

Knowing How to Establish Intellectualism*

Daniele Sgaravatti

*Arché, Philosophical Research Centre for Logic, Language, Metaphysics and
Epistemology*

Department of Logic and Metaphysics

University of St Andrews

And

University of Eastern Piedmont 'Amedeo Avogadro'

ds363@st-andrews.ac.uk

Elia Zardini

*Arché, Philosophical Research Centre for Logic, Language, Metaphysics and
Epistemology*

Department of Logic and Metaphysics

University of St Andrews

ez4@st-andrews.ac.uk

May 15, 2008

1. INTRODUCTION AND OVERVIEW

Is there a unique, basic form of human cognition, to which every other form of human cognition, no matter how apparently dissimilar it may seem in thought and talk, can ultimately be *reduced*? Despite the etymological suggestion, human *cognition*—the wide range of different phenomena in which *reality discloses itself to the human mind*—need not be (and arguably is not) exhausted by *knowledge*: for example, irreducible warrant, perception and understanding—to name but a few epistemic properties that are recalcitrant to any straightforward assimilation to knowledge—may all well occur essentially in a full description of human cognition. Restricting henceforth our attention to (human) knowledge, our question becomes whether there is a unique, basic form of knowledge, to which every other form of knowledge, no matter how apparently dissimilar it may seem in thought and talk, can ultimately be reduced. If there were such a basic form of knowledge,

* Earlier versions of the material in this paper were given in 2007 at the UCLA Epistemology Workshop and at the Arché Knowledge and Language Seminar (University of St Andrews); in 2008, at the Arché Reading Party in Carbost (University of St Andrews) and at a research seminar at the University of Göttingen. We would like to thank all these audiences for very stimulating comments and discussions. Special thanks go to Herman Cappelen, Mikkel Gerken, Felix Mühlhölzer, Nikolaj Pedersen, Adolf Rami, Jonathan Schaffer, Jason Stanley and Crispin Wright. In writing this paper, the first author (the order being merely alphabetical) has benefitted from funding from the University of Eastern Piedmont 'Amedeo Avogadro' and the University of St Andrews, the second author from an AHRC Postdoctoral Research Fellowship.

a good case could be made that this would have to be *knowing that something is the case* (also known as ‘*knowledge-that*’),¹ or so many theorists have traditionally thought.

Many tough challenges await attempts at reducing all knowledge to knowledge-that. For example, *knowing things* (also known as ‘*knowledge-of*’) has long been thought to be irreducible to knowledge-that (see Russell [1912]). Elia knows Daniele, but he would still know him even if he had very different beliefs (and hence knowledge-that states) from the ones he actually has. Maybe it is necessary for him to know Daniele that he know that Daniele exists, but this is certainly not sufficient, since he also knows that Carla Bruni exists, but, much to his own chagrin, does not know Carla. These and similar considerations present a challenge to a reduction of knowledge-of to knowledge-that. More related to the topic of the special issue of this journal, *knowing how things are* (also known as ‘*knowledge-wh*’)² has surprisingly proven to be not so easily amenable to reductions in terms of knowledge-that (see Schaffer [2007] for a recent challenge to such reductionist attempts). In barn-façade country (see Goldman [1976]), Henry would seem to know which barn he’s looking at, but, as the majority of epistemologists have held, Henry would not seem to know that the object is looking at is a barn. These and similar considerations present a challenge to a reduction of knowledge-wh to knowledge-that (without committing ourselves to its final unfeasibility, we will touch later upon some more, hitherto unexplored, difficulties that this reduction must face).

Our leading question in this paper will concern the prospects for reduction to knowledge-that of yet another form of knowledge which is apparently distinct from it: *knowing how to do something* (also known as ‘*knowledge-how*’).³ Indeed, knowledge-how has been understood by many philosophers as consisting in a practical competence that outruns every possible knowledge-that, kicking in exactly when knowledge-that gives out. For example, it has been assumed by many philosophers that one can know every relevant fact concerning how one could and should ride a bicycle, and yet still not be in a position to ride a bicycle: very intuitively, one falls on the ground when put on the bicycle’s saddle not because one does not know that something is the case, but because one fails to put one’s wealth of knowledge-that into action—one still *does not know how* to ride a bicycle. This kind of consideration is prominent especially in the hermeneutic tradition, emerging from its sustained reflection on the notorious problem of the gap between *principles* and their *applications* to particular cases. One may know a certain principle but still fail to

¹ Knowledge-that is very frequently glossed as involving a relation between the relevant subject of knowledge and a *proposition*. Given the extreme plasticity of the notion of proposition in contemporary philosophy, we are very unclear as to what this gloss is supposed to amount to, and hence very unclear as to whether it is even extensionally adequate (see e.g. Lewis [1979] for a view of propositions (and attitudes) according to which not every case of knowledge-that would be a relation between a subject and a proposition). We will thus shun it in the following.

