International Journal of Philosophical Studies Vol. 15(1), 23–51

Does Observational Knowledge Require Metaknowledge? A Dialogue on Sellars

Timm Triplett and Willem deVries

Abstract

In the following dialogue between TT – a foundationalist – and WdeV – a Sellarsian, we offer our differing assessments of the principle for observational knowledge proposed in Wilfrid Sellars's 'Empiricism and the Philosophy of Mind'. Sellars writes: 'For a *Konstatierung* "This is green" to "express observational knowledge", not only must it be a *symptom* or *sign* of the presence of a green object in standard conditions, but the perceiver must know that tokens of "This is green" *are* symptoms of the presence of green objects in conditions which are standard for visual perception.' In the ensuing dialogue, TT argues that it sets the bar too high when knowledge *about* perceptual conditions is required for ordinary observational knowledge – that young children, for example, are implausibly excluded as knowers given Sellars's principle. WdeV defends Sellars's metaknowledge requirement against these charges. Results from developmental psychology are surveyed for what they show about the actual capabilities of young children. The implications of these results for the success of Sellars's principle are debated.

Keywords: Wilfrid Sellars; metaknowledge; observational knowledge; level ascent requirement; genetic epistemology; developmental psychology

1 Must Sellars Deny Observational Knowledge in Young Children?

TT: Given what we know about the perceptual ideas of young children – say typical 4-year-olds – it looks straightforward to me that Sellars has to deny them observational knowledge. Although Sellars includes a reliabilist condition in his theory of knowledge, he also inserts what William Alston has called a 'level ascent' requirement,[1] which stipulates that knowers must have some degree of epistemological sophistication, enough to have some basic metaknowledge. Remember what Sellars says:

> For a *Konstatierung* 'This is green' to 'express observational knowledge', not only must it be a *symptom* or *sign* of the presence of

International Journal of Philosophical Studies
ISSN 0967–2559 print 1466–4542 online © 2007 Taylor & Francis
http://www.tandf.co.uk/journals
DOI: 10.1080/09672550601143086

> a green object in standard conditions, but the perceiver must know that tokens of 'This is green' *are* symptoms of the presence of green objects in conditions which are standard for visual perception.
>
> (*KMG*, p. 247, *SPR*, p. 168)[2]

Let's call this Sellars's Principle of Observational Knowledge (SPOK for short). Now, it certainly seems that a prima facie case can be made for denying that a typical 4-year-old knows that 'tokens of "This is green" are symptoms of the presence of green objects in conditions which are standard for visual perception'. She would seem to lack the sophisticated epistemic concepts required for such knowledge. But, given SPOK, she would then also lack any knowledge even of the observational sort. Denying such knowledge to a 4-year-old would seem to put Sellars in a pretty implausible position. Does he indeed have to deny the young child knowledge, or can something be said in support of the claim that Sellars's requirements for observational knowledge are consistent with a young child knowing that grass is green?

2 A Sellarsian Response to Alston's Criticisms of the Level Ascent Requirement

WdeV: In effect, this is a specific example of one of Alston's major criticisms of Sellars. But Alston's criticism is as complex as it is important.[3] SPOK does indeed include a level ascent requirement: anyone who is a knower must not only be an observer and a reporter, she must be an epistemic agent, with working concepts of symptoms, evidence, standard perceptual conditions, and so on. Alston thinks that this is far too strong a requirement, for precisely the kinds of reasons you've just illustrated.

The level ascent requirement plays a fairly complex role in Sellars's epistemology. He starts with a fundamentally externalist approach to knowledge and corrects it by adjoining a level ascent requirement that entails that observers, if they are to have observational knowledge, must also have knowledge of:

1 the relevant empirical generalizations concerning the connection between responses and the conditions under which they are made;
2 the relevant empirical facts concerning their particular situation;
3 the relevant norms governing the epistemic evaluation of the situation.

This level ascent requirement is absolutely crucial to Sellars's epistemology.[4] It is what rules out the possibility of a level of basic knowledge independent of other empirical knowledge and thus rules out foundationalism, traditional empiricism, and reliabilism. Furthermore, it entails that knowledge must be

acquired wholesale. For particular observation knowledge can be acquired only if the subject already has a great deal of other knowledge, both general and particular. Sellars is barred from using the traditional picture that knowledge begins with simple, basic cognitions and grows by the continued accretion of new bits of basic knowledge plus inferential extensions of the accumulated knowledge base. The level ascent requirement is the linchpin of Sellars's epistemic holism.[5] Yet, important as the level ascent requirement is, Sellars never offers a clear defence of it.

Alston suggests two different mistakes Sellars might be making. One is that Sellars has confused the activity of *justifying* a belief – showing or demonstrating the belief to be justified – with the different concept of a belief's *being* justified, 'where this is some kind of epistemic state or condition of the believer vis-à-vis that belief, rather than something he is or might be *doing*'.[6] If Sellars were committing this error, he would illegitimately have transferred characteristics of the *activity* of justifying a belief (which one does by giving reasons for it) to the *state* of a believer. One must have overt and explicit reasons for a belief in order to engage in the activity of justifying it. But that does not entail that in order for a belief to *be justified* one must have any overt or explicit reasons for it. The mistake would be comparable to thinking that *S*'s knife can be sharp only if *S* sharpens it.

Alston does not dwell on this charge, however, and seems instead to hold that, like a number of other philosophers, Sellars really thinks that only critically reflective thinkers can know anything. The level ascent requirement is, on Alston's view, an expression of this unwarranted prejudice. Alston concludes that the level ascent requirement is not necessary for either justification or knowledge.

Alston's criticism, as I understand it, hinges on his belief that Sellars requires the knower to have performed an actual, explicit level ascent movement at some point in order to be justified in her observational beliefs. I think that this is wrong. What Sellars requires is that the knower be such as to be in a position to make such moves, even if only implicitly.

TT: I'm not sure what you mean by an implicit level ascent move.

WdeV: Consider the idea of the knowledge game, understood as a rule-governed activity in which humans participate from a fairly early age. Moves in the game include making a claim as well as defending, questioning, reconsidering it, and so on.

Obviously, most adult humans play the game and do so at least occasionally at a conscious level. When they doubt testimony or their senses, when they defend their claims, they are playing it overtly. Just as we can participate in the complex, rule-governed activity of language in such a way as to show that we are sensitive to the norms constituting language yet without realizing that it is a complex, rule-governed activity,[7] so are we able to join

in the knowledge game without constantly having to consult general epistemic principles.

TT: Attributing knowledge to infralinguals and animals, however, is very problematic on this view. Without an overt language, it is difficult if not impossible for them to engage in most of the behaviours that would constitute the clearest evidence of being players in the knowledge game. And given his thoroughgoing psychological nominalism in 'Empiricism and the Philosophy of Mind', Sellars seems bound to deny knowledge to the languageless.[8] They can't have even observational knowledge, because such knowledge and the conceptual sophistication it requires are invariably linguistic affairs.

WdeV: Sellars does, in his later works, appear to take a more tolerant attitude, in which he allows internal representational states to infralinguals (though not in exactly the same sense that linguistic beings possess such representations). Be that as it may, the issue here is that you and Alston think that, in 'Empiricism and the Philosophy of Mind', Sellars denies observational knowledge even to some relatively sophisticated linguistic beings like 4-year-old children. And this, I think I can now show, is clearly a misinterpretation.

For if we understand epistemic status and epistemic abilities in terms of the knowledge game outlined above, then the 'confusion' that Alston points to between the act of justifying a belief and the belief's being justified is no longer so straightforward. On Sellars's treatment, a belief's being justified, as a state, is still a matter of the status of that belief state in the knowledge game. Saying that belief *b* is justified for subject *S* is tantamount, according to Sellars, to saying that *b* is a position in the knowledge game to which *S* is entitled. (Note that so far, there is nothing in this claim that is incompatible with foundationalism. It is Sellars's further claim that one's entitlement to a particular [set of] position[s] in the knowledge game always depends on what other positions one occupies that excludes foundationalism by ruling out the idea that one could be entitled to some positions independently of what other positions one occupies.[9])

Under what conditions can we construe *S*'s state or belief as occupying a particular position in the knowledge game? We can do so only if we can construe at least some of *S*'s state-transitions as being legal moves within the game, sensitive to the norms and rules governing the game. Thus, even the state of being justified must be understood in terms of a role, a role that still has its primary mode of being in justificatory argumentation such as is commonly found among humans. When abstracted entirely away from its roots in justificatory argumentation, the notion of justification loses its grounding entirely and is no longer a concept we have good control of. This does not require that we can apply the notion of justification to some being's states or beliefs only if that being is a highly sophisticated logician. It does

mean that if the argumentative capacities of the being are nil, we have taken the notion of justification out of its natural context and stretched its use in a way that may be useful, but that we should be cognizant of as well.

