

John Benjamins Publishing Company



This is a contribution from *Pragmatics & Cognition 13:1*

© 2004. John Benjamins Publishing Company

This electronic file may not be altered in any way.

The author(s) of this article is/are permitted to use this PDF file to generate printed copies to be used by way of offprints, for their personal use only.

Permission is granted by the publishers to post this file on a closed server which is accessible to members (students and staff) only of the author's/s' institute.

For any other use of this material prior written permission should be obtained from the publishers or through the Copyright Clearance Center (for USA: www.copyright.com).

Please contact rights@benjamins.nl or consult our website: www.benjamins.com

Tables of Contents, abstracts and guidelines are available at www.benjamins.com

© 2005. John Benjamins Publishing Company

All rights reserved

Responses

Robert B. Brandom

Preamble

These admirable, insightful essays raise many more useful and interesting considerations than I can possibly address here — or, in some cases, yet know not how to address at all. It is a rich haul. As a result, I have been obliged to confine myself to commenting on a relatively small number of points, accordingly as they seemed most telling and illuminating, and as what I might say in response could turn out to be most helpful.

Pirmin Stekeler-Weithofer offers a fascinating rational reconstruction of the tradition out of which the theory of *Making It Explicit* (MIE) grows, enriching our understanding of both by the context his magisterial meta-narrative provides. I am in equal parts admiring of and instructed by it. He closes by expressing the concern that my account uncritically and unnecessarily (by its own better lights) — although intelligibly, in virtue of some of its antecedents — veers into a kind of naturalistic social behaviorism. This worry points to the fineness of the methodological line the account is committed to pursuing. On the one hand, the project is to *say*, in terms that do not make essential appeal to attributions to practitioners that employ *semantic* or *intentional* vocabulary, what one must *do* in order to count thereby as *saying* something — that is, to do something to which semantic and intentional vocabulary properly apply. That enterprise has a reductionist sound (though I would insist that the fact that what is sought is a non-intentional *pragmatic* metalanguage makes a big difference). On the other hand, in pursuing that project, essential use is made of *normative* vocabulary, and it is asserted that such vocabulary is both ineliminable in this context, and irreducible to non-normative vocabulary. That does not sound so reductive or naturalistic.

While insisting on the irreducibility of the normative to the non-social natural, I am concerned to show that norms implicit in social practice — normative statuses instituted by our practical (initially, non-conceptual) attitudes towards one another — need not be thought of as *supernatural* and (so) spooky. There were no commitments or responsibilities, no authority or responsibility, until hominids started taking or treating each other in practice *as* committed or entitled, authoritative or responsible. Here is an analogy: the status of being *food* is instituted by the practical attitudes of animals, who take or treat something *as* food by *eating* it. Food is what they are treating something as when they eat it. The social pragmatic phenomenalist approach to the kind of normativity that articulates discursive practice accordingly begins by asking: what is to the status of a performance *as appropriate as eating* is to the status of something *as food*? My suggestion is that *sanctions* — rewards and punishments — can play that role. How reductively naturalistic that idea is depends on how one conceives sanctioning. If it is construed as positive or negative reinforcement — responding in ways that raise or lower the probability that the performance being reinforced will be repeated — then the conception of norms is indeed behavioristic.

But in a social setting, it is easy to build on such a basis a notion of normativity that swings free of statistical regularities of behavior. Suppose the hominids Og and Ug build a hut and guard it with clubs, driving off all who attempt to enter it unless the petitioners show a leaf from the special tree at the top of the mountain. They are in a straightforward sense, explicable in wholly behavioristic terms, treating such leaves as *licenses* or *tickets* having the normative status of *entitling* their possessors to do something they would not otherwise be permitted to do: enter the hut. But now imagine that in the same community, a teacher rewards her young charges who respond to differently colored objects by uttering noises that she is in this sense treating as appropriate or correct, by offering them leaves from the hard to reach ticket-tree. Since these confer a *right* that one would not otherwise possess, they are intelligible as *rewards*, that is, as positive sanctions. And that is so even if some or even all of the charges are not in fact behaviorally moved (positively reinforced) by those rewards — intuitively, because they don't care about entering Og-and-Ug's hut. Chapter One of *MIE* describes how chains of *normative* sanctions — rewards and punishments that consist only of alterations of normative status — can be instituted in this way, and further, how in these terms we can make sense of normative sanctions that are not connected in any way to behavioral reinforcements. This is the strategy by which normative pragmatism seeks to distinguish itself from social behaviorism.

Jaroslav Peregrin's thoughtful, wide-ranging essay focuses on one element of the broader context in which Stekeler-Weithofer places the project of *MIE*. He situates it by means of an extended compare-and-contrast exercise relating two broad traditions in philosophical and linguistic theory. I welcome and endorse his characterizations, and find particularly interesting his suggestion that important strands in Chomsky's own thought would underwrite a far more Wittgensteinian, pragmatic, approach to semantics than he or his followers have typically pursued. Summarizing his own assessment, Peregrin says "I think we should settle for a use-theory of meaning, i.e., a theory according to which the meaning of an expression amounts to the role the expression plays within our linguistic transactions". My own view is slightly different. At the most general level, it is expressed in the slogan: "Meaning is to (proprieties of) use as theory is to observation". It is licit to postulate meanings, which need not themselves be thought of (contra Dummett) as themselves aspects of use, in order to codify and explain how it is correct to use expressions. Associating a meaning or content with a sentence or subsentential expression goes hand in hand with specifying in a systematic way how to derive from it various normative features of the use of that expression — for instance, the significance of its assertional, imperative, or interrogative use, its use as the antecedent or consequent of a conditional, and so on. At a more specific level, I suggest a candidate for playing the methodological role of meanings: role in inference, substitution, and anaphoric chains (the ISA approach to semantics).