² As usual, we use ‘wh’ as a place-holder for the interrogative pro-forms ‘who’, ‘whom’, ‘whose’, ‘what’, ‘which’, ‘why’, ‘when’, ‘where’ etc. For the reasons mentioned in fn 3, we don’t wish to include ‘how’ in this list.

³ As we explain in section 2, ‘know how’-ascriptions are at least syntactically on a par with ‘know wh’-ascriptions. Still, since we don’t want the terminology to obfuscate important issues that will be relevant to our discussion, we’re going to use ‘knowledge-wh’ and its like in a way that does not include knowledge-how, even though, of course, this terminological choice leaves fully open that knowledge-how might be identical to knowledge-wh in all relevant respects.

apply it correctly in a particular case. Not even knowledge of a particular way to the effect that that way is a correct way to apply the principle would seem to guarantee that one will apply the principle correctly—one may still interpret that piece of knowledge as requiring a very different application from the one which it in fact does. Only *knowledge of how to apply the principle* would seem to be able to bridge the gap between it and its applications.

The reduction of all knowledge to knowledge-that thus arguably requires the truth of the following *reductionist* thesis concerning knowledge-how:

- (R) Knowing how to do something consists in knowing that something is the case

Going back to a usage introduced by Ryle [1949], in our debate (R) is also known as ‘*intellectualism*’ (see Stanley & Williamson [2001]), while its anti-reductionist denial is correspondingly known as ‘*anti-intellectualism*’.^{4,5} In the following, we intend critically to examine the prospects for (R), focussing—by the way of a case study—on a particularly promising version thereof which has recently been offered and developed by Stanley & Williamson [2001]. In the bulk of the paper, we aim to provide a battery of arguments

⁴ We should like to stress that anti-intellectualism is restricted to ‘know how’-ascriptions taking an *infinitive* rather than finite clause (see section 2). Clearly, Jacques’s knowing how the exam went is no more problematic for (R) than standard knowledge-wh is. We should also like to stress that, even restricted to ‘know how’-ascriptions taking an infinitive clause, anti-intellectualism need not hold an anti-reductive thesis for *every* such ascription. For example, it would seem desirable to allow that Jacques’s knowing how to prove the completeness of first-order logic does consist after all in a certain (possibly very complex) knowledge-that state: if Jacques knows how a proof of the completeness of first-order logic goes (where such a state apparently need not involve anything but a (possibly very complex) knowledge-that state, whose content is something of the form: ‘A proof of the completeness of first-order logic goes by first directly inferring P_i from $P_0, P_1, P_2 \dots P_{i-1}$, and then directly inferring P_{i+j} from $P_i, P_{i+1}, P_{i+2} \dots P_{i+j-1}$, and then directly inferring P_{i+j+k} from $P_{i+j}, P_{i+j+1}, P_{i+j+2} \dots P_{i+j+k-1} \dots$, and then directly inferring the completeness of first-order logic from $P_{i+j+k \dots +m}, P_{i+j+k \dots +m+1}, P_{i+j+k \dots +m+2} \dots P_{i+j+k \dots +m+n}$ ’), this would seem enough to establish that he knows how to prove the completeness of first-order logic. The fact that Jacques might not be able (or know how) to produce a token of the proof (either by vocalising it, or by inscribing it, or by some other means of expression), or that he might not even be able (or know how) to “run through it” in his own occurrent thoughts, would seem to be neither here nor there. While recognising that the issue invites further reflection, we emphasise that what matters for anti-intellectualism is the thesis that at least *some* cases of knowledge-how are not reducible to cases of knowledge-that—for example, paradigmatic cases involving relatively simple material actions such as knowing how to ride a bicycle, knowing how to swim, knowing how to play football etc. In the following, ‘knowledge-how’ and its like will be implicitly understood as carrying this further restriction.