Furthermore, Alston's objection that according to Sellars, *S* knows that *p* only if *S* has the capacity to demonstrate that *S* is justified in believing that *p* can now be seen to have missed its mark. Being able to demonstrate that one is indeed justified in believing that *p* (for almost any *p*) does require a great deal of cognitive sophistication – more than most people can command for the majority of their (justified) beliefs. It is clearly too strong a condition. But there is no evidence that Sellars intends to impose such a strong condition. He does require that the subject *S* know the relevant empirical generalizations and the relevant particular empirical truths about the context – but neither of these imposes an unsupportable burden on the subject of empirical knowledge. Although the subject must *know* these other truths, it is not necessary that they be expressly thematized by the subject. We all know that boa constrictors don't play basketball in the jungle, but most people have never explicitly considered the matter before.

All Sellars requires for a belief *b* of subject *S* to be observational knowledge is that (1) *S* be able to play the knowledge game, (2) *S* occupy appropriate other positions $b_1 \ldots b_n$ (some of which are also observational knowledge, some of which are general knowledge) such that occupying *b* is justified for *S*, and (3) *b* be appropriately caused as part of a reliable mode of belief acquisition.

The requirement that *S* be able to play the knowledge game is not a requirement that *S* be a Grand Master at it, constantly critically reflective on and explicit about her conduct of the game. But it does seem to require that *S* be sensitive to the norms of the game as norms (perhaps as the 6-year-old is sensitive to the rules of grammar as norms – that is, as a basis for assessing moves as correct or incorrect), without necessarily being able to formulate them or make them at all explicit. Much of the subject's knowledge of the game and its rules may well itself be tacit and unconscious. The subject exhibits her knowledge of the rules of the knowledge game in being able to respond appropriately (often enough) to challenges to her claims, in challenging the claims of others at appropriate times, and so on, even if she can't formulate at all explicitly the rules she is tacitly relying on in such behaviour. This is why I don't think Sellars would have to say that the 4-year-old doesn't know that grass is green, or that this is grass, and so on. True, she doesn't have explicit concepts of evidence and would almost certainly not engage in an explicit justification of her assertions by ascending levels, but she is nonetheless quite sensitive to many of the rules of the knowledge game. Certain people and certain experiences in certain circumstances do have greater authority for her than others; she can defend her claims and attack the claims of others, if only in unsophisticated ways; she has learned the concepts of green and

grass, and knows that she has. The 4-year-old may be a rookie at it, but she's still playing the knowledge game.

TT: This is certainly an interesting response to Alston. But, as you've been filling it out, my questions or objections have been accumulating. In the present dialogue, I'd like to explore your attempt to defend SPOK in light of the objection that Sellars's principle sets too high a standard for attributions of knowledge.[10]

3 How Plausible is the Level Ascent Requirement in Light of Data from Developmental Psychology?

TT: I wonder whether your talk of the knowledge game doesn't significantly weaken the conditions Sellars proposes for observational knowledge. For the data from developmental psychology strongly suggest that young children don't satisfy Sellars's requirements for observational knowledge. If that is right, and if the young child is a knower under your 'knowledge game' requirements, then rather than clarifying Sellars's principle, you are proposing a significantly different principle that would have to be examined on its own grounds.

If, as you say, Sellars's principle tells us that the young child knows that this is green, then according to that principle, she has the concept of conditions which are standard for visual perception. But there are serious problems with attributing such concepts to young children, given what we know of their level of understanding. For example, recent work indicates that it is not until between the ages of 4 and 6 that children grasp what developmental psychologist Josef Perner calls the 'aspectuality of knowledge' – the idea that, for example, looking informs only about certain aspects of knowledge and touching only about others. In one telling experiment, if a coloured football is placed inside a tunnel and children have the choice of lifting the tunnel to see the ball or feeling inside with their hands to touch the object (they can only choose one or the other), and they are asked how to tell what colour the ball is, they don't understand that they have to look at it in order to tell its colour.[11] Also, they can't determine that they need to feel and not look at a piggy bank to tell whether it's full of pennies. It is not until after the fourth birthday that children typically acquire this aspectual understanding.

Prior to that developmental stage, it looks deeply implausible to insist that children even have the idea of standard conditions for visual perception, let alone know what those conditions are. *Surely*, in order to understand that there are such conditions and to have the most rudimentary sense of what these are, a person must understand that in order to tell what colour an object is, she must look at it as opposed to feel it. So children who do not grasp aspectuality could not be knowers under Sellars's requirements.

WdeV: You've been referring to typical 4-year-olds. But according to your own source, Perner, it's between ages 4 and 5 that most kids begin to recognize the aspectuality of knowledge. Perner's position is that 'from about 4 years on children develop an *information theory of knowledge*'.[12] And the information theory Perner says children acquire at this age seems to satisfy just the criteria Sellars sets for observational knowledge. He argues that the 4-year-old achieves a

> new understanding [that] arrives with an information-processing view in which knowledge is seen as a mental representation whose function it is to reflect the facts correctly. … [A]t the age of about 4 years the child acquires the necessary concept of representation for such a view. At this age representation can be understood as a medium representing something (referent) as being a certain way (sense). And this conceptual framework provides the basis for understanding that such a medium can be a true representation of the external referent situation only if there exists some reliable, causal mechanism that ensures that the medium is set into correspondence with that external situation.[13]

Perner's findings seem to *support* rather than challenge Sellars.

TT: They hardly support Sellars's position. While Perner does say that an information theory of knowledge is acquired at about age 4, the data Perner cites indicate that most children do not grasp the aspectuality of knowledge until they are between the ages of 5 and 6, in which case 4-year-olds would not qualify as knowers given Sellars's criteria.[14] However, the actual age – whether 4, 5, or 6 – at which children ascend to the status of knowers on your reading of Sellars is itself fairly unimportant, as long as it is clearly *later* than the age at which we are ordinarily willing to attribute knowledge to children. Then we still have a situation that is highly troubling to Sellars's case: we unhesitatingly attribute knowledge that this is heavy or that is green to kids who are younger than 4. But if these kids don't grasp the aspectuality of knowledge, it is highly implausible to claim that they satisfy the conditions of SPOK. Sellars therefore can't attribute to them the knowledge that this is green. But we are perfectly willing to attribute observational knowledge to kids under 4 before they have the epistemological knowledge Sellars requires.

WdeV: For very young children, we are simply wrong on that score, or perhaps we should say we are overly generous and full of hope. But you're right that pinpointing the age is of no philosophical importance if Sellars's criteria diverge significantly from the ordinary in at least some clear cases. I'm willing to stipulate that 4 is a generally important watershed; if it's closer to 6, no big deal, if you can convincingly claim that knowledge is already present at 2 or 3.

TT: The data certainly suggest that knowledge is present at these ages. Here are some examples from the child development literature about 2- and 3-year-olds in different contexts: being asked carefully constructed questions in experimental situations, and making spontaneous remarks in conversation with a parent or another child.

Example 1[15]

AG, *24 months*:	(context: AG had sorted the cylinders and columns by form, leaving color mixed within each group. She had called both the cylinders and columns 'bricks' during testing, and she had named both 'red' and 'blue' appropriately).
Experimenter:	Can you show me a red brick?
A:	Here's one. [Holds red cylinder to Experimenter.] Here's another one. [Holds second red cylinder to Experimenter.]
Experimenter:	Are all the bricks red? [A: Holds *blue cylinder* to Experimenter.] What's that?
A:	Blue brick. Green brick. [Holds the *green half-column* remaining from the oddity-probe to Experimenter.]