Ruth Garrett Millikan, who is Wilfrid Sellars's most eminent and accomplished student, has done us the service, and me the honor, of offering a framework within which to compare and contrast her approach to intentionality with mine, looking at each particularly with respect to its debt and relation to our common intellectual 'father'. Richard Rorty originated the now-traditional division of Sellarsians into a right and left wing. In this botanization, right-wing Sellarsians are most impressed with Sellars's scientific realism, taking their cue from his slogan that "science is the measure of all things, of those that are, that they are, and of those that are not, that they are not". Their manifesto is his famous essay "Philosophy and the scientific image of man". Left-wing Sellarsians are those who take Sellars's greatest philosophical insights to be the critiques of semantic atomism and (so) epistemological foundationalism at the core of his masterwork "Empiricism and the philosophy of mind". (Rorty incidentally expressed the hope that these two schools of thought could work out their differences more irenically than did the right- and left-wing Hegelians, who settled theirs at a strenuous, extended conference otherwise known as the

battle of Stalingrad.) There is something right about this opposition, and along this dimension I think Millikan is slightly to my naturalistic right. But since she has fully absorbed Sellars's lessons in "EPM", in our case the left/right division does not seem to cut at the joints. She has sketched a different context, one that is much more helpful.

One way of thinking about it is to look at what we each take out of another central Sellars essay that we both admire: "Some reflections on language games". Impressed by his account elsewhere of the way worldly regularities can come to be *pictured* by aspects of the activity of creatures making their way through the world, Millikan takes up out of that essay Sellars's notion of *pattern-governed* behavior as a semantically crucial category much richer than merely *regular* behavior, but still far short of (though arguably helpful in the explanation of) fully intentional *rule-governed* behavior. (I talk about this a bit in the twin critiques of what I call 'regularism' and 'regulism' in the first chapter of *MIE*.) In particular, Millikan is inspired by Sellars's idea that the way to get a *naturalistic* story that can fund some of the basic kinds of *normative* assessment required to see *semantic* content as coming into play is to look at the *diachronic* dynamics by which some behavioral patterns can be established and stabilized. The sophisticated, broadly evolutionary story that Millikan tells to elaborate this thought integrates and illuminates the relations between the natural and the normative dimensions of intentionality (roughly, those that impressed the right- and the left-wing Sellarsians, respectively) *far* better than that of Sellars himself — for whom the naturalistic dimension of picturing and the normative dimension of semantic assertibility remained in largely unresolved tension.

By contrast, inspired by Sellars's seminal "Inference and meaning", when I read the language-games essay I am impressed first by the thought that what makes what he there calls "language-*entry*" transitions, in perception, and "language-*exit*" transitions, in action, *language* entries and exits — what gives them specifically *conceptual* content at all — is the *inferential* connections manifested in what he calls "language-language" moves.

As Millikan stresses, these differences of emphasis are by no means irreconcilable. She, for instance, accepts the essential role of inference and public assertion in *linguistic* intentionality — a category to which she devotes a great deal of analytic effort — but is centrally concerned to show how that sort of intentionality is built on more basic, non-linguistic varieties. I think her argument that a single broad form of explanation — roughly, in terms of selectionally stabilized Proper Functions — suffices to account *both* for various kinds of primitive content discernible in a wide variety of purely biological cases

and for an impressive variety of quite sophisticated purely linguistic semantic and intentional phenomena (indexicals, substance-terms, identity locutions, propositional attitude ascriptions...) is a theoretical *tour de force*, and one of the wonders of our philosophical age. I suspect, however, that Millikan tends to overemphasize (in her own characterizations of what she is doing, as opposed to what she actually does) the biologicistic character of her account (for instance, in the purposely provocative title of her path-breaking work: *Language, Thought, and Other Biological Categories*), relative to the crucial (for her, not just for me) *social* element. What I find most admirable about her theory is the way she applies a form of explanation demonstrably significant already at a merely biological level — not just pre-sapient or pre-mammalian, but even at a sub-organismic level — in explaining key *semantic* features of signs and properties of signs in terms of the *pragmatics* of their production and consumption, and especially the *social* relations between producers and consumers. But I do not think that this brilliant *extension* of the apparatus should be thought of as having the *reductionist* consequences she sometimes is pleased to attribute to it.

Sebastian Knell, as the author of the first book-length study of *MIE*, has long studied and fully mastered the theory developed there. In his essay, he offers a very clear account of the understanding presented there of the use of the non-theoretical representational locutions we employ to talk and think about what we are talking and thinking *about*: the ‘of’ of “I am thinking of a number between 1 and 10”, (rather than the ‘of’ of “the pen of my aunt”), and the ‘about’ of “He talked about Hegel”, (rather than the ‘about’ of “the book weighs about five pounds”). He then offers a brilliant, original way of thinking about the philosophical significance of that account: it should be thought of as offering a *deflationary* theory of intentionality, by strict analogy to deflationary theories of truth. This strikes me as an extremely promising line of research.

Sebastian Rödl’s deep and thoughtful essay contrasts the abstract semantic view of objects presented in *MIE* to a richer, temporal view specific to empirical objects. The more general approach treats objects just as what singular terms purport to refer to, and then understands that ‘purport’ in terms of their behavior in substitution inferences. (Existential commitments appear as a special kind of substitutional commitment, with different sorts of existence — e.g., spatio-temporal, numerical, fictional — distinguished by the different classes of singular terms that play the role of “canonical designators” in them.) This notion of *object* (where ‘object’ itself plays the role of a pro-sortal, rather than

of some variety of extremely general sortal) is intended to be general enough to apply, Frege-wise, to numbers, as well as to observable objects. Rödl argues that this notion leaves out something crucial to the identification and re-identification of spatio-temporal objects through empirical change. Taking his cue from Kant, he suggests what is missing: a notion of the temporal form of judgment. I think this is a genuinely profound idea, and I do not pretend to see to the bottom of it. But he may be right.

