

International Phenomenological Society

Review: Knowledge and the Social Articulation of the Space of Reasons

Reviewed Work(s): Knowledge and the Internal by John McDowell

Review by: Robert Brandom

Source: *Philosophy and Phenomenological Research*, Vol. 55, No. 4 (Dec., 1995), pp. 895-908

Published by: International Phenomenological Society

Stable URL: <https://www.jstor.org/stable/2108339>

Accessed: 15-05-2020 19:21 UTC

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at <https://about.jstor.org/terms>



JSTOR

International Phenomenological Society is collaborating with JSTOR to digitize, preserve and extend access to *Philosophy and Phenomenological Research*

Knowledge and the Social Articulation of the Space of Reasons

ROBERT BRANDOM

University of Pittsburgh

In “Knowledge and the Internal” John McDowell presents a deep and interesting argument. I think everything he says is true and important. (Actually, I wouldn’t want that quantifier to be restricted to the claims he makes there; I’m prepared to make this endorsement quite generally—but I won’t try to defend *that* attitude here.) Still, there are a number of points that bear expanding on in order to be properly understood. So I want to say something about his point of departure: the idea of standings in the space of reasons. And I want to fill in further the picture at which he finally arrives, by saying how I think we *ought* to understand knowledge as a standing in the space of reasons, once we have freed ourselves from a prevalent deformed conception of that space. McDowell’s strategy is to show that that conception of the space of reasons is inadequate—that it deserves to be called a ‘deformation’—by showing that it leaves no room for anything recognizable as knowledge. I’ll try to reconstruct that argument by showing what it looks like in the context of a crucial dimension of the space of reasons that McDowell never mentions: its essentially *social* articulation. The effect of this supplementation, I think, is not to turn a bad argument into a good one, but to turn what is already a good argument into one that further illuminates the phenomena with which it deals.

The result may be just another confirmation of the maxim advising us that we ought to be more suspicious of philosophers who think they agree with us than we are of philosophers who think that they don’t. But depending on how successful I am, I hope at *worst* to clarify some of the key concepts and connections that McDowell appeals to, and at *best* to twist his words into a perverted caricature of their intended meaning. The game is worth the candle; for if I read him aright, when all the background he presupposes is made explicit, what he has offered us is nothing less than a generalized argument against all possible forms of epistemological externalism.

I

We can start with the Sellarsian idea that concepts are places in the space of reasons. According to this thought, to talk about the contents of beliefs and claims (the kinds of thing that are candidates for being or expressing knowledge) is to talk about things that can in principle be given as reasons, and for which reasons can in principle be asked.¹ One might think that in giving pride of place in this way to *justification*—by following Sellars in focusing to begin with on issues of what is a good reason for what—McDowell is begging important questions in the context of an investigation of what *knowledge* consists in. After all, it is the hallmark of the school of reliabilist externalism in contemporary epistemology precisely to deny that considerations of justification or what is a reason for what *need* have anything to do with assessments of knowledge. What distinguishes true beliefs that deserve to be called ‘knowledge’ from those that do not is just that they have not been arrived at haphazardly or accidentally (for instance, by coin-flipping). What matters is that they be the outcome of a *reliable* belief-forming mechanism—one whose output is *likely* to be true. Under the right circumstances, making appropriate inferential moves in the space of reasons can qualify as such a mechanism. But it has no privileged status: knowledge can be diagnosed quite apart from any consideration of the space of reasons.

Epistemological externalism of this extreme sort, which regards the space of reasons as an optional superstructure, is one of the four positions on the nature of knowledge that McDowell considers. Even though his main interest lies elsewhere (in the more moderate externalism that sees considerations lying outside the space of reasons as only one element in a hybrid view), it might be thought that his dismissal of extreme externalism with the scornful remark that according to such an approach there is no principled reason not to count thermometers as knowers is a bit cavalier. In fact this is just the right thing to say, and no important questions about knowledge are being begged here. Seeing why this is so will help us to see better the role that is being played by the Sellarsian notion of a space of reasons.