Example 2[16]

Jack (32 months):	(*Points to a [painted] flower on the pot.*) This is like Anne's flower. Like at her house.
Mother:	Where does she have flowers like that at her house?
Jack:	They're in the drawer. In a drawer. I like … underwear … drawers.
Mother:	You like *what*?
Jack:	Underwear drawers.
Mother:	Oh, you mean these flowers are like on her underwear.
Jack:	Those are on her underwear.
Mother:	Right. On her 'big girl' pants.
Jack:	On her big girl pants.
Mother:	Now I know what you're talking about.
Jack:	And another one, Mommy. (*Points to another flower on the pot.*)

Example 3[17]

Conversation between Jack (33 months) and Anne (35 months):
Anne (*on toy telephone*): Hello … Bye.

Jack: Who was that?
Anne: That was daddy.
Jack: What?
Anne: That was daddy.
Jack: Daddy?
Anne: Let's go for a walk with daddy. Let's pretend.
Jack: Okay, let's pretend.
Anne: Let's pretend.
Jack: I will bring this … with me. (Picks up broom.) And this. (Picks up pocketbook.)
Anne: That's mine. Mine. (Takes the pocketbook from him.)
Jack: And I'll carry this. (Picks up Raggedy Ann doll.) Will you / I will carry this dolly.
Anne: I'm gonna take a broom. (Picks up broom Jack has dropped.)

Example 4[18]

Beth (3 yrs 10 months)
Child: Is our roof a sloping roof?
Mother: Mmm. We've got two sloping roofs, and they sort of meet in the middle.
Child: Why have we?
Mother: Oh, it's just the way our house is built. Most people have sloping roofs, so that the rain can run off them. Otherwise, if you have a flat roof, the rain would sit in the middle of the roof and make a big puddle, and then it would start coming through.
Child: Our school has a flat roof, you know.
Mother: Yes it does actually, doesn't it?
Child: And the rain sits there and goes through?
Mother: Well, it doesn't go through. It's probably built with drains so that the water runs away. You have big blocks of flats with rather flat sort of roofs. But houses that were built at the time this house was built usually had sloping roofs.
Child: Does Lara have a sloping roof?
Mother: Mmm. Lara's house is very like ours. In countries where they have a lot of snow, they have even more sloping roofs. So that when they've got a lot of snow, the snow can just fall off.
Child: Whereas, if you have a flat roof, what would it do? Would it just have a drain?
Mother: No, then it would sit on the roof, and when it melted it would make a big puddle.

I think that these examples, together with the study by Perner noted above and studies by others about young children's competence with epistemic concepts, create serious problems for Sellars. For these studies suggest that coming to know *that* one knows, to know that perceptual conditions need to be appropriate, to know that such-and-such conditions are standard – these concepts come relatively late.[19] But while the level ascent to the epistemic is made relatively late, the conceptual competence of children at the observation level – the level of brooms and dolls, shapes and colours, drains and sloping roofs – is remarkably strong and sophisticated very early.

It looks reasonable to attribute propositional knowledge even to children just turned 2. In Example 1, the author reporting this study remarks about the 24-month-olds AG and DD (the latter child was not included in my excerpt from the text):

> These subjects were raising countercases to the experimenters' propositions. When AG was asked if all the bricks were red, she held up a blue block. When DD was asked if all the bricks were blue, she pointed to a red block.[20]

All these examples present situations that are presumably unique for the children in them, and they respond to them, and sometimes reason about them, in flexible, responsive, appropriate ways. Note that one can attribute to children the ability to reason, and knowledge of cause and effect concerning a wide variety of ordinary features of the environment (consider Beth's causal knowledge in Example 4, for example, her understanding the effects of rain and standing water), without thereby attributing competence with epistemic concepts to them. Nothing in a relatively simple understanding of cause and effect requires the understanding of aspectuality, false beliefs, possession of the representational theory of mind, and so on, required for competence with epistemic concepts.

About children who have such observational competence but lack epistemic competence, you and Sellars have to say, in line with SPOK, that they have no knowledge, that our attributions of knowledge are, in your words, 'overly generous and full of hope' – but wrong.

I think that that this judgment will strike most people as implausible.

WdeV: But, remember, I have claimed that Sellars does not and need not require *explicit* metaknowledge. What I have claimed he needs is that the knower at least be responsive to the epistemic norms that constitute the knowledge game. By your own admission here, these children – despite their lack of *explicit* metaknowledge – do exhibit significant responsiveness to epistemic norms. DD may not be able to say, 'I refute you thus!' She may not even command the word 'refute'. But she *knows how* to refute a mistaken claim. When you attribute the ability to *reason explicitly* to these

children, you may be giving Sellars all that he requires, depending (of course) on what you take explicit reasoning to amount to.[21]

And there's another general problem here. When we start accumulating data via very sophisticated empirical techniques, I don't think that we can rely on received knowledge to give us the best interpretation of the data. Otherwise, sophisticated science wouldn't be worth doing. These youngsters show that they have done very well in acquiring the differential response capacities that are the necessary condition of observational knowledge. But it just doesn't follow that these relatively sophisticated response capacities are therefore exercises of *concepts* or *knowledge*.

TT: I grant it doesn't *follow*. We're not doing deductive logic here. I am trying instead to present what I think is strong inductive evidence.

WdeV: Fine, but my point about not relying on received knowledge when we look at this evidence stands. For instance, you blithely assume that one can have a 'simple understanding of cause and effect' without understanding what can be ascertained by touching and what can be ascertained by looking. But that just isn't clear to me at all. Not understanding aspectuality throws into doubt their ability to mobilize evidence in any systematic way. The causal relation, as we philosophers know all too well, is a very complex one. Are you saying that these children have an understanding of such causal relations *as* causal relations? I think that you are being *very* generous and very uncritical in your rush to knowledge attribution in these cases.

TT: You are over-reaching when you claim that the child's failure to understand aspectuality makes doubtful that she is capable of any systematic use of evidence. Aspectuality is the relatively sophisticated and specific idea, concerning conditions for adequate perception, that in order to tell an object's colour, you need to look at it rather than feel it. Well before children understand aspectuality they are able to marshal and compare evidence from a variety of observational situations. Similarly, there's no a priori reason to think that children who don't grasp a perceptual principle like aspectuality are incapable of understanding simple causal relations. Young children are not as confused about the nature of causality as philosophers are! 'The rain will make puddles on the flat roof, unless it has a drain.' If young children can understand such a sentence, they have the sort of causal knowledge I am concerned about. I am not talking about a sophisticated philosophical understanding of causality, so I am not being overgenerous in the way you seem to suggest. The young child's conception may not be very sophisticated, but it already extends to a wide variety of everyday events.

WdeV: There's a difference between expecting things to be wet when they're in the rain and understanding a causal relation between rain and getting wet. Once again, you assume a generous reading of the situation.

Your generosity makes it easy to be generous with the attribution of knowledge as well, even if it violates Skinner's principle of attributing as little mentality or intelligence as possible.

TT: I am not saying merely that expectations are set up in the child. Young children use 'This makes that happen' locutions quite readily. It seems plausible to think that they understand a sentence like 'The ice on the sidewalk made me fall down' in much the same way an adult does. It is not a scientifically detailed or philosophically reflective understanding of causality in the case of the child *or* the ordinary adult. It seems to me hard to explain the child's facility with such locutions in a way that does not grant him any knowledge of causal relations, and it seems quite implausible to suggest that the child can wield such locutions without possessing any knowledge at all.

WdeV: OK, you claim that a child's being able to use common locutions appropriately fairly often requires that we grant the child knowledge of the relevant subject matter. Doesn't this methodological principle turn around to bite you? For even 2- and 3-year-olds often exhibit appropriate use of epistemic terminology. The results about aspectuality and so forth that you've cited are elicited in special laboratory conditions and come as a surprise to most people. In most actual, common situations even young kids use terms such as 'mistake' and even 'know' quite appropriately. (For example, young Jimmy taunts Johnny, 'I know something you don't know.') I think I turn your own words against you with only one change: 'It seems to me hard to explain the child's facility with such locutions in a way that does not grant him any knowledge of epistemic relations, and it seems quite implausible to suggest that the child can wield such locutions without possessing any knowledge at all.'