Friedrich Kambartel very usefully situates my inferentialist pragmatism in the context of the German tradition of philosophical constructivism, to which, among many other traditions, he has made distinguished contributions. Although I was ignorant of the work of this school when I wrote *MIE*, I have since come to admire its many achievements, and to appreciate the many commonalities of motivation that animate our projects — helped immensely by many long and stimulating conversations in which Kambartel, with great patience and forbearance, both educated me and conducted a careful compare-and-contrast exercise. I have been completely convinced of the justice of his principal complaint: that my exclusive emphasis on *inference* overlooks the crucial contribution to their conceptual content that is made by the role expressions play in various *constructions*. The way some more primitive discursive abilities can be systematically arranged and deployed so as to amount to more complex discursive abilities is both an essential feature of the content of the expressions involved, and is not reducible to the role they play in inference. Understanding such constructive elaboration of more basic concept-using abilities will, I think, require close attention to inferential roles, but it is a kind of reasoning, in a broad sense, and yields a kind of conceptual understanding that is not exhausted by those roles as they show up in my account. My lack of attention to this important species of concept-articulating practice is made all the more embarrassing by the fact that a great deal of the explicative work actually undertaken in *MIE* consists precisely of construction, rather than more narrowly inferential argumentation. The whole semantic project of Part Two of the book can be thought of as using the notions of *inference*, *substitution*, and *anaphora* to construct the semantic contents expressed by various familiar sorts of philosophical-linguistic expressions: semantic vocabulary ('true' and 'refers'), intentional or representational vocabulary ('of' and 'about'), propositional-attitude ascribing locutions, singular terms, proper names, expressions of existential commitment, and so on.

Kambartel points out that while both approaches understand conceptual content in terms of practices of justifying claims, a principal difference between

their strategies is that constructivism employs a very strong notion of justification, generalizing from the model of mathematical proof to something like dispositive evidence, while inferentialism employs a very weak notion of justification, construing it as merely probative evidence: something like there being some reasons in favor of the conclusion. I do not fully understand the worry that he expresses concerning the latter. He seems to think that the weak notion of justification, since it can't guarantee *truth*, can't be used to characterize the very same *contents* that are called 'true'. But the truth claim will just inherit whatever kind of content the antecedent of the prosentence has. And that content is *not* characterized just in terms of the dimension of entitlement, which is (as he desires) just one "element in a much more complex conceptual situation". For one of the basic innovations of *MIE* is that its normative pragmatics begins with *two* kinds of normative status: *commitments*, as well as *entitlements* to such commitments. The very weak entitlement-preserving inferences are only one species, with much stronger commitment-preserving inferences and incompatibility entailments (discussed below) also important in articulating conceptual content.

John McDowell's paper is the most recent fruit of a long, thoughtful, and generous (which is not to say, in the end sympathetic) engagement on his part with *MIE*. As a result, his trenchant criticisms are directed for the most part at core elements of the project pursued there.

One of his concerns is the contrast between representationalist and inferentialist orders of semantic explanation, which is invoked both to motivate the adopting the latter, and to situate doing so in a historical context. (The latter enterprise is pursued at greater length in my later book, *Tales of the Mighty Dead: Historical Essays in the Metaphysics of Intentionality*, so I will say no more here about how I understand the pre-20th century semantic tradition.) The most important observation McDowell makes in this connection, I think, is that it need not be the case either that representational relations or concepts should be understood as prior in the order of semantic explanation to inferential ones, or that inferential ones should be understood as prior in the order of semantic explanation to representational ones. It may be, after all, that neither can be understood apart from the other — that reference and inference come as an indissoluble conceptual package that cannot be analyzed reductively, but only relationally. I agree, of course. Looking for a way to get an independent theoretical grip on one range of concepts, and then explicating the other in terms of it is only *one* strategy for illuminating the relations between the representational and inferential perspectives on semantic content. If an account of

this shape *is* possible, however, then we understand the relations between the two perspectives better in terms of it than we would if we did not realize that fact and content ourselves from the outset with a holistic relational picture.

I think we have learned a great deal — indeed, most of what we do understand in semantics — from the pursuit of the representationalist order of explanation, from Frege’s *Grundgesetze* at the outset of the modern history of the subject, through Russell and Whitehead’s *Principia*, Wittgenstein’s *Tractatus*, the Vienna Circle’s logical empiricism, Tarski’s model theory as interpreted philosophically by Quine, down to the programs in naturalistic semantics of our own day, such as Fodor’s. The atomistic, bottom-up, representationalist semantic strategy of beginning with denotational relations between singular terms and objects, and predicates and sets of objects (or tuples of objects), then appealing to set-inclusion relations among them to assign truth conditions to atomic sentences, then to functions taking all these things as arguments and values to assign intensions to logical and other operators (such as modal ones), and finally appealing to set-theoretic relations among those structures to derive inferential relations has had a lot of successes. For example, one difference in the inferential behavior of adverbs is that for some of them, the attributive adverbs, if someone performs action A X-ly, it follows that they perform action A. If Donald butters the toast slowly or in his kitchen, then he butters the toast. But for others this inference does not go through. If Donald butters the toast symbolically, or in his imagination, then it does not follow that he butters the toast. In David Lewis’s generalization of California semantics (in his “General Semantics”), if we take singular terms to represent objects, and sentences to represent sets of possible worlds, then one-place predicates represent functions from objects to sets of possible worlds — intuitively, those in which the objects represented have the properties (thereby) represented. Then adverbs represent functions from functions from objects to sets of possible worlds to functions from objects to sets of possible worlds. In these terms we can say exactly what the difference is between the functions represented by attributive adverbs, which support the inference in question and those represented by non-attributive adverbs, which do not. Seeing the possibility of this sort of reductive representationalist explanation of an inferential phenomenon is genuinely illuminating, and ought to be a part of any account of the relations between the two dimensions. (I go to some trouble in the Appendix to Chapter Six of *MIE* to show that and how a substitution-inferential approach can reproduce just this sort of story — about adverbs and other sub-sentential expressions — working, as it were, from the top down instead of the bottom up.) The aim of *MIE*

is not to say that the inferentialist order of explanation is the only one that can provide semantic illumination. It is to explore what kind of illumination it *can* provide, given that we have learned so much from the contrary linear order. For phenomena not amenable to either approach, we will have to fall back on the merely relational account. But retreating in the face of the mere possibility that a more robust (more committive, hence riskier) explanatory strategy may fail seems prematurely defeatist.