⁵ We should like to stress that, on our view, it is an open question whether the same kind of practical competence that the anti-intellectualist reads paradigmatic ‘know how’-ascriptions as targeting is also targeted by some ‘know wh’-ascriptions in infinitive form (on at least one possible reading). For example, one can know every relevant fact concerning when and where one could and should stamp one’s feet in order to dance flamenco, and yet still not be in a position to dance flamenco: very intuitively, one ends up dancing something that looks rather like a waltz not because one does not know that something is the case, but because one fails to put one’s wealth of theoretical knowledge-wh into action—one still *does not know when and where* to stamp one’s feet in order to dance flamenco (nor, of course, does one know how to do it). Indeed, one could employ some of the arguments that in the following we’ll use to prise apart knowledge-how from both knowledge-that and theoretical knowledge-wh to buttress this intuition. We won’t pursue this issue further here, and will henceforth rule out such cases from our understanding of ‘knowledge-wh’ and its like.

which show that Stanley & Williamson's view is untenable, and which do so without actually relying very much on the peculiarities of their view: hence, even though, for lack of space, we won't attempt to recast our arguments in a more general fashion, we believe that they present an interesting challenge not only to Stanley & Williamson's specific view, but to a wide range of theories subscribing to (R).

The plan for the rest of the paper is as follows. Section 2 introduces the essential features of Stanley & Williamson's reductionist view. Section 3 attacks their identification of knowledge-how with knowledge of an answer, exploiting the case of negative answers. Section 4 uncovers a disanalogy between knowledge-how and knowledge-that concerning their different behaviour in relation to singular contents. Section 5 argues that, as a matter of principle and contrary to knowledge-that states, knowledge-how states are not available to a subject as a basis for acquiring inferential knowledge. Section 6 offers evidence against Stanley & Williamson's presupposition that knowledge-wh is reducible to knowledge-that, rejecting a possible attempt at conciliating the data with the presupposition. Section 7 briefly recapitulates our case and draws some conclusions about the prospects for Stanley & Williamson's reduction and, more generally, for intellectualism.

2. A REDUCTIONIST STRATEGY

In their seminal [2001], Stanley & Williamson have attacked anti-intellectualism claiming that it is in tension with contemporary linguistic theory. They note that, from a *syntactical* standpoint, 'know how'-ascriptions simply belong to the wide class of reports (so-called '*infinitive indirect-question* reports') of the form '*s S WH to A*', where '*s*' is schematic for an appropriate noun phrase, '*S*' for an appropriate attitude verb, '*WH*' for an interrogative pro-form and '*A*' for an appropriate infinitive verb phrase. Examples are:

- (1a) Rudolf remembers whom to give the book to
- (1b) The mafiosi knew where to find the money
- (1c) Some students learnt why to vote for the Communist Party.

As they recall, the accepted syntactic structure of infinitive indirect questions is:⁶

- (2a) Rudolf remembers [whom PRO to give the book to *t*]
- (2b) The mafiosi knew [where PRO to find the money *t*]
- (2c) Some students learnt [why PRO to vote for the Communist Party *t*],

⁶ We follow Stanley & Williamson in ignoring here many syntactic niceties. Square brackets indicate clausal boundaries.

where ‘PRO’ is an unpronounced pronoun and ‘*t*’ the trace of the movement of the relevant interrogative pro-form from its original position.

The syntactic structure of infinitive indirect-question reports is thus no different in the relevant respects from that of so-called ‘*finite indirect-question* reports’. The latter are of the form ‘*s S WH P*’, where ‘*P*’ is schematic for a sentence whose main verb phrase is finite. Examples are:

- (3a) Rudolf remembers whom Gustav had to give the book to
- (3b) The mafiosi knew where the policemen would find the money
- (3c) Some students learnt why Karl and Friedrich voted for the Communist Party.

Hence, the syntactic structure is ultimately not relevantly different from that of propositional-attitude reports (that is, reports taking a ‘that’-clause): in all three cases, the attitude verb is analysed as taking a *sentential* complement.