For the important point has nothing to do with what one thinks of the propriety of the traditional construal of knowledge as *justified* true belief. It has to do with how one distinguishes concept use from nonconceptual activity. What is the difference between a parrot who is disposed reliably to re-

¹ As I use the terms, what can serve as premises and conclusions of inferences are *propositional* contents, which is the fundamental and defining species of *conceptual* contents. Conceptual contents that are not propositional correspond to what is expressed by subsentential expressions, and are to be understood by abstraction from the propositional contents of sentences containing them—which is to say that the contribution a subsentential expression makes to the propositional contents expressed by sentences in which it occurs is to be identified by observing the effects on those contents of substituting other expressions for it.

spend differentially to the presence of red things by saying “Raawk, that’s red,” and a human reporter who makes the same noise under the same circumstances?² Or between a thermometer that responds to the temperature’s dropping below 70 degrees by reporting that fact by moving the needle on its output dial and a human reporter who makes a suitable noise under the same conditions? By hypothesis both reliably respond to the same stimuli, but we want to say that humans do, and the parrots and thermometers do not, respond by applying the *concepts red* or *70 degrees*. The parrot and the thermometer do not grasp those concepts, and so do not *understand* what they are ‘saying’. That is why we ought not to consider their responses as expressing *beliefs*: the *belief* condition on knowledge implicitly contains an *understanding* condition.

The Sellarsian idea with which McDowell begins is that this difference ought to be understood in terms of the space of reasons. The difference that makes a difference in these cases is that for the human reporters, the claims “That’s red,” or “It’s 70 degrees out,” occupy positions in the space of reasons—the genuine reporters can tell what follows from them and what would be evidence for them. This practical know-how—being able to tell what they would be reasons for and what would be reasons for them—is as much a part of their understanding of ‘red’ and ‘70 degrees’ as are their reliable differential responsive dispositions. And it is this *inferential articulation* of those responses, the role they play in reasoning, that makes those responsive dispositions dispositions to apply *concepts*. If this idea is right, then nothing that can’t move in the space of reasons—nothing that can’t distinguish some claims or beliefs as *justifying* or being reasons for others—can even count as a concept user or believer, never mind a knower: it would be in another line of work altogether. And this point is not touched at all by the important observation that something that *is* in this line of work—something that *can* use concepts and have beliefs, something, that is, that can find its way around the space of reasons—can count as having knowledge in particular cases in which it has a true belief that it is not in a position to give reasons for. Extreme, or as I will henceforth feel entitled to call them, *gonzo* externalists mistakenly infer from the fact that issues of justification and reason-giving can be treated as *locally* irrelevant to attributions of knowledge in such cases, that they can safely be treated as *globally* irrelevant. The problem with this form of externalism is not with its construal of the *justification* condition on knowledge, but with its construal of the *belief* condition on knowledge.

These same terms give us some clues as to how we might think about the notion of *standings* in the space of reasons. A typical twenty month old child who toddles into the livingroom and in bell-like tones utters the sentence “The house is on fire,” is doing something quite different from what his

² Perhaps minus the parrot expletive.

seven year old sister would be doing by making the same noises. The young child is not claiming that the house is on fire, for the simple reason that he does not know what he would be committing himself to by that claim, what he would be making himself responsible for. He does not know what follows from it, what would be evidence for it, what would be incompatible with it, and so on. He does not know his way around the space of reasons well enough yet for anything he does to count as adopting a standing in that space. His older sister knows that it follows from her claim that the family is in danger and should flee, and that the kitchen's being full of smoke and flame is evidence for it. She can commit herself, and knows what she would be committing herself to and what might entitle her to that commitment. She has begun to master the *inferential* articulation of such potential positions, statuses, or standings that make up the space of reasons—the things that can stand in the relation *is a reason for* to each other.