McDowell points out that the label ‘inferentialism’ is in some ways misleading, given that the official view of *MIE* is that *inference* and *assertion* are co-ordinate concepts, in the sense that one cannot engage in inferential practices without engaging in assertional ones, and *vice versa*. Specifically *discursive* practices — the ones that really deserve to be called *Sprachspiele* — are characterizable alike as those in which some performances have the significance of making claims and those in which some serve as reasons for others. He urges that the same co-ordinate, non-reductive status be acknowledged for the case of *inference* and *representation*, pointing out, quite correctly, that there can be no inference properly so-called without representation. I think these phenomena — assertion, inference, representation — do come as a package, but that there are nonetheless important insights to be gained by exploiting asymmetries concerning various explicative paths through this constellation. The representationalist tradition has thought of assertion as putting something forward as a fact (and of judging as taking something to be a fact). And in its bottom-up way, it has thought of facts as representations of (in the purest form, pictures of) states of affairs: objects standing in relations. This is alright for “The cat is on the mat”, a bit strained for “Helsinki is the capital of Finland”, and substantially less helpful for “Freedom is better than slavery”. Even after Wittgenstein in the *Tractatus* showed how to get over the difficulties Russell’s logical atomism had had with *negative* and *conditional* facts, *probabilistic*, *modal*, and *normative* facts remained basically beyond the pale of this approach to propositional content. Just how is one saying that objects are related to one another in saying how things probably or possibly are, or how they ought to be? Thinking of propositional contents in the first instance as what can both serve as and stand in need of reasons, that is, as what can play the role both of premises and conclusions of inferences offers a way out of this metaphysical dead-end. For it directs our attention usefully to the distinctive role modal statements play in counterfactual reasoning, and normative ones play in practical reasoning.

Similarly, though *inference* and *representation* are co-ordinate concepts, that fact does not preclude learning something about their relations by explicating

one in terms of the other. The representationalist tradition takes an inference to be good in case the situations represented by the premises are included in those represented by the conclusion — either as a matter of contingent fact, or in a stronger sense, in every possible correct application of those representations. On this line, characterizing propositional content in representational terms underwrites a characterization in inferential terms. The main point of the second half of *MIE* is to show how, beginning with a characterization of propositional content wholly in terms of inferential role, it is possible to elaborate on that basis a characterization in terms of representational content. Insofar as the argument succeeds, it shows that there can be no inferential content without representational content: that neither the inferential and representational dimensions of conceptual content can be present without the other. But the aim of Chapters 5–8 of *MIE* is to say (to make explicit) in inferential terms what that representational dimension is, by specifying the inferential roles of the representational and intentional locutions that make it explicit: terms such as ‘true’ and ‘refers’, and ‘of’ and ‘about’. (Knell’s essay in this volume offers a good account of this bit of the story.) Explicitly inferential talk turns out to be implicitly representational. By explicating the use of the representational terminology by means of the inferential terminology, we gain a better understanding of what we are *doing* when we think and talk about what we are thinking and talking *about*. The stark opposition McDowell assumes between taking concepts to “come in a package, each intelligible partly in terms of the others”, on the one hand, and “conforming to a foundational structure on the other, is too limited to be an exhaustive characterization of the possibilities. Even where *reduction* is not possible, *explication* — *expressing* the content of one vocabulary by the use of another — can be. And it is one of the tasks of *MIE* to make explicit a useful sense of “making explicit”: to *say* in one vocabulary what is *implicit* in the *use* of another.

McDowell frames his essay as a series of criticisms of the *motivations* I offer for the direction of pragmatic and semantic explanation pursued in *MIE*. He notes that in the past I have responded to these by saying that the issue of why, in advance of seeing the details of the project, one might think it a promising one is not of the first importance. For in the end, it should be judged retrospectively on what illumination it actually turns out to provide. The point is not what he calls the ‘advertisement’ (reasons to read the rest of the book), but the product. This is the “proof is in the pudding” response. He says that it would be in order if *MIE* actually delivered a semantic account in inferentialist terms. His most serious criticism is implicit in his claim that it has not — that there is

no “pudding” to be assessed. His principal reason for denying that it has been shown that any significant semantic work can be done in inferentialist terms is his contention that the practices of keeping score on commitments and entitlements and their relations that is related in Chapters 3 and 4 do not suffice to get any intelligible notion of inference, assertion, or (therefore) conceptual content in play. For all that is specified there, he thinks, the practices involved could be *just* a game, involving no genuine claiming or reasoning, no saying how things are at all. In effect, he thinks the practices described there are just a *shadow* of genuine discursive practice, with only as-it-were inferences connecting mock claims. Since genuine conceptual contents, purporting to say how things objectively are, are not successfully put in play to begin with, no amount of formal elaboration of the imitation contents later on can make them real.

This criticism goes to the very heart of the enterprise. So it is important to be clear about how that enterprise is understood to proceed. In the final chapter of *Articulating Reasons* (unpacking some of the claims of Chapter Eight of *MIE*), I argue that any practices recognizable as a game of giving and asking for reasons — and hence (since it is not contested that these phenomena come as a package) as one that accords any performances the significance of being claims or assertions, that is, items with *propositional* content — must involve practically distinguishing between two kinds of normative status: commitments and entitlements. For making a claim must make a difference to what the speaker is committed to, and must have the practical significance both of making the claim available to serve as a reason for other claims, and make the speaker liable to assessment as to the reasons she has for it — that is, whether and how she is entitled to the claim. Keeping practical track of those two kinds of normative status involves distinguishing which commitment-undertaking performances have other commitments as their consequences, and which further commitments they entitle one to and preclude entitlement to. The basic transition from this normative scorekeeping pragmatics to a basic inferential semantics is then made by the claim that corresponding to these consequential relations are genuine inferential relations among what now show up as the *contents* of the claims (the claimables expressed by the claimings):

- i. commitment-preserving inferences, which are a generalization (from the domain of formal to that of material inferences) of *deductive* inferences;
 - ii. entitlement-preserving inferences, which are a generalization (from the domain of formal to that of material inferences) of *inductive* inferences;
 - iii. incompatibility entailments, a generalization of *modal* inferences.
- (Two claims are incompatible in the deontic scorekeeping sense in case

commitment to one precludes *entitlement* to the other. One incompatibility entails another just in case everything incompatible with the conclusion is incompatible with premise.)

