considered and dismissed two possible objections against Tennant's proposal, we presented a counterexample of our own: we constructed freakishly strong quantifiers and showed them to be validated by Tennant's principle of harmony. Next we examined Read's notion of generalized elimination harmony. We showed that it too was vulnerable to P-weak disharmony, but contended that it nevertheless offered a promising framework.

Taking into account the shortcomings of the accounts examined, we proposed a possible emendation of Tennant's account. Tennant's principle of harmony did us good service when it came to avoiding P-weak disharmony, but failed when it came to P-strong disharmony. For intrinsic harmony, we found that the converse was true. Why not therefore put the two together? This line of thought gave rise to our *weakened* principle of $\underline{\text{H}}$armony; Tennant's principle of harmony is constrained by the proviso that only such pairs of rules that admit of a levelling procedure may be selected. Though this principle is adequate in principle, it has the disadvantage of being excessively complicated. Moreover, it does not enable us, given a set of introduction/elimination rules, to determine the (harmoniously) matched set of elimination/introduction rules.

In the final chapter of part two, we developed our own account of harmony, which is free of the defects that plagued our predecessors' accounts. We set ourselves two tasks: the readability task of devising a procedure for determining, for any permissible rule (or set of rules) of inference R, its structural counterpart; and the stability task, which amounted to the task of devising a procedure for extracting, from the structural counterpart of R, the harmonious counterpart to R, R'. Read's system of general elimination rules served as an ideal framework for our enterprise. We began by formulating principles of structural correspondence, i.e. principles that, given a specification of R's characteristic parameters (number of rules, number of premises per rule and number of subproofs per rule), enabled us to determine R's structural counterpart. This is our readability procedure. We then formulated a procedure by which to check which, if any, of the rules determined by the structural counterpart of R admit of a levelling procedure with respect to R. If there are no such rules, we have demonstrated that R has no harmonious counterpart; if there are rules that are levellable together with R, we pick the strongest such rule and have our R'. This is our stability procedure. Because R' admits of a levelling procedure together with R, it does not introduce P-strong disharmony; since we chose the strongest levellable rule, we succeeded in warding off the threat of P-weak disharmony.

Aside from avoiding the problems P-weak and P-strong harmony that bedevilled all existing accounts, our algorithmic approach has the advantage of furnishing a systematic method for reading off the harmonious counterpart R' of a permissible rule R whenever such a harmonious counterpart exists. Where there is no harmonious counterpart, our stability procedure demonstrates this. We have thus given not just a characterization of harmony but also a method for determining harmonious rules and an effective test for harmony that can be applied to any pair of inference rules.

With our notion of harmony in place, part three was then devoted to proof-theoretic arguments. The question was this. Assuming that the anti-realist's case for intuitionistic revisions goes through in the context of natural deduction (as our principle of harmony suggests), is there a way for the classicist to neutralize the proof-theoretic argument by transposing it into an alternative proof system? We identified the classical sequent calculus as the initially most plausible candidate for this purpose. Having shown how the notion of harmony developed in part two can be brought to bear on the sequent calculus, we demonstrated that the classical sequent calculus is harmonious. Moreover, we defended the sequent format against the charge of being ill-suited for the inferentialist project. In our search for possible replies on behalf of the anti-realist, our focus shifted to multiple conclusions more generally as a source of non-constructivity. The anti-realist's hope was that some of these sources of non-constructivity could be shown also to be sources of meaning-theoretic irregularities. Tennant's argument to this effect was found to be viciously circular. However, we succeeded in showing how, drawing on an idea of Dummett's, Tennant's argument could be transformed into a potent argument against multiple-conclusion systems. Since the only way to give any intuitive content to multiple-conclusion systems is by reading them disjunctively, such systems in fact presuppose an antecedent grasp of the meaning of disjunction. As a possible comeback for the realist, we then considered an alternative reading of multiple-conclusion systems based on the notion of denial as a speech act alongside that of assertion. But the denial interpretation turned out to conflict with our meaning-theoretic commitments, in particular with our assumption of the two-aspect model of meaning. The realist rejoinder that seeks to exploit Milne's insight that non-constructivity in multiple-conclusion systems ultimately stems from the fact that classical logic countenances inference rules that violate the principle of separability was also shown to fail to advance the realist's cause. The ensuing question as to whether the principle of separability is tenable was then answered in the affirmative: for the principle of separability was shown to be the only candidate presently at hand that can serve to preserve our fundamental conception of logic, as encapsulated in the principles of autonomy and innocence.