In order to clarify McDowell's argument and its conclusion, I'm going to recast them in an idiom he does *not* use: I'll talk about standings or statuses in the space of reasons in terms of two fundamental categories: *commitments* of a certain kind, and *entitlements* to those commitments. The idea is that occupying the basic sort of standing in the space of reasons is *staking a claim*, that is, undertaking a commitment of the sort that might be expressed by making a claim or assertion. Presystematically we might think of these as commitments to the truth of various propositions, that is, as beliefs. But I think it will be helpful if we keep talk of truth, propositions, and beliefs off-stage for a while. To uphold the fundamental Sellarsian idea about what would be required for these standings to have conceptual content, we must think about them as having two properties. First, it must be part of the conception of these commitments that the issue of one's *entitlement* to such a commitment can arise. Second, it must be possible for one such commitment to *inherit* or *derive* its entitlement from another. Together these mean that commitments can both serve as and stand in need of reasons. That is the sense in which they are being taken to be standings *in the space of reasons*.

The final point I want to make about McDowell's Sellarsian starting-point is that the "space of reasons" that he discusses ought to be understood as an abstraction from concrete practices of giving and asking for reasons. The space of reasons is a normative space. It is articulated by proprieties that govern practices of citing one standing as committing or entitling one to another—that is, as a reason for another. What people actually *do* is adopt, assess, and attribute such standings—and if they did not, there would be no such standings. For in the absence of such normative attitudes of taking or treating people as committed or entitled, there are no commitments or entitlements. They are not part of the furniture of the prehuman world.

When we turn to consider McDowell's diagnosis of deformations in our conception of the space of reasons that threaten to make knowledge unintelligible, it will be useful to keep our eyes on the actual practices of giving and asking for reasons, the practices that give a point to the abstract notion of a space of reasons.

II

McDowell's argument is structured by a botanization that classifies approaches to knowledge as coming in four flavors: sceptical, dogmatic, hybrid, and extreme externalist. We can group these further according to whether they conceive justification and truth as internally or externally related—or as I will say, according to whether they *aggregate* or *segregate* these conditions. The sceptic and the dogmatist take it as a criterion of adequacy on a notion of justification that any claim or belief that is sufficiently justified is true. They are right that if a claim or belief has the status of knowledge, it is guaranteed to be true. But they also take it that justification of a certain sort is what distinguishes knowledge from other belief. If that is right, then justification must be truth-guaranteeing. While *agreeing* on this basic principle, the sceptic and the dogmatist *disagree* about whether a notion of justification meeting this condition is to be had. The sceptic arrives at the *false* conclusion that knowledge is *not* possible by combining the *false* claim that justification must be incompatible with falsehood with the *true* claim that justification that rules out the possibility of falsehood is *not* to be had. The dogmatist arrives at the *true* conclusion that knowledge *is* possible by combining the *false* claim that justification must be incompatible with falsehood with the further *false* claim that justification that rules out the possibility of falsehood *can* be had. McDowell rightly does not rehearse at length the difficulties of these views; their unsatisfactoriness is widely acknowledged.

Where scepticism and dogmatism run the justification and truth conditions on knowledge too closely together, the hybrid and extreme versions of externalism drive them too far apart. Gonzo externalism throws justification and the giving of reasons out entirely, and I've already indicated why this won't do.³ It is by no means obvious, though, why the more moderate externalism

³ I said above that McDowell offers us a generalized argument against any possible form of epistemological externalism. (It might better be thought of as a recipe that, given any such externalist view, shows us how to construct a knock-down argument against it.) In conversation, McDowell points to one case that might be thought to be an exception. Non- or prelinguistic animals do not have status or standing in the space of reasons. So according to the idiom being recommended here, they neither deploy concepts, acquire beliefs, nor count as having knowledge. Nonetheless, it is common to talk about them loosely as though they were capable of some version (usually admitted to be degenerate cases) of these accomplishments. The informational states most closely resembling genuine beliefs that they *do* have (call them *beliefs**), when they both correctly represent how things are and are acquired by a suitable reliable process

of a two-factor or hybrid view can't be made to work. McDowell's core argument is accordingly devoted to showing what is wrong with them. The danger he sees is that if a satisfactory standing with respect to justification or reason-giving is seen as an internal matter, something one can secure all on one's own, while assessments of truth or reliability answer to external standards, then the justification and truth conditions on knowledge are treated as independent of one another. But being justified in holding a belief just is being justified in taking it to be *true*. Such segregationist views, he argues, are inherently unstable and untenable.