For the things (corresponding to equivalence classes of possible performances) that stand in these relations induced by consequential scorekeeping relations to be recognizable as *conceptual* contents, it is claimed, they must also be caught up in three other kinds of practices: *testimonial* inheritance of entitlement to commitment (in which one interlocutor's assertions count as available to serve as reasons for another's), attributions of entitlement as the result of *reliability* inferences concerning language-entry moves in observation (in which the circumstances in which the one taken to be a reliable reporter comes to acknowledge a commitment are taken to entitle him to it), and *practical* inferences (in which commitments entitle interlocutors to non-linguistic performances).

The big, bold claim that ties together the two halves of *MIE* is that when the commitments and entitlements that have been argued to be a *necessary* feature of practices recognizable as involving giving and asking for *reasons* (and hence, assertions, which are what can both be given as reasons and have reasons demanded for them) are elaborated in this practical-consequential structure, the result is *sufficient* for genuinely *discursive* (that is, *conceptual*) practice: that the performances should count as *assertions* and the moves as *inferences*. This is the claim that McDowell rejects. As far as he can see, practices with this structure could be *just* a (non-competitive) game. (I'm not sure I see the relevance of the possibility that Martians might take it to be only such a game. The question is whether or not such a response would be a *mis*-understanding.) And he is certainly right that it is not enough to justify that claim just to start *calling* what goes on there 'assertion' and 'inference'. But although I do start using those words already in Part One, the *justification* for using them is supposed to be provided by Part Two. For the task of Part Two is to start only with the raw materials provided by Part One — that is, only with 'inference' and 'assertion' insofar as those concepts are underwritten by the three kinds of proto-inferential consequential relations among normative statuses (together with the other sorts of interpersonal status inheritance) — and seeing how much recognizably *linguistic* and *conceptual* structure can be elaborated solely on that basis. That is the "pudding" in which the proof is supposed to be found. The test of whether what has been constructed *is* genuine inference-and-assertion is the feasibility of the "collapse of levels" described in Chapter Nine: whether the utterances of a community meeting the conditions in question could be mapped

onto our own so as to make conversation possible (a version of Davidsonian interpretability).

Now McDowell's response is that if what we start with is not really inference-and-assertion, then what is built on that basis is not really, say, singular terms. He does not deny in that connection that my account of what it is to function semantically as a singular term, appealing to substitution inferences, is adequate. But that story is only relevant to the story of Part One if what is constructed there counts as *inference*. And that he denies. Further, it is only on the basis of those prior claims that I can be in a position to appeal to inference without having yet gotten representation on board — since the fact that Part One does not appeal to representational locutions is the *only* warrant for the claim that one could be entitled to the raw materials of the account of singular terms (notions of inference and substitution) without already having to presuppose the applicability of representational locutions. If the conceptual raw materials provided by Part One are only *as-it-were* inferences, then what is later defined in terms of them are only *as-it-were* singular terms: shadows, not the real thing.

My response to this is that while it is true that I am not in a position to *show* the Davidsonian interpretability of practices meeting the conditions laid out in Part One, a *great* deal of the structure that would have to be exhibited for such interpretability *is* on offer. The variety of linguistic-conceptual-logical structure that the apparatus is shown to 'mimic' makes it *extremely* implausible that it is all just a *shadow* of conceptual content. We have *as it were*:

- Conditionals and negation (Ch. 2);
- Language entry moves in perception (Ch. 4);
- Language exit moves in agency (Ch. 4);
- Normative vocabulary (Ch. 4);
- Semantic vocabulary mimicking 'true' and 'refers' (Ch. 5);
- Singular terms and predicates (Ch. 6);
- Identity locutions, anaphoric initiators and dependents, hence deictic expressions, pronouns, bound variables and variable-binding operators such as quantifiers (Ch. 7, and Ch. 5);
- Proper names (Ch. 8);
- Object-dependent indexicals such as 'I' (Ch. 8);
- Propositional attitude ascriptions, including the distinction between ascriptions *de dicto* and *de re* and all arbitrarily iterated combinations of them (Ch. 8);
- Explicitly representational locutions such as 'of' and 'about' (Ch. 8);

It beggars belief that one could reproduce all this structure and not have genuine conceptual contents in play, but a mere shadow of them, a mere *game* that does not involve actually *saying* anything, but nonetheless exhibits all of this structure. If something is missing here, it is something *magical* — for this complaint is like that of the solipsist who insists that whatever *behavior* other humans might exhibit, it is nonetheless not accompanied by what *he* has: a *mind*. If showing the broadly inferential role of all of these locutions is not producing the pudding, what *could* count as doing so?

Notice that the logical locutions that are introduced in Part Two, based on the spare inferential notion of content introduced in Part One have as their expressive job making explicit various aspects of conceptual content. If that is not what the expressions introduced in Part Two do (that is, if they don't serve to make conceptual content explicit because there is no such content in play), then what *are* they doing? How could expressions behave so much like logical locutions and not have that expressive role? But if that *is* their expressive role, then there must be contents in play for them to work on. So, for example, if the conditional and negation introduced in terms of pragmatic incompatibility really *are* a conditional and negation — letting us *say that* two claimables are inferentially related as premise and conclusion, or that they are incompatible — then what those locutions work on must be genuine propositional contents — what appears embedded as the antecedent of the conditional, or is negated. So McDowell has to claim that these are not *really* conditional and negation locutions. And a similar point goes through for the 'singular terms' and 'predicates' distinguish, the 'identity' locutions, and so on. If those really *are* the logical locutions they purport to be, then the contents they operate on and explicate must be genuinely conceptual contents. But if *not*, how is it that they can behave so much like those locutions in their own inferential expressive behavior? The justification for the claim that what is underwritten by the consequential relations implicit in practices of keeping score on commitments and entitlements are genuinely *inferential* relations (of the three flavors mentioned above) is just that the (as-it-were) inferential roles of so many kinds of expressions playing essential expressive roles in natural (and artificial) languages can be elaborated solely on that basis. The extent to which crucial characteristic features of linguistic expressions of many different categories can be reproduced is the best possible reason to conclude that the (mere) intuition that the scorekeeping practices on which they are based could have the complex consequential structure of commitment and entitlement inheritance described, while still being *just* a game, conferring no genuine conceptual contents on

items suitably caught up in it — the intuition on which McDowell relies — is mistaken.