Stanley & Williamson proceed then to apply to ‘know how’-ascriptions the main features of contemporary *semantic* accounts of indirect questions.⁷ According to such accounts, an indirect-question report of the form ‘*s S WH P*’ (with ‘*P*’ in either infinitive or finite form) is true iff,⁸ for every (some) correct answer *Q* to question⁹ as to *WH P*, *s S* that *Q*.¹⁰ For example, ‘George remembers who was at the party’ is true iff, for every (some) correct answer *Q* to question as to who was at the party, George remembers that *Q*; ‘Vinny knows where to buy good Italian food’ is true iff, for every (some) correct answer *Q* to question as to where to buy good Italian food, Vinny knows that *Q* (some reports, like the former, seem most natural under the ‘every’-reading of their truth conditions; others, like the latter, under the ‘some’-reading, see Hintikka [1976]).¹¹ Stanley & Williamson’s application has it then that a ‘know how’-ascription like:

⁷ For simplicity and concreteness, we follow Stanley & Williamson in assuming Karttunen [1977]’s well-known theory of indirect questions. As they note, nothing important for our discussion hinges on the specifics of this choice. We also ignore some features of their proposal that are not relevant to our discussion.

⁸ From now on, some ‘iff’-claims will crucially occur in our discussion of intellectualism. This is however purported to be as strong as a *constitutive* thesis—claiming that knowledge-how *consists in* (is a species of) knowledge-that—and not simply an *extensional* or *modal* thesis—claiming that there is some sort of *correlation* between the instantiation of knowledge-how states and the instantiation of knowledge-that states. The relevant ‘iff’-claims will thus have to be understood accordingly as carrying the force of a constitutive claim (even though our arguments against some of these claims, if successful, will in effect undermine even the weaker correlation readings).

⁹ Throughout, we’ll be rather cavalier in our use of ‘question’ and ‘answer’, sometimes referring to the non-linguistic entities entertained when someone asks something and someone answers something, sometimes referring to the abstract linguistic entities that are usually tokened when someone asks something and someone answers something, sometimes referring to the speech acts of asking and answering.

¹⁰ Note how such a semantics would quite generally reduce knowledge-wh to knowledge-that. See our cursory comments on this in section 1 and our discussion of recalcitrant data that such a reduction has to face in section 6.

¹¹ As these examples might suggest, a lot must be packed into ‘correct’ in order to achieve an at least *prima facie* plausible account, and of course a particular semantic theory of indirect questions which follows this

(4) Diego knows how to play football

is true iff, for some correct answer Q to the question as to how to play football, Diego knows that Q (they observe that for ‘know how’-ascriptions the ‘some’-reading is very plausible). But what does Diego know when he knows any such Q ?

Stanley & Williamson observe that, for ‘know how’-ascriptions, ‘PRO’ is most naturally interpreted as referring to the subject of the attitude, and the modality expressed by the infinitive as being possibility-like. They also assume that, just as one knows a correct answer to the question as to who was at the party iff, for some person p that was at the party, one knows that p was at the party, so one knows a correct answer to the question as to how P iff, for some way w in which P , one knows that [P in w]. Under these assumptions, Stanley & Williamson’s application has it that (4) is true, iff, for some way w in which he could play football, Diego knows that he could play football in w .¹² Moreover, there is a possible situation where Diego is a complete disaster with a football, but, having attentively watched Pelé play football on the television, presumably knows, of the way Pelé played, that he could play football in that way. In order to avoid being committed to saying that, in such a situation, (4) is true, Stanley & Williamson add the epicycle that a ‘knowledge how’-ascription like (4) either semantically or pragmatically implies that, for some way w in which he could play football, Diego knows that he could play football in w by entertaining this proposition¹³ under a so-called ‘*practical mode of presentation*’—that is, a mode of presentation that bears a privileged connection to action just like the first-personal mode of presentation associated with ‘I’ has long been argued to do (Perry [1979]

very broad style of analysis can only be evaluated for adequacy once this and other important details have been settled. In particular, ‘correct’ had better mean something stronger than ‘true’, for otherwise George and Vinny could far too easily know, respectively, who was at the party and where to buy good Italian food, since they know, respectively, that whichever person who could truly be believed to have been at the party was at the party and that one can buy good Italian food wherever one can buy good Italian and Mexican food. How best to specify this stronger sense of ‘correct’ is a notoriously vexed issue (see Boër & Lycan [1986]). We’ll discuss in section 3 another respect in which the delimitation of the set of “correct” answers is very important for our topic.