You want to attribute all kinds of knowledge pretty generously *except* metaknowledge. But you can't have it both ways. If the 3-year-old has knowledge of causality because he can correctly spot a number of causal relations, even though he gets scads of them wrong as well, then the 3-year-old also has knowledge of knowledge, for he can spot correctly many cases of knowledge ('Daddy, Mommy knows where you were last night'), even though he gets many wrong as well. You cannot suddenly jack up the standards when you worry about metaknowledge. If little kids have knowledge of causality by the standard you use, then I think you have to admit that they have knowledge of knowledge, and your objection to Sellars evaporates. I'm willing to apply a high standard for both, but I think that you're applying a double standard.

TT: That's a nice a priori argument against my position, but it just doesn't accord with the empirical data we have about young children. I should note first that, if you are going to count a child's recognizing and communicating about mistakes, for example, as metaknowledge, then I would not claim,

and wouldn't need to claim, that children before age 4 lack *all* metaknowledge. There are many kinds of evidential and epistemic concepts, and it's no surprise if children acquire different ones at different developmental stages. In using 'metaknowledge' earlier I was, perhaps too narrowly, thinking of metaknowledge just in terms of the recognition of aspectuality, false beliefs, and other specific epistemic capacities that studies show are not acquired until after the fourth birthday.[22] I'm happy to accept a broader understanding of 'metaknowledge', if you prefer, but this won't help support your argument.

For what is at issue is whether the specific conditions SPOK requires for observational knowledge are met in young children. And we've seen that the evidence indicates that they are not – children younger than 4 don't grasp aspectuality, and they can't satisfy the conditions for SPOK if they can't understand that. (I trust that, in your references to 'specific laboratory conditions' and surprising results, you don't mean to dismiss or downplay the aspectuality results. Science of course fundamentally depends on controlled experiments, and unexpected results are often the most significant and fruitful for our theoretical understanding.)

By contrast with data indicating a late age for acquisition of aspectuality, the data are very strong that children under four *do* understand causal relations:

> A fairly recent, but already classic, study by Lois Hood and Lois Bloom shows that at about 24 months of age, children's conversations reflect quite a sophisticated awareness of cause and effect. Two- and 3-year-olds make causal statements of various kinds ('I'm putting medicine on the lamb's leg cause he had a booboo.' 'Don't ring the bell either … Jenny will wake up.' 'You can bring me my puzzle … very careful so pieces don't fall down'). Children also ask causal questions ('Why you wrapping it around?'); and they answer causal questions (Adult: 'Why are you taking off your socks?' Child: 'Because it's not cold outside').[23]

So I'm not applying a double standard. The data clearly support different developmental stages for understanding causality on the one hand and aspectuality on the other.

WdeV: I think that your reading of the literature is a bit slanted. Piaget postulated quite a late acquisition of the full concept of causation in children, somewhere around 7, I believe it was, and since then, you are right that the tendency has been to demonstrate ever earlier acquisition of at least parts of the concept. But it is far from unequivocal that children have command of the concept of causation well before they possess the concept of knowledge that we've been worrying about. Consider the following research summary:

Das Gupta and Bryant have produced some evidence which suggests that an important development in children's understanding of cause takes place sometime during the fourth year. They reasoned that to make *genuine* causal inferences, a child must make a comparison between an initial and a final state. … A technique is to show children pictures which represent an initial and a final state (e.g., an intact window and a broken window) and then a set of possible causal objects (e.g., a plum, a stone, and a brush). The children would then be invited to select the likely causal agent. The problem Das Gupta and Bryant identify is that since children know a lot about the words and are given such a choice, they may choose a stone because of what they know about stones and windows rather than paying attention to the initial as well as the final state. The method chosen to explore this possibility and the basis of children's causal reasoning was to present the children with causal sequences which did not obey the normal rules, i.e. those which could not be easily solved by knowledge rather than inference and those in which a choice had to be made. [The experiment is described as follows: 'One example described by Das Gupta and Bryant was to show children, in order, a broken cup and then a wet broken cup. Alternatively they might begin with a wet cup and conclude with a wet and broken cup. The possible causal agents shown to the children included a hammer and water. If the children selected water for the first sequence and the hammer for the second, then one could conclude that they had paid attention to both the initial and final states and made an inference. If the children chose the hammer on both occasions, then the basis for their choice was not likely to have been an inference but rather their previous knowledge of hammers.'] Das Gupta and Bryant found that the 3-year-olds were likely to choose the hammer in both circumstances while the 4-year-olds were unlikely to make this error. The authors conclude,

> Our two experiments demonstrate that by the age of 4 years children do adopt … the more sophisticated … strategy. They can use the difference between an object's initial and final state to work out what happened to it in the meantime. In sharp contrast, the 3-year-old children have grave difficulties in using the difference between initial and final states when making a causal inference.[24]

This shift in the ability to make real causal inferences, rather than exploit conventional causal stereotypes, is important to possession of the concept of causation and seems to occur at about the same time as the blossoming of metaknowledge we've been worrying about. While we have to keep an eye on the empirical results and see how things develop, I don't think there's

anything like clear-cut empirical evidence at this point that either you or I can point to.

TT: But it's a background assumption of Das Gupta and Bryant that 3-year-olds *have* causal knowledge: 'There is evidence that [children] know that knives cut and that water makes things wet.'[25] Their explanation of the 3-year-olds' strategy in their causal inference experiment relies on these children's 'knowledge of the typical effects of particular instruments (e.g., water makes things wet)'.[26] Obviously, the extent of knowledge, attentional focus, and conceptual sophistication increases with age, but the extent of causal knowledge that is attributed to children both in the study you cite and in the one I cite is sufficient for the point I want to make regarding young children's causal knowledge.

WdeV: I grant all this. But it seems to me that these same children have an equally unsophisticated knowledge that, for example, Daddy knows where day care is. In that case the question is where to draw the line between unsophisticated knowledge of *x* (for example, causation or knowledge) and representations that are not yet knowledge. I just don't see the asymmetry between knowledge and metaknowledge here that you really need.

TT: The asymmetry that is at issue here – that between young children's knowledge of causal relations and their *lack* of knowledge of the specific kind Sellars would require for anyone to have even observational knowledge – is to be found in the empirical studies. Perhaps you don't see it because you haven't heeded my caution that 'metaknowledge' is a quite vague and broad term. Your continued reference to metaknowledge, with the wide range of states and competencies that that term can suggest, obscures the real problem for Sellars. (Even regarding metaknowledge thus broadly understood, you may be overstating the competence that 2- and 3-year-olds actually do have. One of your examples of this metaknowledge is 'I know something you don't know', but it's not clear that children at these ages have this locution or the ability to contrast their own epistemic state with another person's different epistemic state.) On the specific question at issue, the studies indicate that children have knowledge of typical cause and effect relationships before age 4 but lack knowledge of aspectuality until the fourth birthday or later. That's the specific asymmetry I need in order to make my point.

Let me turn to an earlier point you made suggesting that young children have 'differential response capacities' that are merely necessary conditions for knowledge, not knowledge proper. This is not merely a terminological debate over the proper extension of the label 'knowledge'. If young children do not have knowledge but only 'differential response capacities', you have to explain what it is that allows the child to act and speak in the ways described.[27] It is not just imitation or parroting. Although we see some of

that in Jack (Example 2), that cannot explain his initial complex action of comparison and memory by which he notices the likeness of a pattern on the pot to a non-present pattern on Anne's clothing, and expresses what he notices. The researcher comments that here Jack 'tried to find a term that would help his mother recall something he had seen'.[28] What sort of explanation could avoid reference to Jack's knowing that this is his mother, that Anne lives in a house, that what he's pointing to is like something in Anne's house?

WdeV: Notice how much generous interpretation his mother provides in Jack's case. If an adult responded the same way to the question 'Where does she have flowers like that at her house?' we would immediately wonder what was wrong and start thinking about scheduling a competency hearing, if there were no evidence of drugs or a recent blow to the head.

TT: But I think that Jack's limitations and his mother's assistance can be explained quite satisfactorily, without denying the knowledge he does seem to have, by noting his limited vocabulary and his lack of *other* sorts of knowledge, which good parents and teachers are always eager to fill in and children are well adapted to absorb. While Jack's mother does help him along in interpreting his words, which are indeed clumsy and gappy, this is like the fact I previously noted that Jack engages in some parroting or imitative behaviour. It's hardly a surprise that we see evidence of continuing behaviour that is relatively primitive given that Jack is just 2 years 8 months old! What's striking is that, in the midst of the more primitive responses, we see behaviour already at this age that looks like propositional knowledge.