What, then, is the moral we are to draw from all these considerations? We have provided answers to the three central questions facing inferentialism. We spelled out in detail the meaning-theoretic prerequisites for logical inferentialism. We proposed a new, rigorous notion of harmony. And we addressed some of the central issues surrounding proof-theoretic arguments. Bringing these three components together we arrived at a coherent picture of logical inferentialism. Does it follow that the logical inferentialist, given his meaning-theoretic commitments and the principle of harmony to which they give rise, must opt for a constructivist logic? Not necessarily. The case that can be made on the basis of our investigations is not watertight, and it was not our aim to make it so. For one, we took it for granted in the third part of this dissertation that proof-theoretic arguments do succeed within the natural deduction setting. This assumption would have to be justified in detail on the basis of the notion of harmony proposed here. A superficial glance at our principle of harmony and at the classical principles required to extend intuitionistic systems into classical ones may give the anti-realist some cause for optimism on this front, but a careful discussion of these matters is left for a future project. Also, as we pointed out, our argument for the principle of separability, although it succeeds in passing the buck to the classicist, nevertheless leaves the realist with some room for manoeuvre. Finally, of course, logical inferentialism is only as strong as the premises on which it rests. The anti-realist's *modus ponens* may just be the classicist's *modus tollens*, at least for the classicist willing to give up the meaning-theoretic assumptions outlined in part one. Certainly, now that assumptions underlying logical inferentialism have been brought to light, these discussions themselves must be subjected to the closest scrutiny. But it would be a mistake, I think, to write the issue off as an irresolvable partisan dispute. Inferentialism has found its share of sympathizers in the classicist camp. And for good reason: the principles that underlie inferentialism capture the deeply rooted notion that meaning must be explained in terms of use. Anyone who dismisses inferentialism must therefore either deny the legitimacy of the meaning-is-use thesis altogether or he must present an alternative account. In light of its overarching significance, I would argue that the inferentialist position and its founding assumptions merit fair-minded and careful consideration.

Bibliography

R. E. Auxier and L. E. Hahn, editors. *The philosophy of Michael Dummett.* Open Court, Chicago, 2007.

N. Belnap. Tonk, plonk and plink. *Analysis*, 23:130–134, 1962.

S. Blackburn. *Spreading the word.* Clarendon Press, Oxford, 1984.

R. Bluhm and C. Nimtz, editors. *Selected papers contributed to the sections of GAP 5, fifth international congress of the society for analytical philosophy.* Mentis, Paderborn, 2004.

N. Bostock. *Intermediate logic.* Oxford University Press, Oxford, 1997.

R. Brandom. Truth and assertibility. *Journal of philosophy*, 73:137–149, 1976.

R. Brandom. *Making it explicit.* Harvard University Press, Cambridge, 1994.

R. Brandom. *Articulating reasons.* Harvard University Press, Cambridge, 2000.

R. Carnap. *Formalization of logic.* Harvard University Press, Cambridge, 1943.

R. Cook. Intuitionism reconsidered. In S. Shapiro, editor, *Handbook of the philosophy of logic and mathematics*, pages 387–411. Oxford University Press, Oxford, 2005.

N. Denyer. The principle of harmony. *Analysis*, 49:21–22, 1989.

K. Dŏsen. Logical constants as punctuation marks. In D. M. Gabbay, editor, *What is a logical system?*, pages 273–296. Clarendon Press, Oxford, 1994.

M. Dummett. *Frege: Philosophy of language.* Harvard University Press, Cambridge, 1973.

M. Dummett. Is logic empirical? In *Truth and other enigmas*, pages 269–289. Duckworth, London, 1976.

M. Dummett. *Elements of intuitionism*. Oxford University Press, Oxford, 1977.

M. Dummett. *Truth and other enigmas*. Duckworth, London, 1978.

M. Dummett. Language and truth. In *Seas of language*, pages 147–165. Oxford University Press, Oxford, 1993a.

M. Dummett. *The logical basis of metaphysics*. Harvard University Press, Cambridge, 1991.

M. Dummett. *Seas of language*. Oxford University Press, Oxford, 1993b.

M. Dummett. 'Yes', 'no' and 'can't say'. *Mind*, 111:289–296, 2002.

M. Dummett. Reply to dag prawitz. In R. E. Auxier and L. E. Hahn, editors, *The philosophy of Michael Dummett*, pages 482–489. Open Court, Chicago, 2007a.

M. Dummett. Reply to John Campbell. In R. E. Auxier and L. E. Hahn, editors, *The philosophy of Michael Dummett*, pages 301–313. Open Court, Chicago, 2007b.