¹² ‘Could’? As Stanley & Williamson are well aware, there are compelling examples which show that knowledge-how does not require the kind of ability usually denoted by ‘can’: for instance, a football player whose legs have been chopped off may still *know how* to play football, even though, in the most natural sense, he *can* no longer do so (this kind of example is due to Ginet [1975]). Similarly, in the most natural sense, it is no longer the case that he *could* do so. Now, suppose that Diego is such an unfortunate football player. Then, in the most natural sense, it is not the case that, for some way w in which he could play football, Diego knows that he could play football in w , for the simple reason that, for no w , Diego could play football in w (because unlucky Diego could no longer play football at all) and knowledge-that is factive. But (4) is arguably still true in such a situation. The lesson is that ‘could’ had better be read in a slightly technical and idealised sense if Stanley & Williamson’s theory is to stand a chance of being correct. But Stanley & Williamson themselves have objected to Ryle [1949]; [1971]’s anti-intellectualist theory by (rather uncharitably) reducing Ryle’s rich and nuanced talk of “abilities”, “skills”, “capacities” etc. to the kind of non-idealised ability that a person in Diego’s situation would lack. Once the Rylean anti-intellectualist is allowed to use the slightly technical and idealised sense of ‘ability’-talk that Stanley & Williamson themselves need in order to save their theory from easy counterexamples, their objection to Rylean anti-intellectualism evaporates.

¹³ Throughout, we’ll be assuming for definiteness a broadly *Russellian* theory of propositions. All the discussion could easily be reframed in the other main theoretical frameworks for thinking about propositions.

is one of the *loci classici* for the latter claim). Given this final analysis, knowledge-how is claimed to have been revealed to be simply a species of knowledge-that,¹⁴ and (R) is sharpened as:

(R^{SW}) s knows how to F consists in, for some w in which s could F , s 's knowing (under a practical mode of presentation) that she could F in w .

3. NEGATIVITY

In section 2 we have seen that, according to contemporary semantic accounts of indirect questions, a ‘know wh’-ascription like:

(5) George knows who was at the party

is true iff, for every (some) correct answer Q to the question as to who was at the party, George knows that Q . While, as applied to ‘know how’-ascriptions and barring the worries mentioned in fn 14, this very broad style of analysis, if correct, would presumably already suffice to establish (R), it is legitimate to ask about the *admissible values* of ‘ Q ’ in order to assess the analysis’ final plausibility. As has emerged in section 2, it’ll be useful to examine whether the *knowledge-of-answer* statements:

(5KA^E) George knows every correct answer to the question as to who was at the party

(5KA^S) George knows some correct answer to the question as to who was at the party

are equivalent with the corresponding *knowledge-of-objects* statements:

(5KO^E) For every person p that was at the party, George knows that p was at the party

¹⁴ As Crispin Wright has emphasised to us in conversation, there is a strong worry that the epicycle about “practical modes of presentation” risks to give the game away. For it implies that, in the final analysis endorsed by Stanley & Williamson (at least taking the option of semantic implication, which seems in other respects much more preferable, given the problem raised by disastrous Diego), knowledge how is *not* simply any old knowledge-that. It is knowledge-that plus a crucial additional feature in the subject’s cognitive system—knowledge that something is the case *plus a privileged connection between this state and action*. Is this something the anti-intellectualist needs to disagree about? Certainly, anti-intellectualism can allow that some form of knowledge-that is a necessary condition for (and possibly, in some sense, “part of”) knowledge-how! There is also a worry about Stanley & Williamson’s *knowledge-that* requirement once “practical modes of presentation” are introduced (see Zardini [2007]). In the following, however, we propose to ignore these pressing worries, pursuing other avenues of trouble for (R^{SW}), which relate more to semantic and epistemic (rather than mental) aspects of our topic.

(5KO^S) For some person p that was at the party, George knows that p was at the party

(we have already had occasion in section 2 to note the ambiguity between the ‘every’- and the ‘some’-reading).