WdeV: The point, it seems to me, is that in this case we *both* have things to explain. You have to explain how Jack can have knowledge despite saying and asking things that would defeat the claim that he knows, if he were an adult. I have to explain why he doesn't have knowledge, despite saying things that would normally, in an adult, be grounds for the attribution of knowledge. You're assuming that the burden of proof here is on me, but I think that in this context, we both have a burden, and we can't turn to untutored common sense to absolve us of that burden. Our ordinary concept of knowledge is pretty well tuned to the case of normal adults, but children deviate in several dimensions from that, and that makes it difficult to apply the concept straightforwardly.

TT: Your specific attempt to create a symmetry or equality of burden doesn't work, for your assumption that attributions of knowledge would be denied to any adult who behaved like Jack is incorrect. If an adult engaged in verbal behaviour at Jack's level, you're right that we would indeed put him under supervision, just as Jack with his limited knowledge and capacities is under supervision. But you can't presuppose that we would deny the adult very simple observational knowledge akin to Jack's knowledge that

Anne lives in a house, that this pattern is like that one, and so on. We wouldn't attribute to the adult the knowledge that most adults have; but there's no reason why we would claim the adult has no knowledge at all, or that such simple knowledge claims are defeated because this particular adult doesn't also have the knowledge we expect of normal adults.

WdeV: You misunderstand my claim. I did not claim that we would deny the adult knowledge, but that the kind of behaviour in question operates as a *defeater* of attributions of knowledge. A defeater, as I understand the term, impugns the attribution, but need not do so absolutely; defeaters can be overridden by other factors. With little kids there are fewer overriders available than with adults.

TT: I think that your point about defeaters gets its purchase only when you remain vague about the types of propositions that are defeated. Those at issue are ordinary observational propositions like 'this is red' or 'this pattern is like that', and there is no reason to think that these are defeated in the first place, in your sense of 'defeater'. The impaired adult is able to identify red things or brooms and to note pattern similarities (by hypothesis, since Jack is able to do these things). Unless one simply presupposes Sellarsian requirements for perceptual knowledge – which is of course what is at issue – it's not clear why we should think that defeaters exist in the case of Jack *or* the impaired adult regarding these simple observational propositions, though of course they will exist for other more complex propositions. If there are no more defeaters in the case of Jack than in the case of the impaired adult with respect to the simple propositions at issue, then the question of overriders or their relative numbers doesn't even arise.

WdeV: You think that simple perceptual identifications are virtually impossible to defeat unless one presupposes a Sellarsian holism. But your assumption that such propositions are virtually immune to defeat seems to me equally question-begging. Do you have non-question-begging, clear way to distinguish the 'simple' observation sentences that are so difficult to defeat from other, presumably less simple ones that are subject to defeat?

TT: I believe that there is such a method. But I note that here again you are trying to shift the burden to me. In cases such as this one, I wish I could accommodate you, but any attempt to address your question with substantive detail quickly gets quite technical, and would clearly take us too far afield. So although I can cite some accounts of epistemic defeaters in the literature that, I believe, do the job you are asking for,[29] I recognize that I cannot say enough here to address your question adequately. But I can at least address the larger question your move raises about who has the burden of proof.

When 'untutored common sense' takes a side on the major question at issue here about whether young children have knowledge, there is a burden

on the other side to explain why common sense is wrong. I believe that common sense would with little or no hesitation say that it looks like these children have knowledge of simple observational propositions. Of course, to say that it 'looks like' propositional knowledge is not to say that it *is*. But I think that there are very strong considerations favouring my attribution of knowledge over your denial of it in such cases. Epistemic terms like 'know', 'remember', 'understand' are deeply entrenched in ordinary language. I am by no means an ordinary-language philosopher of Austinian stripe. But I do believe that when a usage is well established, the philosopher has to have *very* good reasons to break with that usage. The philosopher certainly has a right to do so, and insight is often gained by such a break – there I part company with the ordinary-language school. But the philosopher needs to take seriously that she is arrogating to herself a significant revision of our commonsense understanding of ourselves and our world, and she'd better have damn good reasons for that. I think that attributing knowledge that this is a broom to Anne, that that is a sloping roof to Beth, and so on, is strongly embedded in our commonsense understanding of the world and of what knowledge is. And the sorts of attributions you are recommending in opposition to these common attributions are particularly radical, because it is not just a matter of the philosopher proposing a new distinction *within* an ordinary-language category, where that distinction does not require revising current usage of the broader category (an example of which would be proposing that *intension* and *extension* are two elements that help us refine our understanding of *meaning*). Rather, Sellars is proposing category displacement (or at least partial displacement) rather than refinement: something that we thought was knowledge turns out not to be knowledge at all. And I think Sellars was aware that he was proposing such a partial displacement, as for example with his remark that meeting the requirements for SPOK poses 'a steep hurdle indeed'.[30]

Surely, there needs to be justification for such a proposal.

WdeV: There are plenty of reasons to draw the line between knowers and non-knowers where Sellars does. I think the most fundamental reason is that it is very, very difficult to specify just what it is that the very young know. When we ascribe knowledge to young children, we do so using our own vocabulary. However, because we understand so much more with each proposition than they do, we tend to feel uncomfortable in ascribing such highly determinate 'knowledge' to them. For most purposes this is not a big deal, but we run real risks of error if we just assume that the concepts or propositions *we* would be using in a certain task or situation are also available to the child whose behaviour seems similar. So one reason to distinguish the 'knowledge' of the very young from that of the more mature is the principle that when one has propositional knowledge, one knows some

determinate proposition, and the difficulty we have attributing to infants such determinate representations.[31]

TT: We are not discussing infants but 2- and 3-year-olds. I do not attribute propositional knowledge to children in the first weeks and months of life.

WdeV: I didn't mean just neonates; toddlers too. Furthermore, many of the interesting questions one can ask about knowledge, questions about the development of knowledge, about purported but deviant cases of knowledge, about the causal underpinnings of knowledge, and so on, use a background conception of 'the human adult's normal, typical cognition' as a reference point with regard to which the atypical, the abnormal, the developmental, and the merely causal can be defined. It is this adult, *normal*, typical cognition that is at the heart of both the epistemologist's and the cognitive scientist's concern.[32] The reference-point phenomenon does not need to be construed extremely broadly – indeed, if it is construed too broadly, one then loses the ability adequately to identify the developmental and the atypical. Sellars keeps the focal point of the epistemic nicely delimited.

As we've seen, Sellars adds his level ascent requirement in order to supply a deficiency in what would otherwise be a purely externalist theory of knowledge. If you reject that requirement, you have two choices: accept a purely externalist epistemology or defend some other internalist requirement. If one goes the externalist route, there will be no reason to separate the 3-year-olds from the 4-year-olds. But there won't be any reason to separate them from mice, fish, and lizards either. There is a lot wrong with pure externalism, so that doesn't seem a happy alternative. Notice that Sellars escapes externalism by requiring the cognitive accessibility of the externalist conditions of knowledge. This is really a requirement on the whole cognitive system, not on any particular cognitive state. Another route would be to require the availability of some kind of knowledge that is somehow *intrinsically* internal or internalist – that is, *necessarily* or *essentially* cognitively available to the knower. Beliefs about what appears to one, about one's own sensory state, or about 'what it's like' to have some experience might be thought to be intrinsically internalist in some way. But this raises lots of questions: is there any first-order propositional knowledge that is *intrinsically* internalist? Quite frankly, I doubt it, though we won't have time to argue it here. I think that any first-order propositional representation is such that (1) it could be acquired by 'externalist' means and (2) its having been so acquired would render it capable of supporting other knowledge. Whether it is itself then knowledge depends on the characteristics of the representational system, not on the first-order propositional representation itself. Since I don't think that there are any *internalist facts*, I think that it is exactly right to locate the internalism in the structure of the cognitive system, not in the objects cognized or the particular cognitive states

achieved. The best way to do that is via some reflexivity requirement on the system. If kids don't have that kind of reflexive awareness until their fourth year – well, until then, they are externalist systems, and not yet knowers.