Although his four-part classification lumps them together, McDowell implicitly acknowledges two different forms hybrid views can take. They deserve to be considered together because each *extrudes* from the space of reasons some sort of assessment that is crucial to the attribution of knowledge, segregating it as external to standings in that space. The first version extrudes both considerations of truth and of reliability as statuses distinguished from that of being justified. That is, assessments of what is *true*, and of the *reliability* of various policies for endorsing some claim as true (believing it), are taken to be independent in principle from questions of what is a reason for what. McDowell rightly gives this sort of segregationism short shrift: nothing recognizable as our notion of justification survives if our justificatory practices are forbidden in principle from being criticized and shaped on the basis of assessments of their reliability, that is, the likelihood that reasoning in the ways they sanction will lead to truths. This recognition of the intimate connection between justification and reliability motivates the second version of hybrid segregationism, which seeks to incorporate assessments of reliability into standings in the space of reasons, while still extruding truth. But since the reliability of a belief-endorsing policy just is the likelihood that it will lead to the endorsement of truths, this version of the hybrid approach is no more stable a position than the other.

The underlying fact is that the notions of belief, justification, reliability, and truth are inextricably intertwined in ways that preclude the segregationist approach. I started by talking about how the notion of belief, as a conceptually contentful state, is unintelligible apart from considerations of what is a reason for what—that is, apart from liability to assessments of *justification*. But it is equally essential to our notion of beliefs that they are something for which the question of *truth* can arise: believing *is* taking or treating as *true*. Any adequate account of the contentfulness of beliefs and claims must show why and how these two crucial dimensions of assessment are so intimately linked. As McDowell's argument indicates, a useful way to unpack that link-

may be called *knowledge**. An externalist account of *this* sort of state is all that is to be had. This status has in common with the genuine article what the parrot has in common with the reporter of red things: reliable differential responsive dispositions.

age is by looking at the concept of *reliability*. For, on the one hand, assessment of cognitive reliability makes sense only against a background that includes assessments of truth: a belief-endorsing policy is reliable just insofar as it is likely to lead to truths. And on the other hand, assessments of justification must answer to assessments of reliability. Arguing that a proposed method of justification is not likely to lead to truths is not just *one* way of criticizing such methods—any more than stopping the heart is just one way of killing a vertebrate; it is the form all those ways share, the common conclusion they must lead to if they are to be successful. That is why McDowell can put such pressure on segregationist accounts of knowledge by focusing attention on the notion of reliability: the hybrid approaches are unstable because they can neither adequately construe the space of reasons independently of considerations of reliability, nor adequately construe reliability assessments apart from truth assessments.

We can think of McDowell's argument as coming in three nested parts. The core argument is the one I've just rehearsed, which uses the notion of reliability to underscore that considerations of truth cannot be extruded from the space of reasons. This argument contributes the crucial premise in a wider argument, to the effect that *none* of the four currently available approaches to knowledge is satisfactory: neither scepticism, nor dogmatism, nor hybrid theories, nor extreme externalism. The segregationist approaches drive the justification and truth conditions on knowledge too far apart—in the case of gonzo externalism, at the cost of losing sight of what distinguishes beliefs and claims as conceptually articulated, by ignoring the space of reasons altogether. But the two aggregationist approaches are equally unsatisfactory, running justification and truth together so that their distinctive contributions to knowledge assessments are confounded. The third part of McDowell's argument is the claim that these four approaches share a presupposition concerning the shape of the space of reasons, and that given that presupposition, they exhaust the alternatives. If that is right, then the unsatisfactoriness of these ways of construing knowledge shows that that presupposition must be rejected, and the space of reasons otherwise understood.