McDowell (in company with some of our other authors) takes issue with another explanatory asymmetry asserted in the book, another portrayal of a semantic layer-cake where he sees only the prospect of a homogeneous mixture. This is the claim that we can make sense of creatures who are *rational*, in the sense that they make claims and give and ask for reasons for them (engage in practices of assertion-and-inference), but are not yet *logical*, in that *all* the inferences they endorse are *material* inferences and no expressions are yet in use that play the inference-explicating role characteristic of *logical* vocabulary. On this issue, I think there is room for genuine philosophical debate and disagreement. One of the principal virtues of *MIE*, it seems to me, is the number of questions of this general shape that its constructive methodology raises — quite apart from the virtues of the stands it takes on those issues. Can there be implicitly normative practices of attributing and undertaking commitments without there being specifically *conceptual* norms in play? (*MIE*: Yes.) Can there be normative *statuses* (such as commitment and entitlement, responsibility and authority) apart from any normative *attitudes* (of taking or treating individuals in practice *as* committed or entitled, responsible or authoritative? (*MIE*: No.) Can there be *conceptual* norms in play in a situation in which there are no speech acts with the significance of *assertions*? (*MIE*: No.) Can there be autonomous linguistic/conceptual/discursive practices that do not include *inference*? (*MIE*: No.) Is every *inferentially* articulated speech act *conceptually* contentful? (*MIE*: Yes.) Can there be *inferential* practices that do not include the use of *logical* vocabulary? (*MIE*: Yes.) Can there be implicit *attribution* of normative statuses without the vocabulary needed explicitly to *ascribe* those statuses? (*MIE*: Yes.) Can there be *deixis* without *anaphora*? (*MIE*: No.) Can there be implicit conceptual norms (including practical ones) without the *normative* vocabulary needed to make them explicit? (*MIE*: Yes.) Is what is made explicit by the use of *modal* vocabulary always already implicit in the use of ordinary, non-logical expressions? (*MIE*: Yes.) Could there be *propositional* contents expressed by *sentences* without sub-sentential expressions? (*MIE*: Yes.) Could there be languages without the means to make *existential* commitments explicit? (*MIE*: Yes.) And so on.

Many of these claims have the same form: although as a matter of *methodology*, we *understand* what is implicit in what can be *done* only in terms of the possibility of making it explicit in what can be *said*, nonetheless in general the implicit capacity to *do* what one must do in order to count as *saying*

anything can precede the capacity explicitly to *say* what it is that one is doing. Explicit sayings are to be seen as built on top of an antecedent layer of doings in which what is said remains as-yet implicit. This methodological approach invites an unusual enterprise: demolishing Neurath's boat at sea. The challenge is to see how much of ordinary discursive practice one can detach and throw overboard without sinking the vessel — that is, without so denuding it as to render it unrecognizable as a *discursive* practice, one in which one can make *claims*, *say* how things are, make something *explicit*, assess *reasons* for what one says and does. The stripped-down skeleton of practices that survives as necessary to remain afloat as minimally discursive is unlike our own in many important ways; the loss of semantic self-consciousness that goes with the various forms of expressive impoverishment is radical and significant. But it is illuminating to break it down, and see how the impoverishment consequent upon not having, say, singular terms, is closely related to that resulting from not having conditionals, but quite different from what one misses out on without attitude-ascribing locutions, modal, or normative vocabulary. It deserves to be controversial whether the minimal discursive practices *MIE* invites us to consider really are intelligible, and whether, if so, they deserve to be seen as involving *claiming* and *reasoning*. McDowell doubts that they are. But these are issues it is worth worrying about, and *MIE* at least gives us a new way of framing and addressing them.

Rational but not yet logical creatures, the claim is, would be able rationally to criticize and revise their *doxastic* commitments, on the basis of inferential connections among them (including relations to those they find themselves with non-inferentially), but would be unable to criticize and revise those *inferential* connections themselves. They could *change* their concepts, but not give and demand *reasons* for doing so — since that requires at least being able to *say that* one endorses a certain inferential connection, and so that one be able to deploy conditional locutions. McDowell doubts that any such practitioners should be counted as making *inferences* at all. Semantic *consciousness*, sapience in the sense of applying *concepts*, for him requires the semantic *self-consciousness* afforded only by the use of logical locutions. He may be right. But I do not see that the description of less-capable creatures stuck at a lower level is unintelligible, and what they would be doing has important features in common with full-blooded assertion-and-inference. The question of what sort of content (if not propositional) they *would* be conferring on their utterances seems to me an important one.

McDowell's objection to seeing logical vocabulary as an in-principle-optional superstructure erected on the basis of a more primitive constellation of practices of deploying non-logical vocabulary relies on independently motivated philosophical considerations (in particular, lessons associated with Kant). Daniel Laurier, in the course of an admirably clear and concise rehearsal of some of the darker doctrines of *MIE*, raises a sophisticated, original *internal* objection to the claim that locutions such as 'claims' and 'believes' — which allow the explicit ascription of propositional attitudes — can be seen as a potentially late-coming development of an antecedent practice in which attitudes of acknowledging and attributing commitments (normative statuses) can be *adopted*, but not yet themselves *attributed*. The objection arises in the course of examining a more general issue: the apparent collision of two methodological requirements. Normative statuses (paradigmatically, the propositionally contentful commitments acknowledged by assertions), it is claimed, are instituted by normative attitudes of attributing and acknowledging them. If, at the most basic level, these are construed as themselves propositional attitudes (Laurier's "conceptually contentful"), then the account seems circular. If, on the other hand, they are not so construed, it seems mysterious how they could institute conceptually contentful normative statuses. ("At the most basic level", because once explicit attitude-ascribing locutions are introduced into an already up-and-running discursive practice (rational, but not yet along that dimension logical), it seems clear that attitudes and statuses can alike be thought of as conceptually contentful.)