Last, I do want to hold on to the idea that knowledge is the fruit of virtuous practice, and that the knower should be in some sense *responsible* for her knowledge. There really is a watershed that gets crossed in the fourth year, and it should be marked:

> From all of these experiments we can see that three-year-olds' difficulties are wide-ranging. They do not understand how someone can believe something different from what they [the children] know is actually the case (the false belief task). They do not remember that they themselves once believed something different from what they now know to be true. And they do not understand how something can look different from what it really is. In all of this work great care has been taken to make sure that the particular experimental procedures used do not hide children's understanding. And control tasks have been included to show that three-year-olds can respond to the questions. Even so, they genuinely do not understand false belief.[33]

Someone who doesn't understand the phenomenon of false belief can exercise the epistemic virtues only accidentally.

TT: So there are four reasons you offer for denying knowledge to children under 4: (1) it's difficult to determine what if anything such children know, and their representations are indeterminate compared to ours; (2) Sellars's narrow construal of the epistemic, modelled on the states of normal adult humans, has theoretical advantages in understanding epistemic phenomena; (3) Sellars's specific principle for observational knowledge is the best means of avoiding externalism; and (4) knowing is an exercise of the epistemic virtues, and thus requires a sense of epistemic rights and wrongs unavailable to children under 4. Let me address each of these points in turn.

I find your first point about the indeterminate nature of the young child's states curious. For I've already given quite determinate examples of the knowledge I claim these children have, and other examples of perfectly determinate knowledge seem to leap out from the literature I've cited: Anne knows that this is a broom, Jack knows that Anne lives in a house, he knows that the pattern before him is like the one on Anne's underwear, and Beth knows that her school building has a flat roof. How much more determinate can you get? You also express a concern about adequate evidence for knowledge attributions: you say that we should be cautious in assuming that these children's states are like ours. How far do your evidential worries extend? Should I be similarly cautious in assuming that your epistemic states are like mine? In the cases at hand, the children do not just have the

requisite behaviour for us to attribute knowledge to them, but often they exhibit completely well-articulated, grammatical (if colloquial), and appropriate language. For example, Anne says, 'I'm gonna take a broom', and proceeds to pick up a broom. (And we've already discussed those cases where their language is not up to speed.) As far as the immediate evidence in the situation goes, I have no more grounds for denying that Anne knows that this is a broom than I would for denying this of you if you said and did the same thing.

WdeV: You're trying to read someone's concepts from one piece of behaviour here, and you know that's not going to be adequate. So Anne says 'I'm gonna take a broom' and picks up a broom. Nothing can be read from that by itself. Suppose you ask her to bring the broom from the garage, where you have your push-broom stored, and she reports that there is no broom in the garage, because she recognizes only the typical corn broom as a broom. What does she do when confronted by a whisk broom? Does she recognize that an electric broom is not a kind of broom, though an electric knife is a kind of knife? I find your confidence about the determinateness of the knowledge these children supposedly have itself highly curious. I think that in many of these cases, the determinateness is in the eye of the beholder.

TT: My confidence is quite well placed when we look at the observed capabilities of 2- and 3-year-olds. You suggest that children in this age range will have difficulty recognizing push brooms and whisk brooms as brooms if they have been introduced only to the typical corn broom. While there is undoubtedly a stage at which there exists this sort of inability to subsume varied instances under one category, it appears to occur early and briefly, replaced by the ability to subsume instances prior to the ages we are concerned about.[34] Children age 2, and indeed younger, are very quick at pattern recognition – think again of Jack's recognition of pattern similarities. It is indeed these children's great flexibility, adaptability, and ability to recognize new instances that suggest that Anne won't be tied to one shape or one set of materials in broom recognition. Of course, some instances may elude her correct categorization, but if you want to insist that *all* instances must be correctly categorized if there is to be knowledge, we are all in trouble. I was not myself aware, until you noted it, that an electric broom was not a kind of broom. I trust you aren't about to deny *me* knowledge that there is a broom in my closet!

You have relatedly mentioned that an adult such as yourself has so much more knowledge than a 2-going-on-3-year-old like Anne does. But I don't think that this expanded knowledge set provides any grounds for claiming that young children have *no* knowledge. Anne doesn't *fail* to know just because you have knowledge of more scope and depth. Suppose I know the basics about red wine, and you're a connoisseur. You know much more than I about its defining conditions, provenance, grape varieties, choice vintages,

and so on. There might be drinks I would mistakenly think to be red wines that would not confuse you at all. Granting all this, it hardly follows from your greater general knowledge that I can't know, when I'm drinking some at a party, that this is red wine. And it hardly follows from our greater general knowledge that Anne can't know, when she and Jack are playing with the broom, that this is a broom.

WdeV: I don't think I'm inferring that they don't know because they don't know as much as I do, but I do think that one can fail to know because one doesn't know enough. Anne can't know only in a specific instance that something is a broom: the term or concept BROOM must be general in its application for her. If she knows it's a broom, the instances to which she applies the term or concept, both actually and counterfactually, have to mirror those to which the term applies in the language generally. If she gets only a certain subclass of brooms right, say brooms in the traditional corn-broom shape, then there's a real question about just what concept she has: her response propensities in this regard are so organized that CORN BROOM-SHAPED SWEEPING TOOL operates as a genus for her, whereas in the cultural framework it is highly specific and subordinated to a broader genus. Given that the child's usage differs from that in the larger linguistic community, how do we find a 'best match' that determines just what concept she has? I think that in general we give kids the benefit of the doubt, even when very young. These kids often respond correctly in the actual cases, but I worry about their responses to counterfactual cases.

Remember the general vision here: we acquire a number of differential response propensities piecemeal, then as we develop, we not only acquire still more response propensities, but we begin knitting them together, forging connections among them. Responding to a broom with the vocable 'broom' is a manifestation or exercise of the *concept* BROOM only when one's response propensities mirror to a sufficient degree (where sufficiency must remain vague and open-textured) the socially sanctioned uses of the term. Attributing to someone the concept BROOM is, in part, an *evaluation* in which it is claimed that the person's behavioural response propensities reflect the socially sanctioned organization of such propensities.

TT: It's interesting to see the principles you make explicit in arguing for your denial of knowledge to young children. I don't think it would be easy to support your principle that, in order to know a proposition containing a term or concept, 'the instances to which [one] applies the term or concept, both actually and counterfactually, have to mirror those to which the term applies in the language generally'. For one thing, there will not always even *be* one set of instances which language users in general will agree constitutes the proper application of a term, even if one restricts oneself to normal adult language users. Suppose that the populations of the eastern and western regions of the US differ systematically in what they classify as hills and

mountains – Westerners denoting as hills what Easterners would classify as mountains. It doesn't follow from this lack of consensus that either or both groups fail to know that Mt McKinley, which all agree is solidly within the mountain category, is a mountain. But even when there is a social consensus regarding concept application, your principle doesn't seem right. Suppose that just a few of us had the more generous understanding of what counts as a mountain while the rest of the population would categorize the smaller of these as hills. It still doesn't impugn the subgroup's knowledge that Mt McKinley is a mountain. The existence of confusions or disagreements at the borderline areas of concept application does not undermine knowledge of paradigm and non-borderline cases. Otherwise, no one would have any empirical knowledge.

WdeV: Sure, it is clear that our terms are usually vague and that there is at best an overlapping consensus (to borrow a term from Rawls) about what instances count as proper applications of a term. But that's not the point. My point is rather that in attributing a concept (or a propositional attitude that employs the concept) to someone on the basis of some behavioural evidence such as a verbal response, there are two distinct levels of evaluation involved. The first – which we often take for granted – is that the behaviour in question is properly *conceptual* in nature. For instance, if you think that the kid is merely parroting or mimicking someone else in using a word, you won't even bother to ask 'what he means by' the word. Assuming that this prior question is answered positively, we can then raise the question about *which concept* the behaviour evinces. I don't think you pay enough attention to the first of these evaluations. You tend to assume that the verbal behaviour of the young exhibits at most only a difference of degree in sophistication from the behaviour of adults. I think that a more sceptical attitude is called for, and that there is a period during which children have individual responses that can, in isolation, appear very similar to the responses of an adult, but are not yet sufficiently connected with other response propensities (to counterfactual cases) that they constitute expressions of a concept.