We have serious doubts about the correctness of almost all these putative equivalences. By way of building up to our main argument in this section, we start with the putative entailments from (5KO^E) to (5KA^E) and from (5KO^S) to (5KA^S). We first note that such putative entailments (in particular, the one from (5KO^E) to (5KA^E)) stand a chance of being correct only given an appropriately restrictive understanding of what the range of correct answers to the question as to who was at the party is. Such an appropriately restrictive understanding is required in order to rule out counterexamples of the following form: suppose that, interestingly enough, all and only the former US Presidents (those who are still alive!) were at the party, but that George, even though knowing, of each person (former US President) p that was at the party, that p was at the party, does not know that p is a former US President. Then, under some extremely natural assumptions, George does not know that all and only the former US Presidents were at the party, even though that all and only the former US Presidents were at the party may contextually be an excellent answer to the question as to who was at the party.

However, sometimes what would have to be an appropriately restrictive understanding of what the range of correct answers to the question as to who was at the party is in order for the entailments from (5KO^E) to (5KA^E) and from (5KO^S) to (5KA^S) to be guaranteed is actually of dubious adequacy. Already with regard to the previous example, there are possible contexts where the interest in whether former US Presidents were at the party is so salient that an utterance of (5) in such contexts would be infelicitous and indeed false. To strengthen this point, consider a situation where the party is a masquerade with, unbeknownst to George but under his carefully observing eyes, Bill Clinton attending masked as Abraham Lincoln. Suppose further that the day after George is watching a documentary on former US Presidents and George is told that one of them was attending the party the night before. Suppose finally that, for no other person p that was at the party, George knows that p was at the party. It is very plausible that George’s careful observation of the person masked as Lincoln gives him knowledge of Clinton that he was at the party, even though, in the context of the day after the party, George knows no correct answer to the question as to who was at the party (and even though an utterance of (5) in the context of the day after the party is infelicitous and indeed false).

Even more straightforward and related to our upcoming argument from negativity, the putative entailment from (5KO^E) to (5KA^E) is plainly incorrect in a case where no one was at the party and George does not have the faintest idea about the party’s attendance (because, say, being completely disconnected from the relevant people and groups involved in the party, he is not even aware of the fact that that party had been planned). In such a case, (5KO^E) is vacuously true (since no person was at the party), but (5KA^E) (and (5)) is clearly false (since George has no clue about the party’s attendance). Notice that this

is so independently of one's favoured representation of the *quantificational structure* of (5KO^E), either as involving a *unary* universal quantifier along the lines of:

(5KO^{EU}) For every person p (p was at the party \supset George knows that p was at the party)

or as (more plausibly) involving a *binary* universal quantifier along the lines of:

(5KO^{EB}) For every person p (p was at the party, George knows that p was at the party)

(see the next paragraph for discussion of logically stronger, existentially committing quantifiers). Although we have no space to undertake a full survey of the options available, we should mention that vacuous truth will still result on most representations (extensional or intensional) of the *conditionality* possibly implicit in (5KO^E), and that many intensional representations under which (5KO^E) would come out false are dubiously adequate as general representations of the kind of statements to which (5KO^E) belongs (in the sense that they have the opposite flaw of also counting as false many statements of the kind to which (5KO^E) belongs that should rather come out true). Opting e.g. for a representation in terms of *strict implication* (denoting it with ' \square ') would certainly make:

(5KO^{E□}) For every person p (p was at the party \square George knows that p was at the party)

false in our last example (since at least one person p is such that it is metaphysically possible that [p was at the party and George does not know that p was at the party]), but the existence of the very same metaphysical possibility would also make (5KO^{E□}) false even if George knew full well who was at the party.

It is important to observe how the last counterexample to the putative entailment from (5KO^E) to (5KA^E) cannot be avoided by simply beefing up (5KO^E) in such a way that it carries an existential commitment to someone's having been at the party, as in:

(5KO^{ES}) Some person was at the party and, for every person p that was at the party, George knows that p was at the party.