J. Etchemendy. *The concept of logical consequence*. Harvard University Press, Cambridge, 1990.

J. Fenstad, editor. *Proceedings of the 2. Scandinavian logic symposium*. North-Holland, Amsterdam, 1971.

G. Frege. *Grundlagen der Arithmetik*. Meiner, Hamburg, 1884/1988.

D. M. Gabbay and F. Guenthner, editors. *Handbook of philosophical logic*, volume III. Reidel, Dordrecht, 1986.

D. M. Gabbay and F. Guenthner, editors. *Handbook of philosophical logic*, volume II. Kluwer, Dordrecht, 2001.

D. M. Gabbay and H. Wansing, editors. *What is negation?* Kluwer, Dordrecht, 1999.

P. Geach. Assertion. *The philosophical review*, 74:449–465, 1965.

G. Gentzen. Investigations into logical deduction. In M. Szabo, editor, *The collected papers of Gerhard Gentzen*, pages 68–128. North Holland, Amsterdam, 1934/1969.

G. Gentzen. The consistency of elementary number theory. In M. Szabo, editor, *The collected papers of Gerhard Gentzen*, pages 132–213. North Holland, Amsterdam, 1936/1969.

M. Gómez-Torrente. The problem of the logical constants. *The bulletin of symbolic logic*, 8:1–37, 2001.

M. Green. The status of supposition. *Noûs*, 34:376–399, 2000.

I. Hacking. What is logic? *Journal of philosophy*, 76:285–319, 1979.

P. Hájek, L. Valdés-Villanueva, and D. Westerståhl, editors. *Logic, methodology and philosophy of science: Proceedings of the twelfth international congress.* King's College Publications, London, 2005.

V. Halbach. Axiomatic theories of truth. *Stanford encyclopedia of philosophy*, 2005. URL `http://plato.stanford.edu/entries/truth-axiomatic/`.

G. Harman. *Change in view: Principles of reasoning.* M.I.T. Press, Cambridge, 1986.

J. Ketland. Deflationism and Tarski's paradise. *Mind*, 108:69–94, 1999.

W. Kneale. The province of logic. In H. D. Lewis, editor, *Contemporary British philosophy*, pages 237–261. George Allen and Unwin, London, 1956.

M. Kremer. Logic and meaning: The philosophical significance of the sequent calculus. *Mind*, 47:50–72, 1988.

N. Kurbis. *A pluralist justification of deduction.* PhD thesis, King's College London, London, 2007.

H. D. Lewis, editor. *Contemporary British philosophy.* George Allen and Unwin, London, 1956.

J. MacFarlane. *What does it mean to say that logic is formal?* PhD thesis, University of Pittsburgh, Pittsburgh, 2000.

J. MacFarlane. Frege, kant, and the logic in logicism. *The philosophical review*, 111: 25–65, 2002.

A. Miller, editor. *Essays for Crispin Wright: Logic, language and mathematics.* Oxford University Press, Oxford, forthcoming.

P. Milne. Classical harmony: Rules of inference and the meaning of the logical constants. *Synthese*, 100:49–94, 1994.

P. Milne. Harmony, purity, simplicity and a "seemingly magical fact". *Monist*, 85: 498–534, 2002.

S. Negri and J. von Plato. *Structural proof theory.* Cambridge University Press, Cambridge, 2001.

F. Paoli. Quine and Slater on paraconsistency and deviance. *Journal of philosophical logic*, 32:531–548, 2003.

D. Prawitz. *Natural deduction.* Dover, Mineola, NY, 1965/2006.

D. Prawitz. Ideas and results in proof theory. In J. Fenstad, editor, *Proceedings of the 2. Scandinavian logic symposium*, pages 237–309. North-Holland, Amsterdam, 1971.

D. Prawitz. On the idea of a general proof theory. *Synthese*, 27:63–77, 1974.

D. Prawitz. Meaning and proofs: On the conflict between classical and intuitionistic logic. *Theoria*, 43:2–40, 1977.

D. Prawitz. Proofs and the meaning and completeness of the logical constants. In J. Hintikka, editor, *Essays in mathematical and philosophical logic*, pages 25–39. Reidel, Boston, 1978.

D. Prawitz. Review of *The logical basis of metaphysics*, by Michael Dummett. *Mind*, 103:373–376, 1994.

D. Prawitz. Pragmatist and verificationist theories of meaning. In R. E. Auxier and L. E. Hahn, editors, *The philosophy of Michael Dummett*, pages 455–481. Open Court, Chicago, 2007.

A. Prior. The runabout inference ticket. *Analysis*, 21:38–39, 1960.

A. Prior. Conjunction and contonktion revisited. *Analysis*, 24:191–196, 1964.

H. Putnam. The logic of quantum mechanics. In *Mathematics, matter and method: Philosophical papers volume I*, pages 174–197. Cambridge University Press, Cambridge, 1975a.

H. Putnam. *Mathematics, matter and method: Philosophical papers volume I.* Cambridge University Press, Cambridge, 1975b.

W.V.O. Quine. *Mathematical logic.* Harvard University Press, Cambridge, 1940.

