What harmony could and could not be[*]

Florian Steinberger

Abstract: The notion of harmony has played a pivotal role in a number of debates in the philosophy of logic. Yet there is little agreement as to how the requirement of harmony should be spelled out in detail or even what purpose it is to serve. Most if not all conceptions of harmony can already be found in Michael Dummett's seminal discussion of the matter in *The Logical Basis of Metaphysics*. Hence, if we wish to gain a better understanding of the notion of harmony, we do well to start here. Unfortunately, however, Dummett's discussion is not always easy to follow. The following is an attempt to disentangle the main strands of Dummett's treatment of harmony. The different variants of harmony as well as their interrelations are clarified and their individual shortcomings *qua* interpretations of harmony are demonstrated. Though no attempt is made to give a detailed alternative account of harmony here, it is hoped that our discussion will lay the ground for an adequate rigorous treatment of this central notion.

Keywords: Inferentialism, harmony, logical constants.

1. Harmony's meaning-theoretic basis

The idea of harmony has its source in Michael Dummett's conception of the general shape a theory of meaning should take. Let us assume with Dummett that a broadly use-theoretic approach is on the right track. A question we then face is whether an account of this form can aspire to systematicity. Dummett urges that it can. The key to the answer lies in what we might call Dummett's *two-aspect model of meaning*. According to this model, the core principles governing the use we make of a sentence fall into two main categories:

> 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 the hearer, including, in the case of assertion, what he is entitled to infer from it if he accepts it.
>
> [Dummett 1973: 396]

The meanings of sub-sentential expressions are explained in terms of the contribution they make to determining the two aspects of the assertoric use of the sentences in which they occur. The former category—let us call these features *I-principles*—corresponds to a broadly verificationist conception of meaning. The latter category of principles—call such features *E-principles*—represents what we can do with a sentence containing the said expression in virtue of having asserted it (including behavioural consequences). An emphasis on this feature of use is the hallmark of a pragmatist theory of meaning.

So far, so good. But what does this have to do with harmony? Dummett insists that both aspects of meaning—I-principles and E-principles—cannot properly be fixed independently of one another. For a linguistic practice to be in good working order there ought to be a 'certain consonance' [idem: 397] between both aspects of meaning. Why should this be the case? Well, assume someone introduces us to a novel expression that is to extend our linguistic practice (which we may assume to be functioning satisfactorily). Imagine, moreover, that the principles governing the use of our novel expression lack the desired balance.[1] Invoking the expression may then enable us to make inferences that we should not have

[*] I am most grateful to Arif Ahmed, Yoon Choi, Luca Incurvati, Julien Murzi and Michael Potter for helpful comments on earlier drafts of this paper. I also benefited from the excellent comments of two anonymous referees.

[1] Note that strictly speaking sub-sentential expressions do not, in and of themselves, have grounds or consequences; only sentences containing them do (perhaps relative to contexts). Dummett, following the context principle, takes it that a speaker's mastery of the meaning of a sub-sentential expression consists in knowledge of the grounds for asserting a certain representative range of sentences in which the said expression figures in canonical ways and of the consequences that follow from so asserting them. The molecularist Dummett insists, moreover, that any such representative range cannot extend to the whole of language. The notion of a meaning-determining representative

been in a position to make prior to the introduction of the expression. Dummett's much discussed example of the pejorative 'boche' may serve to illustrate the point at hand: in the presence of 'boche' we can mediately infer from someone's being German (the circumstances of the term's application) to certain of the person's character traits—that he is prone to cruelty (the consequences of its application), for example. We should not have been able to draw the same conclusion *im*mediately for want of a direct inferential link from the I-principles governing 'German' to those of 'prone to cruelty'. By creating such a link 'boche' in effect alters the meanings of 'German' (by adding to its E-principles) and of 'prone to violence' (by adding to its I-principles). Meaning-theoretically speaking, the expression 'boche' suffers from a type of disharmony in that its E-principles outstrip the associated I-principles. More generally, in the absence of a harmony requirement, any newly introduced expression regulated by disharmonious principles of use runs the risk of perturbing the extant linguistic practice.

For Dummett a practice that displays this type of disharmony—even if the practice is in use and hence learnable—is defective. Thus, there is a sense, according to him, in which a linguistic practice can be subject to criticism.[2] A language may be criticized if the conventions regulating the use of its constituent expressions are incoherent, i.e. its component conventions are such that the linguistic practice 'cannot be rendered comprehensible by any systematic ascription of content—and this means [that] it would be impossible to construct a meaning-theory that accorded with all features of its practice' [Dummett 1991: 241]. Harmony is therefore a prerequisite for the Dummettian project of constructing a theory of meaning: a theory that 'aims at a systematic means of ascribing *content* to the expressions and sentences of the language, in terms of which accepted modes of operating with it (including the rules of inference observed) can be justified' [idem].

Here is not the place to evaluate Dummett's larger meaning-theoretic enterprise. Suffice it to say that even someone sympathetic to its aims may nevertheless doubt that a notion of harmony can be carved out applicable to the whole of language—a formidable task, as Dummett himself concedes (see e.g. [Dummett 1973: 455n]). More importantly, it is not clear that a language conforming to Dummett's non-conservativeness constraint—his semi-formal correlate for this global notion of harmony—is workable. Not only does it seem plausible that such a constraint would be violated even by seemingly innocuous every-day expressions, conceptual progress in the sciences standardly (and indeed necessarily) involves the introduction of novel vocabulary that leads to non-conservative extensions of our theories. As Robert Brandom has pointed out, the fact that newly introduced theoretical expressions induce non-conservativeness arguably may merely show that such expressions have 'substantive content [...] that is not already implicit in the contents of the other concepts being employed' [Brandom 2000: 71].[3]

2. Harmony for the logical constants

Be that as it may, our concern in this paper is exclusively the concept of harmony insofar as it applies to the logical constants, setting aside the problem of formulating harmony for the remainder of language. Whatever misgivings one may have about Dummett's wider project, a strong case can be made for a logic-specific harmony requirement. The reason for this stems from the role logic plays in our assertoric practices. On the use-theoretic view sketched above, the meanings of non-logical sentences (sentences not containing any logical operators) are thought to be given by their I- and E-principles. Logic, in addition to the direct grounds for assertion given by the appropriate I-principles, offers indirect grounds for asserting non-logical sentences: we may assert a non-logical sentence if it can be correctly deduced from a set of accepted premises. But for these indirect deductive routes to assertibility to be not only legitimate but to have the unassailable reliability we require of logical inference, our logical modes of inference must

range raises a number of important questions for general theories of meaning and for accounts of the meanings of the logical constants in particular. Unfortunately, these questions are beyond the scope of this paper.