III

I said at the outset that I think this is a good argument. I've now indicated how I think the first two steps work: the argument for the instability of hybrid views and, based on it, the argument for the unsatisfactoriness of any of the four kinds of epistemological segregationism and aggregationism he distinguishes. Two large issues remain: the diagnosis of these approaches as generated by a shared erroneous conception of the space of reasons, and the recommendation of an alternative. It is at this point that I would like to offer

what I regard (though McDowell may not) as a friendly amendment to or clarification of his account.

Early on in his paper, McDowell gives us the following characterization of the conceptual pathology he takes to have generated the shuttling back and forth between unsatisfactory positions (whether aggregationist or segregationist) characteristic of contemporary and classical epistemology:

The deformation is an interiorization of the space of reasons, a withdrawal of it from the external world. This happens when we suppose that we ought to be able to achieve flawless standings in the space of reasons by our own unaided resources, without needing the world to do us any favors.

Now I don't want to disagree with this, but I do want to insist that this diagnosis should come at the end of a story, not at the beginning. You may have noticed that although here and there I helped myself to McDowell's imagery of what is conceived as internal or external to the space of reasons, in my exposition of his core arguments I did not find it necessary to say *anything at all* about *interiorizing* the space of reasons in this sense—did not need to talk at all about what candidate knowers are supposed to be able to do *all on their own* as opposed to what they can only do by grace of *favors from the world*. The deformed and defective conception of the space of reasons that I see as underlying the various forms of epistemological aggregationism and segregationism in play so far is prior to, and, I want to argue, explanatory of, the one McDowell focuses on.

For I want to claim that the mistake is to begin with to *individualize* the space of reasons. The complaint I want to make about McDowell's discussion is that he makes nothing of the essential *social* articulation of that space. The passage above is typical: he says the thought is that *we* ought to be able to achieve flawless standings in the space of reasons by *our* own resources, without needing the world to do *us* any favors; but for all he says here or elsewhere, this *us* could be *each* of us, individually or by ourselves, rather than *all* of us collectively. But this difference makes all the difference.

The best way I know to make this clear is to try to indicate in more detail than McDowell does just what an account looks like that *does* construe factive statuses such as knowledge as standings in the space of reasons, while respecting the lessons of externalism. I said above that it is important to remember that our abstract talk about reasons and the space of reasons has to be grounded in an appreciation of the concrete practices of giving and asking for reasons, namely in what people actually *do*. I also suggested that what McDowell calls *standings* in the space of reasons should be thought of in terms of *commitments* and *entitlements* that are practically acknowledged by those engaging in such practices. What I want to claim now is that if we recognize that giving and asking for reasons is a constellation of essentially *social* practices, and that the commitments and entitlements those practices

involve are accordingly essentially *social* statuses, we will be in a position to understand factive locutions such as *knowledge* and warrantive locutions such as *reliable* as attributing standings in the space of reasons. Furthermore, we will be able to make appropriate sense of the different roles of assessments of truth and of justification in attributions of knowledge, as the aggregationist approaches of the sceptic and the dogmatist could not, without disjoining those roles so severely as to engender the problems we saw with the various segregationist approaches, in particular the embarrassment that hybrid theories have concerning the notion of reliability.

The key to understanding knowledge as a standing in the space of reasons is to focus on the practical attitude adopted by one who is assessing a candidate for such a standing: What is someone who attributes knowledge *doing*? For these purposes we can continue to be guided, as we have throughout, by the traditional conception of knowledge as justified true belief. Construed as a standing or status, belief will correspond to some sort of commitment, while justification (being justified) will correspond to some sort of entitlement to that commitment. So taking someone to have a justified belief will be understood as attributing two sorts of standings: a commitment and an entitlement.