The general response to this worry has already been indicated in discussing McDowell. In the most primitive discursive scorekeeping practice, what one is practically attributing (taking or treating another as exhibiting) is commitment to a kind of *doing* that need *not* yet be characterized in terms of its conceptual content: *consequential* commitment, for instance, to acknowledgingly-uttering 'q' if one acknowledgingly-utters 'p' (and is suitably queried or otherwise prompted). This, the claim is, is a kind of practical attitude that belongs in a box with others that can evidently be adopted by non-concept-users, who can practically take or treat someone to be consequentially committed for instance to bringing back a bit of meat for others if the hunt is successful. If these consequential relations between commitments, and those between entitlements, and those relating commitments to entitlements, have the right structure, then practically treating each other as having such consequential statuses *is* in fact (if the theory is right) attributing *propositionally* (hence conceptually) contentful commitments — even though one does not yet have the *concept* of doing so

and so cannot yet be *aware* of oneself *as* doing that (and won't be until one has at least ascriptional locutions, conditional locutions, and normative locutions on board). In the most primitive case, to adopt scorekeeping attitudes is not yet to adopt *intentional* attitudes, except in the minimal sense of practical attitudes that are directed at others. They are only implicitly, and not explicitly, propositional attitudes, that is, attitudes towards conceptual contents.

In the course of considering this general question, however, Laurier raises a much more specific worry: whether the account of the *objectivity* of conceptual commitments — that *what* one is committed to, the *content* of one's conceptual commitments, is not determined by *anyone's attitudes*, not one's own, nor those of the whole community — applies to the primitive-rational case of creatures who can practically *attribute* commitments, but not yet explicitly *ascribe* them, that is, *say* that they are doing that (attributing commitments). For the claim is made in *MIE* that the capacity to attribute *attitudes*, rather than *statuses* — take someone not just to be committed, but to *attribute* commitments, not just keep score on her, but take her to be herself a scorekeeper — is a *logical* capacity, in that it depends on the availability of explicating *ascriptional* locutions. Here the threat to the intelligibility of the pre-logical but still discursive practices comes in effect from Davidson's argument in "Thought and talk": To be a believer, he argues, one must have the *concept* of belief, at least in the sense that one treats beliefs as the sort of thing that can be true or false independently of one's own commitment to what is believed. One must acknowledge some standard of assessment of the correctness of beliefs beyond the fact that one is oneself committed to them. The discussion of this sort of objectivity in Chapter 8 of *MIE* is conducted in terms of the *de dicto* and *de re* ascriptions that make explicit the differences of scorekeeping perspective that ultimately underwrite the account of what it is in this sense to treat one's commitments as answering to objective (attitude-transcendent) standards. Laurier doubts that account can be translated into the terms available in the pre-ascriptional phase.

From one point of view, as he notes, it would not be a big problem if he were right. The claim that ascriptional locutions are in-principle late-coming expressions that make explicit features of practices intelligible as autonomously discursive in advance of their introduction could be acknowledged to be mistaken without upsetting the general methodology or architectonic of the project. But Laurier has constructed here a clear argument, based only on other commitments acknowledged in *MIE*, for an important case of the first-rational/discursive-then-logical layer-cake picture McDowell objects to much more generally. So it is important to consider Laurier's objection. There is no

incompatibility (as he seems to think) between the claim that ascriptional locutions have the expressive role of making explicit what is implicit in practical attributions, on the one hand, and that introducing them gives their users new expressive capacities. For the first claim just requires first that the capacity to use them depends on no capacities one does not already have in engaging in the non-logical practices on which they are based, and second that once they are introduced they let one *say* what it is that one was then *doing*. Being able to say what before one could only do may bring many further capacities with it. Thus, in the index case that McDowell cares most about, if I can sort inferences into materially good and bad ones, and make assertions, I can do everything I need to in order to introduce conditional assertions, to which I'm committed just in case I endorse the corresponding inferences, and whose endorsement commits me to the goodness of those inferences. But once I've got such locutions, I can now *criticize* inferential commitments, giving and asking for reasons for *them*.

The key demand that Laurier properly makes is rather that we be able to say what it is for someone without access to ascriptional locutions to "be able to make (*in practice*, since *ex hypothesis* no ascriptional locution is yet available) a distinction between what he takes someone else to be committed to and what this someone takes himself to be committed to". Laurier immediately paraphrases this "that is to say, between the commitments he attributes to someone and the commitments he takes to be *acknowledged* by this someone". Since acknowledging a commitment is adopting an *attitude* toward it, this would indeed require an attitude toward an attitude — attributing an acknowledgment — which is what it is claimed one cannot do in the absence of ascriptional locutions (which permit that, since one can attribute an attribution by attributing commitment to an ascription). But it is enough that I can distinguish between the commitment I attribute to X and what X is in fact committed to. Thus attributing *knowledge* is doing three things (corresponding to the three elements of the justified true belief account of knowledge): attributing a commitment (corresponding to the belief condition), attributing an entitlement (corresponding to the justification condition), and *undertaking* the commitment myself (corresponding to the truth condition). These are all attitudes toward *statuses* (of commitment and entitlement), not toward *attitudes*. Since each interlocutor can adopt all these statuses, each one can take it that someone else is committed to something that, though they may be entitled to that commitment, is not *true*. That doing this requires comparing commitments one *attributes* to those one *undertakes* is the distinction of social perspective that is

the practical basis of the institution of a dimension of *objective* correctness into the assessment of commitments. In this way, each pre-ascriptional interlocutor has the wherewithal to distinguish in practice between what someone is committed to and what is correct (true). Ah, but since he does not by hypothesis yet have the explicit *concept* of truth or objective correctness (since, as Sellars put it, “grasp of a concept is mastery of the use of a word”, and he doesn’t have the use of the relevant ascriptional words), can he apply this distinction *to himself*? He can *in effect* distinguish between attitude and status for others (only “in effect”, since he does not have the concept of an attitude until he can not only attribute and acknowledge but ascribe statuses), but how can he do that in his own case? This will be possible if he can *attribute* to himself different commitments than he *acknowledges*. This is evidently possible when he changes his mind, alters the commitments he acknowledges; the distinction between past and present commitments admits the implicit distinction of social perspective between status and attitude. But can we make sense of the pre-ascriptional deontic scorekeeper applying his practical distinction between status and attitude to his own *current* commitments? I admit that it is hard to see how. And if such an interlocutor cannot do that, what are we to say of the practical conception of objectivity of conceptual content that he exhibits? Is it robust enough to meet Davidson’s condition on having anything recognizable as beliefs? Have we at this point thrown overboard what turns out to be an absolutely essential component of Neurath’s boat and left ourselves with nothing that deserves to count as a *discursive* (that is, concept-using) practice?