[2] '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 for revision, than our social, political, or economic practice' [Dummett 1991: 214-15].

[3] I am grateful to an anonymous referee for stressing this point.

respect the conditions under which the (direct) assertion of non-logical sentences is justified. That is, logical inference alone may not license the assertion of non-logical sentences that we should not have been in a position to assert directly (at least in principle). Let call us call this the *principle of innocence*: it should not be possible, solely by engaging in deductive logical reasoning, to discover hitherto unknown (atomic) truths that we would have been incapable of discovering independently of logic. An alternative way of making the same point is to say that our principles of logic ought not tinker with the meanings of non-logical expressions. It is the principle of innocence that ensures the applicability of logic to our talk about the world.[4] How can we make sure that innocence obtains? This is where harmony comes in. The primary purpose of harmony is precisely to secure the innocence of logic. In short, it falls to harmony to prevent our deductive practices from going off the rails.

A moment's reflection reveals that harmony is not only an adequate measure, but that it seems entirely natural that any measure designed to guarantee the holding of the requirement of innocence should take the form of a harmony requirement. After all, our aim is to ensure that the meanings of the logical constants are fixed in such a way as not to perturb the non-logical regions of language. The best way to do this (at a local level) is by requiring that the introduction and elimination rules that govern the meanings of logical constants be exactly commensurate in strength. Why? Well, because when such an equilibrium between I-rules and E-rules obtains, we can rest assured that our deductive practices will not, as it were, *create* novel grounds for asserting non-logical sentences (as in the case, for example, of *tonk*). The requirement of harmony thus seems to be an eminently reasonable and natural safeguard for the principle of innocence.

The considerations of the preceding two paragraphs have already given us a glimpse of what harmony amounts to in the case of the logical constants: a pair of inference rules is harmonious just in case it displays a certain kind of equilibrium. We can make this somewhat more precise. Let $ be a logical constant. (For simplicity we may assume $ to be a binary connective.) Harmony obtains when the canonical grounds for asserting a sentence containing $ as its principal operator are appropriately counter-balanced by the deductive consequences of accepting a sentence of this form. Moreover, that nothing more and nothing less may be deduced from an assertion of *A*$*B* via $-E than can already be deduced from the premises of the corresponding I-rules. Put yet another way, the rules for $ should be balanced, in the sense that the E-rules ought to exploit *all* and *only* the inferential powers that the I-rules have bestowed upon it.

Let us refer to this refined informal notion as *general harmony*. General harmony is a relational property of pairs of introduction and elimination rules. By extension, a logical constant can be said to be harmonious or not depending on whether there are harmonious rules of inference governing it. Though insufficiently rigorous for theoretical purposes, general harmony is still precise enough to serve as our target notion, setting the standard against which formal accounts of harmony are evaluated. As I will try to show, a surprising number of interesting conclusions concerning the nature of harmony can be drawn simply on the basis of this informal characterization.[5]

It follows from our characterization of general harmony that disharmony can arise in one of two ways: a set of E-rules could be too weak or too strong relative to their corresponding I-rules. In the former case, our E-rules bar us from drawing all of the deductive conclusions that the premises of the corresponding I-rule would have licensed. Its unduly restrictive E-rules deprive us of certain inferences and so prevent us from making assertions to which we are in fact entitled. Call this type of disharmony *E-weak disharmony*.[6]

In the latter case, our E-rules are excessively permissive in the deductive inferences they allow. While E-weak disharmony deprives us of certain deductive inferences, this form of disharmony, *E-strong*

---

[4] Dummett's examples of counting and addition, and of Euler's solution to the Königsberg Bridges Problem serve to make this very point [Dummett 1991: 219].

[5] Alternative accounts of harmony like those of Neil Tennant (see [Tennant Forthcoming] for the latest version and [Steinberger 2009] for critical discussion) and Stephen Read [2010] can both be seen to be attempts to codify our notion of general harmony.

[6] Saying that E-weakly disharmonious rules deprive us of certain inferences does not mean, of course, that the assertions that we are prevented from making in this way cannot be arrived at by other means.

disharmony, may illicitly produce new information: the elimination rules of $\$$ make available inferences that we would perhaps not have been in a position to make had we not introduced $\$$ via one of the I-rules. For future reference, let us note the two types of disharmony as follows:

- **E-strong disharmony**: the E-principles are unduly permissive (relative to the corresponding I-principles);
- **E-weak disharmony**: the E-principles are unduly prohibitive (relative to the corresponding I-principles).

Accordingly, Dummett himself frequently distinguishes two distinct notions. Harmony *tout court* is the requirement that a logical expression's E-rules not outstrip the corresponding I-rules, but leaves open the possibility that E-rules be too weak fully to exploit the associated I-rules. In other words, harmony *tout court* guards against what we have called E-strong disharmony, but not E-weak disharmony. In order to block E-weak disharmony as well, we must ensure that *stability* obtains [Dummett 1991: 287]. This will be the case whenever the E-rules fully exhaust the grounds for asserting a sentence with the expression in question as stated by the corresponding I-rules.

A comment is in order here. Consider the connective § whose introduction rule matches that of conjunction and whose elimination rule is just the standard rule for $\vee$-elimination. § appears to be E-weakly disharmonious because §-E is too weak relative to §-I. But what was our justification for this diagnosis? After all, who is to say that it is the E-rule that is at fault? Could one not, by the same token, privilege the E-rule and instead reject the I-rule on account of its excessive strength?[7] From this angle, §'s defect would be rather more aptly described as *I-strong* (as opposed to E-weak) disharmony. An analogous question could be asked with respect to the parallel case of I-weak versus E-strong disharmony.

Whether one chooses to treat a given disharmonious pair of rules as a case of I-type or as a case of E-type disharmony (i.e. whether one blames I-rules or E-rules for the breach of harmony) depends on one's stance on the question which of the two—I-rules or E-rules—are the primary determinants of the meaning of the constant in question. A number of authors have followed Gerhard Gentzen in privileging introduction rules. I-rules, according to this view, fix the meaning of the logical constant; E-rules, being functionally dependent on the I-rules, serve to explicate that meaning. However, advocates of this 'I-rules first' view rarely supply arguments for their position. But why should not E-rules in some or all cases determine the meanings of the logical operators at least in part? At least in the case of the conditional [Rumfitt 2000: 790], the universal quantifier [Dummett 1991: 275] and the disjunction operator [Tennant 1987: 90], it has been argued that it is in some sense more natural to accord meaning-theoretic primacy to elimination rules. In the following, we will remain neutral on the question of the meaning-theoretic priority of I- and E-rules. In order to allow for the possibility that the 'E-rules first' view may be correct (at least in some cases), we supplement the above list of kinds of (E-type) disharmony with the two further variants of I-type disharmony:

- **I-strong disharmony**: the I-principles are unduly permissive (relative to the corresponding E-principles);
- **I-weak disharmony**: the I-principles are unduly prohibitive (relative to the corresponding E-principles).