What about the truth condition? To take someone to have the status of a knower one must take it that the justified belief in question is also *true*. What is it to do that? Taking the belief in question to be true is not a matter of *attributing* a commitment, but of *undertaking* one—endorsing the claim oneself. For taking-true is just believing, that is, committing oneself, adopting a standing or status.⁴ What sort of case leads us to distinguish justified beliefs that are true from those that are not? If you are standing in a darkened room and seem to see a candle ten feet in front of you, I may take you to have good reason for believing that there is a candle ten feet in front of you, and so to take you to be entitled to your commitment. But that may be my attitude even if I know, as you do not, that there is a mirror five feet in front of you, and no candle behind it, so that I am not in a position to endorse or commit myself to what you are committed to.

Thinking of things this way, assessing someone as having successfully achieved the status or standing of a knower involves adopting three different attitudes: *attributing* a commitment, *attributing* an entitlement, and *undertak-*

⁴ The root notion of truth is just what the tradition always took it to be: saying of what is that it is. The cases in which I take it that *p*, that is where I believe or am committed to the claim that *p*, are just the cases in which I take your belief that *p* to be true. The mistake of metaphysical conceptions of truth, including any substantive correspondence theory, is to assimilate what I am doing when I take your belief to be true to what I am doing when I take you to believe it, or to be justified in doing so. For in those cases I *attribute* a commitment and an entitlement, respectively. Metaphysical theories of truth are theories of the property I am attributing to your commitment when I take it to be true. But in taking it to be true I am not attributing *any* property to that commitment, I am endorsing it myself.

ing a commitment. There is nothing in principle mysterious about such assessments, nor, therefore, about the standing being assessed. Knowledge is intelligible as a standing in the space of reasons, because and insofar as it is intelligible as a status one can be taken to achieve in the game of giving and asking for reasons. But it is essentially a *social* status, because it incorporates and depends on the *social* difference of perspective between *attributing* a commitment (to another) and *undertaking* a commitment (oneself). If one *individualizes* the space of reasons, forgetting that it is a *shared* space within which we adopt attitudes towards *each other*—and so does not think about standings in the space of reasons as socially articulated, as potentially including the social difference of perspective between attributing and undertaking commitments, that is, between your standing and mine—then one will not be able to understand knowledge as a standing in the space of reasons. One will then have either to try to get some individualized standing to do the work of the socially articulated factive, as the aggregationists do, or to extrude some components of it from the space of reasons entirely, as the segregationists do. One will then be doomed either to lose the crucial distinction between the belief and justification conditions, on the one hand, and the truth condition on knowledge, on the other, as the aggregationists do, or to lose the crucial connections between them, as the segregationists do.

The distinction of social perspective between attributing a standing and adopting it keeps the truth condition from being run together with the others, and so makes it possible to understand assessments of something as having the standing of a justified belief that is not true. But McDowell's core argument indicates that the danger of distinguishing these elements too firmly—losing the crucial connections—manifests itself in difficulties with notion of *reliability*. What can we say about this test case?

A fundamental point on which broadly externalist approaches to epistemology are clearly right is that one can *be* justified without being *able* to justify. That is, one can have the standing of being *entitled* to a commitment without having to *inherit* that entitlement from *other* commitments inferentially related to it as reasons. A paradigm case is that of noninferential reports. If you are a generally reliable noninferential reporter of lighted candles in darkened rooms, then you can be entitled to your claim that there is a candle in front of you in cases where that claim or commitment arose by your exercise of that reliable differential disposition to respond to such candles by making such reports (undertaking such commitments, adopting such standings). And this can be the case even if you are not able yourself to cite your reliability in such matters as a reason for the belief you acquired.⁵

⁵ Where reliabilist counterexamples undermine the *necessity* of the JTB account of knowledge, Gettier-style counterexamples undermine its *sufficiency*. Justifications that depend essentially on false premises (even ones the candidate knower is entitled to