If I understand the thrust of Peter Grönert’s final concern it is that the account of the practical-inferential role of normative vocabulary in Chapter Four does not jibe with the use made of normative vocabulary in the account of deontic scorekeeping that is the framing theory within which that account is given. Since *MIE* aspires to expressive completeness — specifying in terms of the theory the conceptual contents of the expressions used in expounding the theory — this would be a failure to meet acknowledged criteria of adequacy. It is true that the account of normative vocabulary does not explicitly deal with ‘commitment’ (either of the doxastic or the practical sort), or ‘entitlement’, addressing instead various senses of ‘ought’. But the intent is that the sketch offered there can easily be extended to those cases. In particular, locutions such as “...is practically committed to...”, which is offered as functionally analogous to “...intends to...”, will be read so as to support the same sorts of inferences to ‘shall’ (which explicitly acknowledges practical commitments in a way hooked up to non-linguistic activity by reliable differential responsive dispositions) in

the first-person case, and 'should' (which attributes such commitments) in the third-person case. So the Chapter Four account does underwrite the inferences Grönert sees going missing. As for the connection to the language of the theory itself, a scorekeeper's acknowledging a commitment (whether doxastic or practical) is to have the same consequences for her behavior — consequences such as in turn acknowledging the inferential consequences of the claim she acknowledges commitment to — and that scorekeeper's attributing a commitment to another similarly has practical consequences for what she takes it that interlocutor should go on to do (is committed to going on to do). The as-it-were practical reasoning (only "as it were" since the scorekeeper need not have the explicating concept of *commitment*, and so is not explicitly engaging in practical reasoning) on the part of scorekeepers that is made explicit by the theorist's use of 'commitment' and 'entitlement' is just *consequential* scorekeeping: a matter of what further commitments or entitlements one *shall* or *should* acknowledge or attribute, given that one has acknowledged or attributed some others.

Carlo Penco's "Keeping track of individuals" concludes with the worry that the sort of semantic holism that inevitably goes along with inferentialism leaves us with conceptual contents that are, because of their responsiveness to speakers' collateral beliefs, ever-changing and only rarely and accidentally shareable. This is, of course, a kind of concern familiar already from discussions of the virtues and vices of the semantic holism Quine endorses late in "Two dogmas of empiricism". Two structural features of the way the account in *MIE* differs from Quinean holism address this issue. First, concepts are thought of as public norms by which speakers bind themselves by using words whose significance is not up to them. It is up to me whether I call the coin copper, but if I do, it is not up to me that I have committed myself to its melting at 1083.4° C — and I do not need to *know* that I have undertaken that commitment in order to have done so. Thus you and I may have quite different collateral beliefs about copper, and still be using the same concept, still binding ourselves by the same inferential norms when we use the word. The second feature of the account then concerns the relations between different speakers and hearers potentially idiosyncratic *understandings* of the common conceptual contents they deploy. Although the same things *do* follow from my calling something 'copper' and you calling it 'copper', we may be disposed to draw quite different conclusions from our claims, and to accept quite different sorts of evidence for or against them. It is these differences in social perspectives on the common content that are expressed explicitly in the *de re* and *de dicto* styles of propositional-attitude

ascription. The capacity to characterize the contents of claims from these two different perspectives is the explicitation of what is implicit in the fundamental practical ability to navigate between different doxastic points of view: to keep two sets of books on each interlocutor, and move more or less smoothly back and forth between what each one is *really* (whether she knows it or not) committed to, and what she merely *takes* herself to be committed to. The fact that the very same commitment can have its content specified either *de re* or *de dicto* shows that there is one conceptual content shared by both speaker and scorekeeper, but specifiable differently, depending on the doxastic perspective each occupies. The measure of the success of this perspectival account of what is and isn't shared by different interlocutors is the extent to which the account of *de re* and *de dicto* propositional-attitude ascriptions is adequate to *their* contents.

The principal sort of defective discourse that *MIE* addresses is concepts that are defective, according to a scorekeeper, in that the inference from their circumstances of application (committive or permissive) to their consequences of application (committive or permissive) is not a good one. The claim is that one of the basic functions of logical vocabulary is to make it possible for such implicit inferential commitments to be made explicit in the form of claims (paradigmatically, conditional claims), whose own evidential credentials can then be queried. Kevin Scharp thinks that when further kinds of defective discourse are considered, the counsel of wisdom is to adjust the basic model of scorekeeping in various ways, so as to incorporate at the ground-level the critical and conversational mechanisms by means of which we cope with such defects. I applaud his extension of the apparatus. It will be interesting to compare in detail what can be done in his terms with what logical vocabulary makes room for.

Author's address

Robert B. Brandom
Department of Philosophy
University of Pittsburgh
Cathedral of Learning
Pittsburgh, PA 15260
USA

rbrandom@pitt.edu
www.pitt.edu/~rbrandom/

About the author

Robert B. Brandom is Distinguished Service Professor of Philosophy, Fellow of the Center for Philosophy of Science at the University of Pittsburgh, and Fellow of the American Academy of Arts and Sciences. His books include *Making It Explicit: Reasoning, Representing, and Discursive Commitment* (1994), *Articulating Reasons: An Introduction to Inferentialism* (2000), and *Tales of the Mighty Dead: Historical Essays in the Metaphysics of Intentionality* (2002). He is the editor of *Rorty and his Critics* (2000), and the author of a number of important articles in different journals. Currently he is working on a long-planned book on Hegel's Phenomenology.