That being said, in the last analysis it will not much matter how one chooses to classify a given pair of disharmonious inference rules. As Dummett repeatedly stresses, both I-rules and E-rules are required in order fully to determine the meaning of a logical constant; irrespective of whether one takes a constant's I-rules or E-rules as one's point of departure, either set of rules (provided that they are permissible) ought to determine its harmonious counterpart uniquely (see e.g. [Dummett 1983: 142]).

Having thus characterized the central notion of general harmony, let us now return to Dummett's project. Despite harmony's larger meaning-theoretic significance for Dummett, a good deal of his

[7] I am grateful to an anonymous referee for pressing me on this point.

treatment of it is devoted to the quest for a logic-specific version of harmony. In part this can be explained by the central role the logical constants play in Dummett's meaning-theoretic enterprise: only if a satisfactory account of harmony can be given for the logical constants, does Dummett's wider project of formulating a language-widely applicable notion of harmony stand any chance of getting off the ground. Also, he takes the task of formulating a harmony constraint to be more tractable in the restricted realm of the logical operators [Dummett 1991: 215]. But in attempting to elaborate a specifically logical version of harmony, Dummett is also driven by intrinsically logical aims. It is here that Dummett's project aligns with the goals standardly pursued in the literature on harmony. We can identify four principle aims. First, as we have seen, harmony is deployed as a safeguard for the requirement of innocence.[8] Second, harmony is at the heart of the inferentialist approach to the meanings of the logical constants—the position according to which, roughly, the meanings of the logical constants are determined by the rules of inference they obey—where it serves as an adequacy condition for logical constanthood, warding off unruly pseudo-constants like *tonk*.[9] Third, harmony lays the ground for an attack on classical logic developed in Chapters 12 and 13 of *The Logical Basis of Metaphysics*.[10] Finally, Dummett considers (and rejects) the idea that harmony might serve as a criterion of logicality [Dummett 1991: 272].[11]

3. Conservativeness and its discontents

The notion of harmony in its logic-specific form thus serves several purposes for Dummett. In particular, as we have just seen, it serves as an adequacy condition for logical constanthood, and it serves as a *desideratum* in disputes between classical and constructivist logicians. Interestingly, Dummett frequently employs Nuel Belnap's notion of a conservative extension for both purposes. The aim of the current section is to elucidate the notion of conservativeness and to assess its suitability as a technical correlate for the notion of harmony.

Belnap deploys the notion of a conservative extension in his well-known response [Belnap 1962] to Prior's attack on logical inferentialism [Prior 1960]—the view that what determines the meanings of the logical constants are the rules of inference that govern their deductive behaviour. *Au fond* what is objectionable about Prior's connective *tonk*, according to Belnap, is that it perturbs our existing deductive practice. Indeed its effect on any consistent existing practice is so deleterious that the entire system is reduced to triviality. What is needed, therefore, is a principled constraint to the effect that the introduction of a new logical operator not interfere with our established logical practice.

Conservative extensions are standardly introduced in the context of formal theories. Let T and T' be theories formulated in the languages L and L' respectively, such that $T \subseteq T'$ and $L \subseteq L'$. T' can then be said to extend T conservatively if T' allows us to prove no hitherto unprovable sentences expressible in the restricted language L; any newly provable theorems must be sentences involving new vocabulary. In other words, for every sentence $A \in L$, if $T' \vdash A$, then $T \vdash A$. Call this *theoretical conservativeness*.

In the context of systems of logic, conservativeness amounts 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 the languages L and L' associated with the deductive systems S and S' respectively, where $L' = L \cup \{\$\}$ for a newly introduced logical operator \$. Systematic conservativeness then requires that if $\Gamma : A$ is a sequent in L

[8] Indeed, Dummett takes this to be the central task of the notion of harmony, calling it 'the *point* of the requirement of harmony' [Dummett 1991: 290, Dummett's emphasis].
[9] See also e.g. [Belnap 1962], [Dummett1973] and [Tennant 1987].
[10] Other authors who have relied on harmony as a *desideratum* in disputes over the validity of the Law of Excluded Middle include [Dummett 1991], [Prawitz 1977] and [Tennant 1997].
[11] For an interesting elaboration and defence of the view that logicality should be equated with expressibility in terms of harmonious rules of inference see [Tennant 1997].

(and so does not involve \$), we have $\Gamma|\text{-}_{S'}A$ only if $\Gamma|\text{-}_{S}A$.[12]

The difference between theoretical conservativeness and Belnap's 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-tautologous 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.[13]

But Dummett's contention that systematic conservativeness aligns with general harmony cannot be quite right. General harmony, it may be hoped, implies conservativeness in the sense that the introduction of a logical operator obeying generally harmonious I- and E-rules into a deductive practice that is in good working order (i.e. a practice featuring only logical operators that are likewise governed by generally harmonious pairs of inference rules) will result in a conservative extension (see section 7 for further discussion of this point). 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 *E-weak disharmony*. Conservativeness guards against excessively permissive principles of inference; it offers no protection against excessively restrictive ones.

This thought may lead one to suspect that the introduction of a logical constant \$ governed by disproportionately weak E-rules is not only compatible with conservativeness, but that it even entails conservativeness. For suppose the meaning of \$ is given by a pair of rules that displays E-weak disharmony. In that case the I-rules will invest \$ with a meaning that the E-rules are too weak fully to exploit. And that, one might think, implies that the E-rules associated with \$ license fewer inferences not more than we are already committed to if we are in a position legitimately to assert a sentence containing \$ as its principal connective. It would then seem to follow that the addition of a logical expression governed by E-weakly disharmonious inference rules 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 E-weak and E-strong disharmony are mutually exclusive. But this is not so. The rules governing *tonk*, for instance, display E-weak disharmony (*A*, the premise needed to obtain *AtonkB* via *tonk*-I, cannot be recovered from *AtonkB* by applying *tonk*-E) *and* E-strong disharmony.[14]

Interestingly, it does not entail conservativeness *even if* we assume the constant *not* to be E-strongly disharmonious. There are cases where adding a perfectly harmonious logical constant to certain deductive systems that contain E-weakly disharmonious (though not E-strongly disharmonious) expressions results in a non-conservative extension (see section 5 for an example of this kind).