TT: You are now focusing on an evidential question. Of course there is a first stage of evaluation at which we determine whether the child's behaviour is complex enough to count as conceptual. We don't disagree about that. What's at issue between us on this score is whether there is enough behavioural evidence reasonably to attribute concepts to children before age 4.

I agree with you that before we can begin to attribute concepts to a child, she must be familiar with and respond flexibly to actual and counterfactual cases, but I think the evidence indicates that young children *do* have the requisite familiarity and flexibility with respect to many concepts about everyday objects, events, and actions.[35] Indeed, you now acknowledge that 'these kids often respond correctly in the actual cases', though this seems in tension with your previous suggestion that Anne won't recognize various

kinds of brooms. Regarding your 'worry about their responses to counterfactual cases', this takes us back to our earlier debate about the child's understanding of causal relations, for if young children understand causal relations, to that extent at least they are able to understand and respond to counterfactuals.

So far, we have just been considering your first reason for denying knowledge to young children. Since you said that it was the most fundamental reason, it has been worth devoting considerable discussion to it. I can more briefly address your three additional reasons.

Your second reason is that scientists and philosophers, in examining knowledge, need to set as a reference point for knowledge the standard of 'the human adult's normal, typical cognition'. There is a valuable methodological procedure behind this idea, but I don't think that it goes any way toward supporting the conclusion that young children do not have knowledge. The valid methodological procedure is to start with the typical, normal, uncontroversial cases, then see what else counts as genuine knowledge. In so proceeding, scientists will not be claiming that young children do not or cannot have knowledge. Rather: Here is clear-cut knowledge; now, do young children also have knowledge? And the answer increasingly seems to be: Yes, even earlier than we originally thought. Developmental psychologists have been surprised at the extent of cognitivity in young children. This is the so-called 'cognitive revolution' in psychology that has taken place in recent decades.[36] And although this revolution attributes some knowledge to children ages 2–3 and even earlier, it also has discovered stages in the development of cognitive capabilities, whereby the epistemic knowledge required by Sellars – about false beliefs, aspectuality, and so on – does not emerge until later. We can discuss the legitimacy of this claimed revolution on some other occasion, but surely it would be dogmatic apriorism to conclude that young children are not knowers on the grounds that normal adult humans set a valuable reference point as knowers.

WdeV: I think that this is far too simple an understanding of what's happening in cognitive science and developmental psychology. Many researchers are refining our understanding of what adult, normal, typical cognition consists in. The question developmental psychologists ask is not 'Do children have knowledge?' but 'In what respects are the cognitive capacities of children similar to or different from those of the normal and typical adult?' I do not think that it is necessarily the job of the psychologist as such to make the evaluative judgment that such and such a manifestation of such and such a capacity is knowledge; that is rather for the general community of knowers to decide on the basis of the data discovered by the psychologist (and of course, the psychologists belong to that community, and so may make such a judgment in that role). We're finding all kinds of different ways in which children are both similar to and different from adults, and surely

some of the most striking are the ways in which they are similar despite their youth. I'm not being at all dogmatically aprioristic in my claims: I just want to make sure that, in making the overall evaluation of the cognitive status of children, you appreciate fully their differences from adults and don't simply get wowed by the similarities. I think that a balanced judgment ends up with a different conclusion from yours.

TT: I didn't imply – what *would* be an oversimple conception of the discipline – that answering the knowledge question is the *only* goal of developmental psychology. But, contrary to your claim, it *is* a question which developmental psychologists ask and try to answer, and not just as members of the general community but as part of their line of work.[37]

A balanced judgment must not get wowed by the similarities between adults and children, but also it must not discount the similarities because of theoretical commitments that require the denial of knowledge in young children. Your claim about adult humans as paradigms of knowledge in methodological studies only counts as evidence against young children as knowers if that is what these studies end up suggesting.

On the road to your third point – that SPOK is the best means of avoiding externalism – I think that you provide important insights in filling out what is very likely Sellars's underlying motivation for developing SPOK and including within it the specific conditions he does. But in terms of the issue at hand – whether we should attribute knowledge to young children – I don't think that you've made a case for a negative answer here. You have merely argued by assertion that you find other internalist options besides SPOK to be doubtful or problematic.

WdeV: Nor could I do otherwise in this context, unless we want to spend a great deal of time looking at all the alternatives.

TT: I grant that, but there have been, since 'Empiricism and the Philosophy of Mind', several internalist options proposed which I find significantly less doubtful and problematic than SPOK itself. I see SPOK as brilliant for its time, but, in comparison to what has been proposed since then, it comes across as a rather coarse-grained – and much too strong – attempt to avoid externalism. Subsequently, epistemologists have discovered more nuanced and, to my mind, more plausible ways to do this. An argument that we should abandon our attribution of simple observational knowledge to young children because SPOK is the only reasonable means of avoiding externalism needs to come to terms, these days, with at least a couple of the most developed internalist alternatives to SPOK.[38]

Regarding your fourth point, I'm also unimpressed with your assertion that a core requirement for knowledge is the informed ability to exercise epistemic virtues that are beyond the ken of young children. On what basis could such an assertion be so powerful, evident, and inviolable that we must

forgo attributing simple observational knowledge to children like Jack, Anne, and Beth in the face of it?

For one thing, the claim that all knowledge is an exercise of the epistemic virtues is controversial and problematic.[39] And even if epistemic virtue theory were correct, why should we think that the epistemic virtues consist only of those that you apparently have in mind? There may be some epistemic virtues short of those required by SPOK that a young child can exercise, such as paying attention. If you both (1) build into the exercise of the epistemic virtues sufficient cognitive complexity, and (2) require that every knowledge state be an exercise of such virtues, then of course something like SPOK is required for perceptual knowledge. But both claims would need to be further explicated and defended before they could offer a compelling reason for accepting SPOK. And your fourth point in effect simply asserts that these SPOK-like conditions must obtain. Therefore it begs the question at issue.

Taken individually or collectively, I don't believe that your four points offer good enough reason to abandon the attribution of observational knowledge to young children.

WdeV: I don't think of paying attention as itself an epistemic virtue. It is a capacity the exercise of which may be epistemically virtuous in some conditions. And I've never claimed to offer any list, much less an exhaustive list, of epistemic virtues. My guess is that we probably disagree at important points about the nature and role of the virtues, and that we could argue on this theme for hours. So let me throw out a quick and dirty analogy, and we'll have to conclude for now. We can distinguish prudential virtues from moral virtues, but they are not going to be mutually exclusive: a morally virtuous person will also have to be able to exercise the prudential virtues. But I also think that in the hands of a moral agent the prudential virtues are transformed. Whatever proto-cognitive virtues 3-year-olds may have get transformed in the increasingly reflective context of the older child and adult, just as their personal virtues are transformed by the dawning of moral responsibility in their hearts.

University of New Hampshire, USA

Notes

1 William Alston, 'What's Wrong with Immediate Knowledge?', in *Epistemic Justification* (Ithaca: Cornell University Press, 1989), p. 64.

2 Citations from Sellars's 'Empiricism and the Philosophy of Mind' have dual page references. '*KMG*' refers to the reprint of Sellars's essay that is included in our commentary on it: Willem deVries and Timm Triplett, *Knowledge, Mind, and the Given: Reading Wilfrid Sellars's 'Empiricism and the Philosophy of Mind'* (Indianapolis: Hackett, 2000), pp. 205–76. '*SPR*' refers to Wilfrid Sellars, *Science, Perception and Reality* (London: Routledge & Kegan Paul, 1963), pp. 127–96.

3 We discussed Alston's objection in Chapter 8 of *Knowledge, Mind, and the Given*, but much more can be said – one of the motivations of this dialogue.

4 In his study guide to 'Empiricism and the Philosophy of Mind', Robert Brandom is vague on this, but he seems to think that Sellars's level ascent requirement goes too far, and that one could just as well back off and remain with an externalist epistemology. See Robert Brandom, 'Study Guide', in Wilfrid Sellars, *Empiricism and the Philosophy of Mind* (Cambridge, MA: Harvard University Press, 1997), p. 159. Whatever the value of the point philosophically, Sellars clearly rejected pure externalism.