Now it would be wrong to conclude from the fact that a piece of knowledge can be acquired noninferentially even where the knower is unable to justify it that reasons need not be in play at all. For to begin with, you must be capable of making the claim or acquiring the belief in order to be a candidate for knowing it. And that requires that you *understand* it: that you have at least a rough practical mastery of its inferential role, the know-how to discriminate some things that follow from it from others that don't, and some things that would be evidence for it from others that would not. In "Empiricism and the Philosophy of Mind" Sellars unfortunately takes it that in order to secure *this* claim, he must insist that one is not justified unless one *knows* one is justified—in particular, that noninferential reports should be accorded the status of knowledge only in cases where the knower can cite her own reliability as a reason, from which the correctness of the noninferential report could be inferred. This response is excessive; there is no reason to deny the externalist insight that, once one is capable of achieving standings in the space of reasons—for instance capable of committing oneself to the claim that there is a candle in the room—one can become entitled to such standings without being able to give reasons for them. But Sellars' overreaction also contains an important insight: reliability matters to assessments of knowledge precisely because of the *inferences* it can support.

The key point to understanding reliability as a warrantive standing in the space of reasons is that the notion of reliability itself is essentially an *inferential* notion: a matter precisely of what is a reason for what. What must be kept in mind if one is to talk (a variant of) the traditional language of justification as *internal* entitling and reliability as *external* entitling is that what they are internal or external *to* is not the practice of giving and asking for reasons, and so not the space of reasons, but rather the individual whose standings in that space are being assessed. For reliability is precisely a matter of a *socially* articulated inference. For me to take you to be a reliable reporter of lighted candles in darkened rooms is just for me to endorse a particular pattern of reasoning; in particular it is for me to endorse the inference that could be made explicit by saying:

If in a darkened room S noninferentially acquires the belief that there is a lighted candle, then (probably) there is a lighted candle there.

believe) may be assessed as inadequate support for attributions of knowledge. Although the point cannot be pursued here, this phenomenon can be accommodated in the social-perspectival framework presented here, by looking at the relation between what would be reasons for the one to whom a candidate piece of knowledge is attributed, on the one hand, and what would be reasons for the one who attributes them (from whose point of view the truth of those premises is assessed).

Translated into the language I have suggested for discussing standings in the space of reasons (that is, statuses one can acquire in the game of giving and asking for reasons) this is an inferential connection between a suitably noninferentially acquired commitment *attributed* to you and a corresponding commitment that *I undertake*. It is treating your commitment as a (defeasible) reason for my own.

The externalist epistemologist who takes reliability to warrant the attribution of noninferential knowledge in the absence of justification relies precisely on this essentially *interpersonal* pattern of inference. If we like, we can say that Sellars' point is reinstated at one remove of social perspective: although it is enough that the subject of knowledge *be* reliable to be entitled to a belief (without having to be able to *cite* that reliability as a reason for it), the *attributor* of knowledge must be able to cite that reliability as such a reason.⁶

My conclusion is that if we keep firmly in mind that the space of reasons is founded on practices of giving and asking for reasons—practices in which standings or statuses can not only be adopted but attributed—then we can understand truth and reliability, no less than justification, and hence in the end even *knowledge*, as socially articulated standings in that space. According to such a picture, the insights of externalism are accommodated as pointing to features of the essentially social practices of giving and asking for reasons, undertaking and attributing inferentially articulated commitments and entitlements. So issues of justification, on the one hand, and of truth and reliability, on the other, are not severed from one another, and the instability that McDowell diagnoses in what he calls 'hybrid' views is avoided. Knowledge and reliability, involving as they do essential reference to truth, *are* in a certain sense *hybrid* statuses on this account. For they are made intelligible by appeal to *two* different social perspectives, that of the one to whom a status is *attributed* and that of the one *attributing* it.⁷ But the crucial difference is that this is a distinction of perspectives *within* the space of reasons, not a distinction between what is within it and what is without it. That is why the instability McDowell point to does not arise.