So far, then, we have seen that the notion of conservativeness failed to guard against the threat of E-weak disharmony. But there are further, more fundamental difficulties with the notion of conservativeness. As Belnap points out, whether a given constant satisfies the demand for conservativeness depends on whether the rules we lay down for it are consistent with our prior assumptions; a constant, even if it is acceptable, is always acceptable only 'relative to our characterization of deducibility' [Belnap 1962: 133]. But this means that one and the same logical constant may result in conservative extensions of some systems, but lead to non-conservativeness when adjoined to others.[15]

---

[12] I shall explicitly distinguish the two notions of conservativeness—theoretic and systematic—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.

[13] For Dummett there is in fact a further type of conservativeness—we might call it *linguistic conservativeness*—which amounts to an adaptation of the notion of conservativeness to language in general. Linguistic conservativeness is proposed as a way of cashing out the global notion of harmony [Dummett 1991: 219]. It might be considered an attractive feature of Dummett's approach that (essentially) one and the same notion (i.e. conservativeness) applies, albeit in context-specific forms, to language as a whole and to the restricted fragment of the logical constants.

[14] I thank an anonymous referee for simplifying my original example.

[15] Classical negation, for example, can be conservatively added to the system $\{\wedge, \vee\}$, but leads to non-conservativeness when adjoined to a system also containing $\supset$.

There would be nothing objectionable about *tonk* so long as we are careful to introduce it only into contexts where it does not do any damage. Hence, if we were to identify harmony with conservativeness, harmony would become a property ascribable to pairs (*S*, ($-I, $-E)) where *S* is a base system and ($-I, $-E) is a pair of inference rules governing the logical constant that conservatively extends the language associated with *S*. But this seems wrong on two counts. For one, whether a given logical constant is harmonious ought to depend on the rules that determine its meaning, not on the context within which it is deployed. Second, conservativeness simply is not a relational property of the right kind: general harmony, we have said, is a property of pairs of inference rules; conservativeness is a property of systems and pairs of inference rules. Unlike the notion of conservativeness, the notion of general harmony we are after is a *local*, not a *global* property. Consequently, systematic conservativeness could not possibly offer a suitable formalization of the intuitive notion of general harmony.

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. We will return to the role conservativeness plays in Dummett's account of harmony in sections 6 and 7.

4. Levelling local peaks and normalizability

In the course of his discussion of harmony as it applies to the logical constants, Dummett offers three distinct requirements: total harmony, intrinsic harmony and stability. *Total harmony* is simply another name for the requirement of conservative extendibility, which, as we have just seen, has proved to be ill-suited for our purposes. Dummett's notion of *intrinsic harmony*, by contrast, is 'a property solely of the rules governing the logical constant in question' [Dummett 1991: 250] and so *prima facie* appears to be a more promising candidate. Intrinsic harmony is based on what Dag Prawitz called the *inversion principle*. In Prawitz's words, 'an elimination rule is, in a sense, the inverse of the corresponding introduction rule: by an application of an elimination rule one essentially only restores what had already been established if the major premiss of the application was inferred by an application of an introduction rule' [Prawitz 1965: 33].

If an elimination rule is really just a device for 'undoing' a primitive inferential move effected by an application of an introduction rule, it should not be possible, simply by introducing and subsequently eliminating a logical constant, to arrive at new conclusions about the world. Insofar as intrinsic harmony prevents the creation of new information it will ensure that constants enjoying this property will not disrupt an extant practice into which they are introduced, or so one would expect.

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 enabling us to *level* local peaks or to *eliminate* them. 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 [Dummett 1991: 250].

Let us briefly pause to illustrate intrinsic harmony with the aid of a standard example. Consider the case of disjunction with its familiar introduction rules

$$\frac{\begin{matrix}\Gamma\\ A\end{matrix}}{A \vee B}$$

(similarly for the case where *B* is the premise) and the corresponding elimination rule

$$\frac{\begin{array}{ccc}\Gamma_0 & \Gamma_1, A & \Gamma_2, B\\ A\vee B & C & C\end{array}}{C}$$

Suppose we have a local peak featuring ∨ of the following form

$$\frac{\begin{array}{ccc}\Gamma_0 & & \\ \Pi_0 & \Gamma_1, A & \Gamma_2, B\\ \underline{A} & \Pi_1 & \Pi_2\\ A\vee B & C & C\end{array}}{C}$$

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)

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

Here $\Delta = \Gamma_1 \cup \Gamma_2$.[16] The rules for ∨ thus 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.[17]

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 such a way as to create a maximum. This standardly requires so called *permutative reductions* that allow us to rearrange the order of application of the inference rules involved in the proof.[18] 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

---

[16] It is worth noting that the application of levelling procedures will in many cases yield a logically stronger result. In our example this will be the case when only a subset of the assumptions $\Gamma_1 \cup \Gamma_2$ is required to establish the result, i.e. when Δ can be chosen such that $\Delta \subset \Gamma_1 \cup \Gamma_2$. I thank an anonymous referee for pointing out an inaccuracy in an earlier formulation.

[17] This formulation holds for intuitionistic logic. Things are not quite as neat in classical logic, see e.g. [Prawitz 1965: 42].

[18] I am here borrowing Dummett's terminology. For more on permutative reductions and other types of procedures required for normalization, see [Dummett 1977: 112].

auxiliary procedures just mentioned) can be applied to it. A system is *normalizable* if any proof within it can be converted into normal form.

Of chief importance for our purposes is the observation that normalizability is *a property of a proof system as a whole*. Whether a system normalizes depends on the specific configuration of the inference rules it contains, thus making normalizability a *global* property of a system. Levelling, on the other hand, is a *structural feature* of pairs of introduction and elimination rules; unlike the (global) property of admitting a permutative reduction procedure, the property of possessing a levelling procedure is a *local* property of a pair of rules. The question then is how the notion of intrinsic harmony fares as a potential interpretation of harmony. In the following section we discuss an example of Dummett's involving the quantum-logical disjunction operator that demonstrates the insufficiency of the formal notion of intrinsic harmony as an explication of the pre-formal notion of general harmony.