5 There is another way into epistemic holism for Sellars, namely via his coherence theory of concepts. But then one must both buy Sellars's theory of concepts and overcome Roderick Firth's arguments that the coherence theory of justification is independent of a coherence theory of concepts. See Roderick Firth, 'Coherence, Certainty, and Epistemic Priority', *Journal of Philosophy*, 61 (1964), pp. 545–7. No matter what one's theory of concepts, the level ascent requirement in SPOK forces a form of epistemic holism.

6 Alston, 'What's Wrong with Immediate Knowledge?', p. 70.

7 'Good Heavens! For more than forty years I have been speaking prose without knowing it.' *Le Bourgeois Gentilhomme*, Act II, Scene 4.

8 Psychological nominalism is the doctrine that '*all* awareness of *sorts*, *resemblances*, *facts*, etc., in short, all awareness of abstract entities – indeed, all awareness even of particulars – is a linguistic affair' (Sellars, 'Empiricism and the Philosophy of Mind', *KMG*, p. 240, *SPR*, p. 160).

9 Sellars thinks of the knowledge game as a norm-governed activity that, like all norm-governed activities in his view, depends for its existence on social, usually language-mediated interaction. Such norm-governed social forms of activities are essentially holistic in their mode of being, Sellars believes, because their elements are defined by their place in the whole pattern of activity. Individual pieces of behaviour fit into or exemplify these forms of activity only as parts of *patterns* among actual and counterfactual behaviours having certain causal roots. Thus, one could agree with Sellars that it is illuminating to liken knowledge to a 'game', but disagree, for instance, about the status of norms. If it is a 'game' we are *by nature* equipped and compelled to play, then our earliest, infantile explorations might as well be treated as moves in the game. (This might lead one to some form of naturalized epistemology.) Again, one could use the 'game' metaphor, but deny the holistic treatment of games, and thus hold on to the empiricist picture of knowledge acquisition by gradual accretion from simple beginnings.

10 An important additional question concerns WdeV's appeal to implicit abilities in order to account for the cognitive abilities Sellars requires of knowers. While WdeV here takes the position that Sellars can countenance a subject's having only implicit knowledge of, for example, standard perceptual conditions, without being able to make explicit what these conditions are, he acknowledges that more needs to be said in defence of this position. TT disagrees with WdeV's position, but this is unfortunately a debate that cannot be pursued here.

11 Josef Perner, *Understanding the Representational Mind* (Cambridge, MA: Bradford/MIT Press, 1991), pp. 160–2.

12 Ibid., p. 158.

13 Ibid., p. 155.

14 Ibid., pp. 161–2. See especially figure 7.2, p. 161. Since for Perner grasping the aspectuality of knowledge is a necessary condition for having an information theory of knowledge, Perner's claim that the latter theory is acquired 'at the age of about 4' is in tension with the data he cites indicating a later age for the

understanding of aspectuality. Unfortunately for our ability to resolve this discrepancy, the studies that produced the data Perner cites are unpublished.

15 Susan Sugarman, *Children's Early Thought* (Cambridge: Cambridge University Press, 1983), pp. 205–6.

16 Catherine Garvey, *Children's Talk* (Cambridge, MA: Harvard University Press, 1984), pp. 68–9.

17 Ibid., p. 154.

18 Barbara Tizard and Martin Hughes, *Young Children Learning* (Cambridge, MA: Harvard University Press, 1984), p. 124.

19 These studies include, in addition to the Perner study already cited, J. H. Flavell, E. R. Flavell, and F. L. Green, 'Development of the Appearance–Reality Distinction', *Cognitive Psychology*, 15 (1987), pp. 95–120; J. H. Flavell, F. L. Green, K. E. Wahl, and E. R. Flavell, 'The Effects of Question Clarification and Memory Aids on Young Children's Performance on Appearance–Reality Tasks', *Cognitive Development*, 2 (1987), pp. 127–44; Janet Wilde Astington, *The Child's Discovery of the Mind* (Cambridge, MA: Harvard University Press, 1993), see especially p. 118. All these studies indicate that there are limitations in children's *meta*knowledge, that is, their knowledge about knowledge, about evidential requirements for beliefs, about conditions under which one should attribute a false belief to another or to oneself, and so on. Prior to around age 4, children appear to lack such metaknowledge.

20 Sugarman, *Children's Early Thought*, p. 206.

21 Here's where Brandom's work does hook into Sellars adroitly. See Robert Brandom, *Making it Explicit* (Cambridge, MA: Harvard University Press, 1994).

22 See n. 19 above.

23 Michael Schulman, *The Passionate Mind* (New York: The Free Press, 1991), pp. 108–9. The study referred to in the quoted passage is L. Hood and L. Bloom, 'What, When, and How About Why: A Longitudinal Study of Early Expressions of Causality', *Monographs of the Society for Research in Child Development*, 44 (1979), pp. 1–47. See Schulman's footnote containing this citation for additional studies supporting an early age for the acquisition of causal concepts.

24 P. Das Gupta and P. E. Bryant, 'Young Children's Causal Inferences', *Child Development*, 60 (1989), pp. 1138–46, quoted passage from p. 1145.

25 Ibid., p. 1139.

26 Ibid., p. 1145.

27 I understand that Sellars does not see knowledge as different in kind from differential response capacities. For Sellars, a knowledge state *is* a differential response capacity of sufficient complexity to be accorded its normative evaluation as a knowledge state. (TT)

28 Garvey, *Children's Talk*, p. 69.

29 See Robert Audi, *The Structure of Justification* (Cambridge: Cambridge University Press, 1993), pp. 199–205; Roderick Chisholm, *Theory of Knowledge*, 3rd edn (Englewood Cliffs, NJ: Prentice-Hall, 1989), p. 55; Paul Moser (using the locution 'contravener'), *Knowledge and Evidence* (Cambridge: Cambridge University Press, 1989), pp. 101, 105.

30 Sellars, 'Empiricism and the Philosophy of Mind', *KMG*, p. 247, *SPR*, p. 168.

31 See Roderick Chisholm and Wilfrid Sellars, 'Intentionality and the Mental', in *Minnesota Studies in the Philosophy of Science*, Vol. 2, ed. H. Feigl, M. Scriven, and G. Maxwell (Minneapolis: University of Minnesota Press, 1957), pp. 507–39. When Chisholm objects that 'infants, mutes, and animals' can't have beliefs according to Sellars (p. 524), one of Sellars's concerns in his response is the determinateness of the representations of such beings (see pp. 527–8).

32 See Barbara von Eckardt, *What is Cognitive Science?* (Cambridge, MA: Bradford/MIT Press, 1993), pp. 6ff., for a discussion of the centrality of ANTCOG (Adult Normal Typical Cognition).
33 Astington, *The Child's Discovery of the Mind*, pp. 118–19.
34 'Cognitive developments in the second year of life [i.e. prior to the age range we are here concerned about] … include, at least, the ability to form categories and concepts.' Lois Bloom, *The Transition from Infancy to Language* (Cambridge: Cambridge University Press, 1993), p. 241. See also Schulman, *The Passionate Mind*, pp. 38–9.
35 See n. 34 above.
36 See Jerome Bruner, *The Culture of Education* (Cambridge, MA: Harvard University Press, 1996), p. xii.
37 See ibid., pp. 50–2; also Bloom, *The Transition from Infancy to Language*, p. 25.
38 I particularly recommend Paul Moser's comprehensive and detailed internalist foundationalism as developed in *Knowledge and Evidence* (Cambridge: Cambridge University Press, 1989). Compare Roderick Chisholm, 'A Version of Foundationalism', *Midwest Studies in Philosophy*, 5 (1980), pp. 543–64. It is also noteworthy that former critic of foundationalism Laurence BonJour has come to develop an internalist version of foundationalism. See his 'Foundationalism and the External World', *Philosophical Perspectives*, 13 (1999), pp. 229–49 (TT).
39 One can be a normativist about knowledge states without being an epistemic virtue theorist. For a good discussion of virtue theories in epistemology and objections to them, see John Greco, 'Virtues in Epistemology', in *The Oxford Handbook of Epistemology*, ed. Paul K. Moser (New York: Oxford University Press, 2002), pp. 287–315.