⁶ Assuming sufficient expressive resources are available to formulate the reliability inference explicitly in a conditional (which makes it available in a form suitable to serve as a premise in further inferences).

⁷ When I talk about "the social articulation of the space of reasons" I mean that standings in that space must be understood in terms of these two kinds of socially related perspective. I do *not* mean that the community is privileged in some way relative to individuals. So it would be a mistake to think of my remarks as suggesting a super-individual sort of interiorization, in which the community as a whole plays the role formerly played by particular individuals.

IV

My response to McDowell's paper has come in two parts, one constructive, and one critical. Constructively, I have indicated how knowledge can be construed as a standing in the space of reasons. The key question is what I must be *doing* in order to take you to have that standing. And the answer is, in line with the JTB account of knowledge, that corresponding to the belief condition (which includes an understanding condition) I must *attribute* a propositionally contentful commitment, that corresponding to the justification condition I must *attribute* also entitlement to that commitment (whether inferentially or noninferentially grounded), and that corresponding to the truth condition I must also myself *endorse* or *undertake* the same propositionally contentful commitment. The possibility of such *truth* assessments is already implicit in the attribution of a commitment corresponding to belief—for propositionally contentful commitments are essentially, and not merely accidentally, things for which the question of truth can arise. Likewise, and for that reason, the possibility of *reliability* assessments is already implicit in the attribution of an entitlement corresponding to the justification condition on knowledge. But assessing your reliability is a matter of whether to endorse a certain pattern of interpersonal inference: the inference from a commitment I *attribute* to you to one I *undertake* myself. In this way what is expressed by the use of *factive* locutions such as 'believes truly', *warrantive* locutions, such as 'believes reliably', and *cognitive* locutions such as 'knows', which include both factive and warrantive dimensions, can all be understood as standings in a socially articulated space of reasons: standings that incorporate what are with respect to *individual* knowers internal and external epistemic considerations in the form of the distinct *social* perspectives of attributing and undertaking commitments.

It is the different perspectives provided by different sets of commitments that make it possible to *triangulate* on objective states of affairs. Our practices of comparing, assessing, and correcting different repertoires of commitments one with respect to another—those we attribute to others and those we undertake ourselves—are what make them intelligible as *perspectives*, views of something, ways in which a perspective-independent reality can *appear*.⁸ To individualize the space of reasons is to interiorize it. To ignore the social articulation of standings in the space of reasons is to leave out what makes it possible to understand such standings as answerable for their correctness to how things actually are. And such an interiorized rendering must in the end fail, as McDowell insists, even to be recognizable as *belief*. For what an individualized construal leaves out is what makes statuses such as knowledge

⁸ I develop this view, and the other constructive suggestions offered in the second half of this paper, in greater detail in *Making It Explicit: Reasoning, Representing, and Discursive Commitment*, Harvard University Press, 1994.

and reliability intelligible as standings in the space of reasons. Factives, like *believes truly* (because of their relation to truth-assessment) warrantives, like *believes justifiedly* (because of their relation to reliability), and so cognitives, like *knows*, testify at once to the way in which objective facts (concerning how things really are, not just how they are taken to be) are incorporated in the space of reasons, and equally how the social articulation of that space makes such incorporation so much as intelligible.

My only complaints against McDowell have been accusations of sins of omission—a matter of what he has *not* said. Such complaints are often unfair: one can't say everything. But I have not reproached him for saying nothing about the effects of the discovery of silver in the New World on the spread of the Hussite heresy in Central Europe, even though he has indeed been silent regarding this important topic. I have reproached him for saying nothing about the social articulation of the space of reasons in the context of a discussion of a deformed conception of the space of reasons that makes it impossible for us to understand how knowledge and reliability are related to such standings. For that disastrous interiorizing of the space of reasons results precisely from individualizing it.

So let me end as I began: everything McDowell says is true and important—but sometimes he leaves stuff out.