5. What's wrong with quantum-or?

Central to Dummett's example is the restricted disjunction elimination rules favoured by certain quantum logicians. The quantum-disjunction elimination rule distinguishes itself from the standard disjunction elimination rule by disallowing collateral hypotheses in the minor premises (i.e. Γ_1 and Γ_2 in the rule schema above must be empty). Let us denote the quantum-logical disjunction operator whose meaning is given by the usual introduction rules and the restricted elimination rule by 'ü'. Note that since the procedure for eliminating local peaks we gave for standard disjunctions in the previous section applies equally in the case of ü, the rules for ü are intrinsically harmonious by Dummett's standards. Suppose we situate ourselves in a system containing only, say, ∧ and ü. Dummett now observes that the addition of ∨ with its usual unrestricted elimination rule collapses quantum-logical disjunction into standard disjunction, since one can use ü-E and ∨-I to deduce $A\vee B$ from AüB:

$$\frac{A\text{ü}B \qquad \begin{array}{c}[A]\\ \hline A\vee B\end{array} \qquad \begin{array}{c}[B]\\ \hline A\vee B\end{array}}{A\vee B}$$

It follows that any grounds sufficient to warrant the assertion of AüB will be *ipso facto* sufficient to assert $A\vee B$, and will therefore justify the application of the unrestricted elimination rule. The fact that ü collapses into ∨ indicates that the rules governing quantum-logical disjunction, though intrinsically harmonious, do not succeed in conferring a stable meaning on ü. What this shows is that intrinsic harmony is insufficient for general harmony. In a sense, this should not come to us as a surprise. After all, levelling procedures guarantee only that operators cannot—locally—create new facts about the world; that is, they block E-strong disharmony by requiring that anything deducible from a statement containing the operator in question in a dominant position could have already been deduced directly on the basis only of the premises of the corresponding introduction rules. Yet intrinsic harmony on its own offers no protection against the opposite vice, E-weak disharmony. And it is precisely on these grounds that the weakened ü-elimination rule may be criticized; the ü-elimination rule is too weak fully to exploit the deductive powers that the introduction rules have bestowed upon ü.

What is striking about the example we are considering is that it demonstrates how an E-weakly disharmonious connective, ü, can destabilize a system. Even a perfectly well-behaved logical constant like the standard disjunction operator can, when introduced into a defective base system {∧, ü}, lead to undesirable consequences. In our case, the system {∧, ü, ∨} is a *non*-conservative extension of the base system: in particular, the law of distributivity $A\wedge(B\text{ü}C)\vdash(A\wedge B)\text{ü}(A\wedge C)$—characteristically invalid in quantum logic—becomes readily derivable for ü once ∨ is introduced.

What is more, the interplay of ∨ and ü may lead to irreducible plateaux. Assume that we were to extend the above proof by adding an application of ∨-elimination. We would then have created a plateau: the instance of ü-elimination following the application of the ∨-introduction rule delays the subsequent ∨-elimination. However, if we now try to apply our permutative reduction procedure in order to reduce the plateau to a local peak, we arrive at the following:

$$\dfrac{A\text{ü}B \qquad \dfrac{\dfrac{[A]}{A\vee B} \quad \dfrac{\Gamma_1, [A]}{C} \quad \dfrac{\Gamma_2, [B]}{C}}{C} \qquad \dfrac{\dfrac{[B]}{A\vee B} \quad \dfrac{\Gamma_3, [A]}{C} \quad \dfrac{\Gamma_4, [B]}{C}}{C}}{C}$$

Here the final application of ü-elimination is not in general permissible—it is legitimate only in special cases where $\Gamma_{1\text{-}4}$ all happen to be empty. Therefore, in our system the application of reduction procedures may lead us from genuine deductions to ill-formed ones. What this means is that plateaux of this kind are not eliminable and hence that the system {∧, ü, ∨} is not normalizable (whereas the system {∧, ü} is).

Summing up, Dummett has produced a system composed exclusively of intrinsically harmonious pairs of inference rules that is nonetheless not normalizable and with respect to which ∨ does not display total harmony. The lesson to be learned here is that intrinsic harmony is an inadequate characterization of general harmony on its own and therefore requires strengthening. What is needed is an additional provision for ruling out E-weak disharmony [Dummett 1991: 287]. It is to this end that Dummett advances the notion of *stability*, intrinsic harmony's symmetric counterpart (unfortunately, however, without ever fully developing it). Jointly, intrinsic harmony and stability yield an adequate formalization of our intuitive notion of general harmony, or so one would hope. Let us, in this hopeful vein, refer to the conjunction of levelling procedures and a stability requirement designed to ward off E-weak disharmony as *ideal* harmony. I will make no attempt to devise a notion of stability here. Instead we will try to gain a better understanding of the role the global notions of total harmony and of normalizability play in Dummett's account. What is Dummett's reason for invoking them and is he justified in doing so? Also, we must ask how total harmony (i.e. harmony-as-conservative extendibility) and normalizability relate to ideal harmony and to each other. What is Dummett's motivation for introducing both local and global forms of harmony requirements in the first place?

This task takes on a particular urgency in the light of the somewhat perplexing course Dummett's discussion takes. Given its pivotal role in the elaboration of the notion of harmony, one would expect 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. He tells us that '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 that fragment' [Dummett 1991: 290, Dummett's emphasis]. Ideal harmony, 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 it is the latter—total harmony—that matters in the end. Dummett, after wavering on the issue, thus comes down on the wrong side and privileges total harmony over ideal harmony.

So what leads Dummett to prioritize total harmony? The idea underlying the demand for total harmony appears to be the intuition that the meaning of a constant should be fully determined by the rules of inference it obeys; the presence or absence of further connectives in the same system should not affect that meaning in any way. In particular, then, the addition of new logical vocabulary should not make available new theorems involving only the old vocabulary. For if it did, this would imply that the meanings of the old vocabulary had not been fully determined by their inference rules after all; i.e. that there had been ingredients of their meanings that had only been 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 one operator at a time. Put another way, we are assuming that the inference rules associated with a particular connective need mention only that and no other connective. This is the content of the *principle of separability*.

Now, there may be some intuitive plausibility to the principle of separability. Should it not be possible to learn the meanings of the logical constants one at a time in whatever order one wishes? Yet, as Peter Milne [2002] has pointed out, one searches in vain for arguments in favour of separability that go beyond such rather vague plausibility considerations even in the writings of authors who put considerable weight on the principle. Dummett himself seems conflicted on the issue. On the one hand, he explicitly singles out the logical constants 'as a uniquely simple case' of expressions that can be learned quite independently of other expressions in the same class [Dummett 1991: 223]. On the other hand, commenting on the idea that all introduction rules should be pure like those of Gentzen's original natural deduction formulation, Dummett makes the following claim:

> 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 we could so order the logical constants that the understanding of each depended only on the understanding of those preceding it in the ordering.
> [Dummett 1991: 257]

Let us call this more relaxed restriction, the requirement of *well-foundedness*.[19]

We make no attempt here to settle the question as to which kinds of dependencies among the meanings of the logical constants are permissible. What we can say, is that molecularism in no way entails separability or even well-foundedness. As far as language in general is concerned, the molecularist allows for dependence relations among expressions: the understanding of one expression may necessitate a prior grasp of the meanings of several others. Even semantic clusters—sets of expressions where the understanding of any one expression requires an understanding of all the others in the set—are admissible. What the molecularist opposes is the holist notion that all of these clusters collapse into one all-encompassing master cluster, language as a whole.[20] So far, then, no reason has been given for why the class of logical operators should not itself constitute such a cluster. After all, why should the set of the logical constants not be comparable to other semantic clusters the meanings of the members of which must be learned in one fell swoop? In the absence of an argument for separability or at least well-foundedness, the demand for total harmony is unjustified.[21] All the molecularist could reasonably require is that the fragment of the logical constants be separable *as a whole*; that is, that the logical fragment form a stable isolable semantic unit that does not interfere with the meanings of non-logical expressions.[22]