S. Read. *Relevant logic.* Blackwell, Oxford, 1988.

S. Read. Harmony and autonomy in classical logic. *Journal of philosophical logic*, 29:123–154, 2000.

S. Read. Identity and harmony. *Analysis*, 64:113–119, 2004.

S. Read. Harmony and necessity. In C. Dégremont, L. Kieff, and H. Rückert, editors, *Dialogues, logics and other strong things: Essays in honour of Shahid Rahman*, pages 285–303. College Publications, London, 2008.

G. Restall. *An introduction to substructural logics.* Routledge, London, 2000.

G. Restall. Multiple conclusions. In P. Hájek, L. Valdés-Villanueva, and D. Westerståhl, editors, *Logic, methodology and philosophy of science: Proceedings of the twelfth international congress*, pages 189–205. King's College Publications, London, 2005.

I. Rumfitt. 'Yes' and 'no'. *Mind*, 109:781–823, 2000.

I. Rumfitt. Unilateralism disarmed: A reply to Dummett and Gibbard. *Mind*, 111: 305–322, 2002.

I. Rumfitt. Knowledge by deduction. *Grazer philosophische Studien*, 77:61–84, 2008.

P. Schroeder-Heister. On the notion of assumption in logical systems. In R. Bluhm and C. Nimtz, editors, *Selected papers contributed to the sections of GAP 5, fifth international congress of the society for analytical philosophy.* Mentis, Paderborn, 2004.

S. Shapiro. Induction and indefinite extensibility: The Gödel sentence is true, but did someone change the subject? *Mind*, 107:597–624, 1998a.

S. Shapiro. Proof and truth: Through thick and thin. *Journal of philosophy*, 95: 493–521, 1998b.

S. Shapiro, editor. *Handbook of the philosophy of logic and mathematics*. Oxford University Press, Oxford, 2005.

D. Shoesmith and T. Smiley. *Multiple-conclusion logic*. Cambridge University Press, Cambridge, 1978.

T. Smiley. Rejection. *Analysis*, 56:1–9, 1996.

T. Smiley. Multiple-conclusion logic. *Routledge encyclopedia of philosophy*, 1998. URL `http://www.rep.routledge.com/article/Y046SECT2`.

F. Steinberger. Tennant on multiple conclusions. *Logique et analyse*, 2008.

F. Steinberger. Not so stable. *Analysis*, 69:655–661, 2009.

F. Steinberger. Inferentialism, structural rules and the truth about *ex falso*. unpublished.

G. Sundholm. Proof theory and meaning. In D. M. Gabbay and F. Guenthner, editors, *Handbook of philosophical logic*, volume III, pages 471–506. Reidel, Dordrecht, 1986.

G. Sundholm. Systems of deduction. In D.M. Gabbay and F. Guenthner, editors, *Handbook of philosophical logic*, volume II, pages 1–52. Kluwer, Dordrecht, 2001.

M. Szabo. *The collected papers of Gerhard Gentzen*. North-Holland, Amsterdam, 1969.

N. Tennant. *Natural logic*. Edinburgh University Press, Edinburgh, 1978.

N. Tennant. *Anti-realism and logic*. Oxford University Press, Oxford, 1987.

N. Tennant. The law of excluded middle is synthetic a priori, if valid. *Philosophical topics*, 24:205–229, 1996.

N. Tennant. *The taming of the true.* Oxford University Press, Oxford, 1997.

N. Tennant. Negation, absurdity and contrariety. In D. M. Gabbay and H. Wansing, editors, *Negation.* Kluwer, 2004a.

N. Tennant. A general theory of abstraction operators. *Philosophical quarterly*, 54: 105–136, 2004b.

N. Tennant. Rule-circularity and the justification of deduction. *Philosophical quarterly*, 55:625–648, 2005a.

N. Tennant. Relevance in reasoning. In S. Shapiro, editor, *Handbook of the philosophy of logic and mathematics*, pages 696–726. Oxford University Press, Oxford, 2005b.

N. Tennant. Inferentialism, logicism, harmony and a counterpoint. In A. Miller, editor, *Essays for Crispin Wright: Logic, language and mathematics.* Oxford University Press, Oxford, forthcoming.

A. Troelsta and H. Schwichtenberg. *Basic proof theory.* Cambridge University Press, Cambridge, 2000.

D. van Dalen. *Logic and structure.* Springer, Berlin, 1997.

P. van Inwagen. *The problem of evil.* Oxford University Press, Oxford, 2006.

J. von Plato. Gentzen's proof of normalization of natural deduction. *Bulletin of symbolic logic*, 14:240–257, 2008.

A. Weir. Classical harmony. *Notre Dame journal of formal logic*, 27:459–482, 1986.

Index