If the primary aim of a logic-specific harmony requirement is to ensure the *global* well-functioning of the logical fragment within the wider reaches of language in this sense (*and* if we *were* to follow Dummett in allowing that such a constraint may be global), then it seems that it is the property of *normalizability*, not that of conservativeness that would be our best bet. Normalizability ensures that no

---

[19] The reason Dummett settles for well-foundedness, one may reasonably suspect, stems from the perceived need to accommodate the intuitionistic rules for negation, which he takes to flout the principle of separability because they invoke the sign for absurdity '⊥', which Dummett takes to be as a logical constant endowed with its own meaning-giving rules [Dummett 1991: 295]. It should be noted, however, that ⊥ need not (and indeed *should* not, in my view) be conceived of in this way. Tennant has shown us that ⊥ is better conceived of as a dispensable 'structural punctuation marker within deductions', rather than as a logical or propositional constant [Tennant 1999: 199]. On Tennant's view, then, the rules for intuitionistic negation fully conform to the requirement of separability.

[20] Cf. Tennant's discussion of this point: 'what is crucial for the *molecularist* is whether, as we trace along, the 'covering umbrella' of ancestral concepts reached from any given concept does not mushroom out, well-foundedly or not, in such a way as eventually to take in as basis the basis of the whole conceptual system' [Tennant 1987: 37, Tennant's emphasis].

[21] The argument given here is not intended as an objection to either molecularism or the requirement of well-foundedness both of which are, in my opinion, attractive positions. My claim is simply that molecularism about meaning alone is insufficient to motivate the requirement of well-foundedness for the logical constants (and hence *a fortiori* insufficient to motivate the requirement of separability).

[22] Nothing excludes the possibility in principle that the logical fragment could, as it were, semantically 'bleed into' the non-logical regions of language, but that stability is achieved at a higher level. But while such a configuration cannot be ruled out, it does seem unlikely: unlikely that there should be such a fragment encompassing the logical fragment and unlikely that such a language should enjoy the surveyability required by a Dummettian theory of meaning.

logical operator in the system creates content; any conclusion reached by invoking a logical operator not itself occurring in the conclusion (or in the premises) can be shown to be deducible without appeal to the said operator. It guarantees, therefore, that there can be no semantic spill-over from the logical to the non-logical regions of language. In other words, it guarantees the requirement of innocence.

However, the choice does not present itself in the same way for Dummett. For Dummett claims that total harmony is entailed by normalizability and hence that we incur a commitment to total harmony already by adopting the requirement of normalizability. But this, I would argue, is not so. Before we can assess Dummett's claim that normalizability entails total harmony (i.e. systematic conservativeness) we must ensure that the two properties in question are properties of the same type: properties of a system. To ensure comparability, we must slightly generalize the requirement of conservativeness, making it into the property we might call *full conservativeness*: A system *S* is *fully conservative* if and only if for each logical constant \$ occurring in it, *S* is conservative over *S*-{\$}. Dummett's claim can then be put as follows: for any system of logic *S*, if *S* is a normalizable system of logic, then *S* is fully conservative.

6. Normalizability and conservativeness

Dummett offers the following argument for his claim. Suppose a system *S* is normalizable. Take any constant \$ of *S*: 'if we have a proof whose final sequent does not contain \$, any sentence [occurring in the proof] containing \$ 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 1999: 250]. Therefore, it may seem that taking any constant \$ of *S* and adding it to *S*-{\$} 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: $((A\supset B)\supset A)\supset A$.[23]

Consider the fragment of classical logic $S=\{\wedge, \supset\}$ with the familiar introduction and elimination rules for these connectives and with the rule of *ex falso sequitur quodlibet* (*EFQ*): $\bot/A$. The normalizability of *S* is obvious. In the absence of the elimination rules for ∨ and ∃ that necessitate additional reduction procedures, all detours must be instances of local maxima and so we can cheerfully level them away. As for *EFQ*, we can show that any derivation in which *EFQ* is used to infer a complex formula can be reduced to one in which all applications lead to atomic conclusions. This is a routine procedure (see e.g. [van Dalen 1997: 207]).

Now take $S' = S\cup\{\neg\}$, which is obtained by adding introduction and elimination rules for ¬,

$$\begin{array}{c}\Gamma, [A]\\ \vdots \\ \underline{\bot}\\ \neg A\end{array} \qquad\qquad \frac{A \qquad \neg A}{\bot}$$

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

$$\begin{array}{c}\Gamma, [\neg A]\\ \underline{\bot}\\ A\end{array}$$

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 ¬-I and ¬-E rules and by showing that all applications of *CRAA* can be reduced to instances where the conclusions are atomic. To

[23] 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.

start with the levelling procedure for ¬, consider a ¬-maximum:

$$\dfrac{\begin{array}{c}\Gamma_0\\ \Pi_0\\ A\end{array} \qquad \dfrac{\begin{array}{c}\Gamma_1, [A]\\ \Pi_1\\ \bot\end{array}}{\neg A}}{\bot}$$

Such a local peak can be levelled as follows:

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

The demonstration that all applications of *CRAA* can be reduced to applications on atomic formulas is again routine (see e.g. [Prawitz 1965: 40]). 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 ∧-maximum;
2. a local peak involving a ⊃-maximum;
3. a local peak involving a ¬-maximum;
4. the result of a non-atomic application of *EFQ*.
5. the result of a non-atomic application of *CRAA*.

Cases 1.-3. are taken care of by corresponding levelling procedures. 4. and 5. are adequately dealt with by the reduction procedures for *EFQ* and *CRAA* respectively. 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 straightforwardly proved in *S*'. This disproves Dummett's claim.

The reason why 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 (in the classical case) is not sufficient for full conservativeness, and so refuted Dummett's claim.

7. Ideal harmony and conservativeness

In the previous section we have followed Dummett in assuming that harmony is best represented by a global formal constraint. But this we did only 'for the sake of the argument' in order to disprove Dummett's claim that normalizability entails total harmony. But it is far from clear why global is the way to go here. After all, as we discovered in section 3, the question whether the desired balance between I- and E-principles 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. For Dummett the chief motivation behind the notion of harmony, its *raison d'être*, is to ensure the good functioning of language as a whole. If we are right, therefore, to take general harmony as our target notion, our task must be to formulate *a local principle in such a way as to avert global disaster*. In practice this amounts to showing that the formal correlate of general harmony, ideal harmony, entails normalizability in the sense that any system composed exclusively of ideally harmonious logical constants will be normalizable. Alas, in the absence of a worked out notion of ideal harmony, it will be difficult to assess.

Dummett himself also holds that ideal harmony should ensure the good *internal* functioning of the logical fragment, by which he means that total harmony should reign. And we too have mentioned that conservativeness is generally taken to be a necessary condition for harmony. While it is true that Dummett's quantum-logical excursion showed us that intrinsic harmony was no sufficient condition for total harmony, Dummett conjectures that adding the extra local requirement of stability *does* give us a sufficient condition for systematic conservativeness, i.e. that ideal harmony entails total harmony [Dummett 1991: 290]. The conjecture amounts to the claim that for all systems *S*, if *S* is composed exclusively of ideally harmonious rules of inference and a new constant $ also governed by stable rules is added to *S* to form *S*', then *S*' is systematically conservative over *S*. This amounts to the claim that any system composed of ideally harmonious operators is fully conservative. Let us dub this claim *Dummett's conjecture* for future reference.

One may wonder whether this conjecture is not already disproved by our demonstration in the previous section? 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 grade. Therefore, Dummett's conjecture that ubiquitous stability (along with intrinsic harmony) implies systematic conservativeness (total harmony) is not necessarily imperilled.

But this is not the only reason why one may have doubts about Dummett's conjecture. Both [Prawitz 1994: 374] and [Read 2000: 127] attack Dummett on this point, arguing that a ready-made counterexample can be found in the 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.[24] Take the following form of the *T*-schema given by the following straightforward rules:[25] $[A]/Tr(\ulcorner A \urcorner)$ and $Tr(\ulcorner A \urcorner)/A$. Prawitz and Read assume that no matter what the exact details of our account of ideal harmony are, if any rules are to qualify, then the rules for the truth predicate (as just presented) do. Hence, if the truth predicate can be found to induce non-conservativeness when adjoined to a well-functioning system, Dummett's conjecture will be refuted.

However, a closer look reveals that Prawitz's and Read's counterexample has no teeth. 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 1999: 76] and [Halbach 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.[26]

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.[27] Let *A* be a logical truth and *B* a contradiction, and let Π be a proof of *A* and Σ a

[24] See e.g. [Ketland 1999] and [Shapiro 1998b] for details.
[25] I take it that appropriate precautionary measures have been taken to ward off paradox.
[26] We have shown that the alleged putative counterexample poses no threat even if we accept that the truth predicate is governed by ideally harmonious inference rules. As we will now see, however, this assumption is unsustainable.
[27] I owe the proof to follow to an anonymous referee. It constitutes a significant simplification of my original proof.

disproof of *B*. Then consider

$$\dfrac{\dfrac{\begin{array}{c}\dfrac{\dfrac{\dfrac{\begin{array}{c}\Pi\\ A\end{array}}{Tr(\lceil A\rceil)} \qquad [\lceil A\rceil=\lceil B\rceil]}{Tr(\lceil B\rceil)}}{B}\\ \Sigma\\ \bot\end{array}}{\neg\lceil A\rceil=\lceil B\rceil}}{\dfrac{\exists y\neg\lceil A\rceil=y}{\exists x\exists y\neg x=y}}$$

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.[28] So does this not show that an ideally harmonious operator can still be the cause of non-conservativeness?

I do not think it does. Whatever form an account of ideal harmony might take in the end, it is clear that a constant obeying the following trivial rules must qualify: *A*/$(*A*) and $(*A*)/*A*. $ displays a deductive equilibrium if ever there was one. But clearly $ is not the truth predicate. For one, as we mentioned above, some restriction is required in order to block paradox. This means that the inferential behaviour of a restricted truth predicate might deviate from that of our rather bland $-operator. So at least in cases where the necessary restrictions affect the rules for the truth predicate, we cannot be sure that the restrictions can be implemented while preserving ideal harmony. More importantly still: 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). And this additional operator must figure in the rules of inference governing the truth predicate. We therefore find, as before, in the case of the adjunction of the Tarskian theory of truth to **PA**, that the fault does not lie with the inferential behaviour of the truth predicate (or rather of $) itself but with the additional apparatus required to put it to work. It is to the term-forming operator that we must attribute the phenomenon of non-conservativeness.[29]

So why exactly does our proof above not constitute a counterexample to Dummett's conjecture? Recall that Dummett's conjecture states that intrinsic harmony implies total harmony in a 'context where stability prevails' [Dummett 1991: 250]. This should be read as, 'Take a system in which all of the rules of inference are ideally harmonious. Adding a further ideally harmonious inference rule to such a system will result in a systematically conservative extension'. Therefore, any putative counterexample to Dummett's conjecture must exhibit a system in which ideal harmony 'prevails' and an ideally harmonious logical constant, which, when adjoined to the system, demonstrably results in a non-conservative extension. Have we supplied such an example? I do not think we have. The truth predicate, we have said, only has a place in a system that contains a name-forming device. But the behaviour of this name-forming device must be characterizable independently of the truth predicate, since we can form expressions involving it but not the truth predicate. For the example to work we would need a name-forming device that is itself governed by ideally harmonious introduction and elimination rules. Is such an ideally harmonious name-forming operator likely to be found? It is hard to see what the introduction and elimination rules for such an operator might look like (let alone harmoniously matching ones). An obvious place to look would be Tennant's general inferentialist theory of abstraction operators, in which he shows how term-forming

[28] This is just an instance of the commonplace—though not, it should be noted, universally accepted (see e.g. [Tennant 1997] for a dissenting view)—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.

[29] In the proof above it is the term-forming operator that is really doing all of the work. The truth predicate is needed only to derive $\neg\lceil A\rceil=\lceil B\rceil$ from which it follows that there are at least two objects.

operators can be viewed as obeying introduction and elimination rules [Tennant 2004]. 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 is meta-theoretical in nature and not expressible in the object language. Only if the quotation operator could be shown to be expressible in terms of (ideally) harmonious inference rules would the above example pose a threat to Dummett’s conjecture.

Dummett’s conjecture is thus not affected by these alleged ‘counterexamples’. 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. Moreover, insofar as harmony can be taken to be at least a necessary condition for logicality, we are led to the conclusion that the truth predicate, insofar as it incorporates a (seemingly non-logical) name-forming operator, does not qualify as a logical expression (at least in the absence of an ideally harmonious name-forming operator). Therefore, Stewart Shapiro is mistaken 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: 617].

That being said, we have not, of course, ruled out the existence of genuine counterexamples to Dummett’s conjecture. Indeed both Prawitz and Read also suggest that ‘higher-order concepts’ [Read 2000: 127] might be a possible source of non-conservativeness. The thought, I take it, is that the resources of third-order logic enable us to construct a truth predicate for second-order arithmetic. With the resources of third-order logic we can define a truth-predicate that satisfies the Tarskian compositional axioms, which enables us to prove the consistency of second-order **PA** and thus induces non-conservativeness.[30] The truth predicate, though not harmonious in and of itself (for the reasons just considered), is definable purely in terms of third-order logic. The question thus arises whether instances of higher-order non-conservativeness like the one exemplified by third-order logic over second-order arithmetic constitute a counterexample to Dummett’s conjecture. The answer is that they would have to, *provided* that higher-order logics really are logics *and* can be given harmonious inference rules. On views like Tennant’s that equate logicality with characterizability in terms of (ideally) harmonious introduction and elimination rules, all boils down to the question as to whether higher-order quantifier rules satisfy ideal harmony. However, one might wish to separate questions of logicality and of harmony, denying that expressibility in terms of harmonious I- and E-rules is not only necessary but also sufficient for logicality. In this case a scenario where higher-order quantifiers, though expressible in terms of harmonious introduction and elimination rules, do not qualify as properly logical concepts would be imaginable.[31] Insofar as Dummett’s conjecture is presumably intended to apply only to the logic-specific notion of harmony—ideal harmony is, after all, a logic-specific notion—such a view would not tell against the conjecture. However, in the absence of worked out accounts of logicality and ideal harmony we can do little more than speculate at this point.

8. Conclusion

Let us briefly recapitulate the conclusions we have reached. We began by identifying *general harmony* as our informal characterization of the notion of harmony. We then noted the inadequacy of conservativeness constraints as formal counterparts to general harmony. In particular, we pointed out Dummett’s mistake of identifying harmony for the logical constants with the global property of *total harmony* (i.e. conservative extendibility). However, Dummett’s local logic-specific constraint, *intrinsic harmony*, also missed the mark; in order to block E-weak disharmony, we must supplement intrinsic harmony with a *stability*

---

[30] See [Wright 2007: 167]. I am grateful to Julien Murzi for alerting me to this literature and for helpful discussions on this topic.

[31] In particular, one would wish to refrain from equating logicality with expressibility by harmonious inference rules if one has independent grounds for believing that logic should not, on its own, have existential import. It is the present author’s view that the question as to whether logic must by its very nature be ontologically neutral is a deep and important question in the philosophy of logic that should not be treated as a merely definitional matter.

requirement. Together, they give rise to the elusive demand of *ideal harmony*. Having shown that the salient global property that ideal harmony should guarantee is *normalizability*, not total harmony, we proceeded to clarify the relationships between these global notions and the notion of ideal harmony. First, we demonstrated that Dummett was wrong in assuming that normalizability implies conservativeness. Second, we offered a preliminary defence of Dummett's conjecture that ideal harmony entails conservativeness against attacks by Read and Prawitz. We found that the fate of Dummett's conjecture will depend on whether higher-order logics count as properly logical and are expressible in terms of ideally harmonious inference rules. Our discussion strongly suggests that ultimately an adequate formulation of harmony will have to be a local constraint that must incorporate an account of stability so as to entail normalizability.

Queens' College, University of Cambridge

References

Belnap, Nuel 1962. Tonk, Plonk and Plink, *Analysis* 23/6: 130-134.
Brandom, Robert 2000. *Articulating Reasons*, Cambridge: Harvard University Press.
Dummett, Michael 1973. *Frege: Philosophy of Language*, Cambridge: Harvard University Press.
Dummett, Michael 1977. *Elements of Intuitionism*, Oxford: Oxford University Press.
Dummett, Michael 1993. Language and Truth, in his *Seas of Language*, Oxford: Oxford University Press: 147-165.
Dummett, Michael 1991. *The Logical Basis of Metaphysics,* Cambridge: Harvard University Press.
Halbach, Volker 2009. Axiomatic Theories of Truth. *Stanford Encyclopedia of Philosophy* (Winter 2009 Edition), ed. Edward N. Zalta, URL= <http://plato.stanford.edu/entries/truth-axiomatic/>
Ketland, Jeffrey 1999. Deflationism and Tarski's Paradise, *Mind* 108/429: 69-94.
Milne, Peter 2002. Harmony, Purity, Simplicity and a 'Seemingly Magical Fact', *Monist* 85/44:
498-534.
Prawitz, Dag 1965/2006. *Natural Deduction*, Mineola: Dover.
Prawitz, Dag 1977. Meaning and Proofs: On the Conflict Between Classical and Intuitionistic
Logic, *Theoria* 43/1: 2-40.
Prawitz, Dag 1994. Book Review of *The Logical Basis of Metaphysics*, by Michael Dummett, *Mind* 103/411: 373-376.
Prior, Arthur 1960. The Runabout Inference Ticket, *Analysis* 21/2: 38-39.
Read, Stephen 2000. Harmony and Autonomy in Classical Logic, *Journal of Philosophical Logic* 29/2:
123-154.
Read, Stephen Forthcoming. General-Elimination Harmony and the Meaning of the Logical Constants, *Journal of Philosophical Logic*.
Rumfitt, Ian 2000. 'Yes' and 'No', *Mind* 109/436: 781-823.
Shapiro, Stewart 1998a. Induction and Indefinite Extensibility: The Gödel Sentence is True, But Did
Someone Change the Subject?, *Mind* 107/427: 597-624.
Shapiro, Stewart 1998b. Proof and Truth: Through Thick and Thin. *Journal of Philosophy* 95/10: 493-521.
Steinberger, Florian 2009. Not So Stable, *Analysis* 69/4: 655-661.
Tennant, Neil 1987. *Anti-Realism and Logic*, Oxford: Oxford University Press.
Tennant, Neil 1997. *The Taming of the True*, Oxford: Oxford University Press.
Tennant, Neil 1999. Negation, Absurdity and Contrariety, in *Negation*, ed. Dov M. Gabbay and Heinrich Wansing, Amsterdam: Kluwer: 199-222.
Tennant, Neil 2004. A General Theory of Abstraction Operators, *Philosophical Quarterly* 54/214:
105-136.
Tennant, Neil Forthcoming. Inferentialism, Logicism, Harmony and a Counterpoint, in *Essays for Crispin Wright: Logic, Language and Mathematics*, ed. Alexander Miller, Oxford: Oxford University Press.
van Dalen, Dirk 1997. *Logic and Structure,* Berlin: Springer.
Wright, Crispin 2007. On Quantifying into Predicate Position, in *Mathematical Knowledge*, ed. Alexander Paseau, Michael Potter and Mary Leng, Oxford: Oxford University Press.