This is so because, even if—setting aside the earlier worry illustrated by the Clinton/Lincoln example—it may now be granted that (5KO^{ES}) is strong enough as to entail (5KA^E), (5KO^{ES}) is in effect so strong as not be any longer plausibly entailed by (5KA^E). For suppose that no one actually was at the party but that this time George knows full well about this circumstance. Then George does have an excellent answer to the question as to who was at the party—namely, 'No one'! We may further suppose that George knows of all the people who possibly could have gone to the party (Janet, Mary, Moira etc.) why they did not go to the party, and so that he is also in a position to give an even more informative answer along the lines of 'Neither Janet nor Mary nor Moira etc. did, because such-and-such was the case'. Indeed, whatever the range of correct answers to

the question as to who was at the party may exactly amount to in this case, the scenario can clearly be enriched (if it need be) in such a way as to have George in possession of all such answers. In such a situation, $(5KA^E)$ is true but $(5KO^{ES})$ is false, since its first existentially committed clause is not satisfied.¹⁵

Moreover, not only is $(5KA^E)$ true in the situation just described—in it, (5) itself is arguably true. As in many cases of untypical truth, this semantic judgement may be obfuscated by some pragmatic noise. Yet, it can be made to emerge in its strength in various ways. Firstly, and most simply, all the natural *negations* of (5)—such as ‘George does not know who was at the party’, ‘It is not the case that George knows who was at the party’, ‘The following is not true: George knows who was at the party’—just sound plainly false. Secondly, the implication that someone was at the party can be explicitly *cancelled* in such a way that (5) becomes even more clearly acceptable, as in: ‘George knows well who was at the party: he has seen with his own eyes that no one was there!’, ‘I don’t know whether anyone was at the party, but ask George—he knows who was there’ and ‘No matter whether one, or ten, or billion or no people at all are at the party: George always knows who is there’. Thirdly, (5) is straightforwardly entailed by some perfectly acceptable *implicit comparison*. Suppose for instance that we are comparing George, whose epistemic state is as just described, with James, who does not have the faintest idea about the party’s attendance. Given such comparison, ‘George is the one who knows who was at the party’ is compelling. Fourthly, (5)’s truth is entailed by the most natural way of *connecting* indirect questions with direct questions and answers, which requires that (5) be true (under the ‘every’-reading) iff George knows every answer to the question as to who was at the party (with suitable restrictions on the domain of the universal quantifier). The appeal of such a direct connection is even stronger once one considers that statements like

¹⁵ As suggested by e.g. Groenendijk & Stokhof [1982], in some contexts the truth of (5) may well demand a stronger condition than $(5KO^E)$, requiring knowledge of both positive *and negative* cases:

$(5KO^{E\pm})$ For every person p that was at the party, George knows that p was at the party, and, for every person p that was not at the party, George knows that p was not at the party.

For instance, if there is an early-morning lecture the day after the party and it is important to determine how many handouts to print out, but George falsely believes that Janet, Mary and Moira were at the party (while he also knows of every person that was at the party that she was at the party), an utterance of (5) could naturally be interpreted as false, which provides good evidence that its truth condition is something along the lines of $(5KO^{E\pm})$ (of course, with suitable restrictions on the domain of its universal quantifiers). $(5KO^{E\pm})$ is not existentially committing to someone’s having been at the party, and so does not fall afoul of our criticism to $(5KO^{ES})$. Nor does it fall afoul of our vacuous case where no one was at the party and George does not have the faintest idea about the party’s attendance. Still, we submit that $(5KO^{E\pm})$ represents an implausibly strong truth condition for many uses of a sentence like (5): for instance, if all that matters is to thank people for coming to the party the day after the party and George knows of every person that was at the party that she was at the party (while he is also agnostic about a couple of people who actually did not show up), an utterance of (5) is most naturally interpreted as true. Moreover, $(5KO^{E\pm})$ ’s analogues for indirect questions fronted by other interrogative pro-forms are even less plausible truth conditions in most contexts of use. In any event, we can afford to remain neutral here about the extent to which $(5KO^{E\pm})$ and its analogues provide correct truth conditions for ‘know wh’-ascriptions, since our argument from negativity will go through even under the exacting conditions required by them. There are other truth conditions stronger than $(5KO^E)$ which share the relevant features of $(5KO^{E\pm})$ discussed in this footnote—the problem of which of these truth conditions are suited to interpret which ‘knowledge wh’-ascriptions is known in the literature as ‘the exhaustiveness problem’ (see Groenendijk & Stokhof [1997]).

‘George knows the answer to the question as to who was at the party, but ‘George knows who was at the party’ is not true’, ‘George knows that no one was at the party, but ‘George knows who was at the party’ is not true’ and ‘I know that George knows the answer to the question as to who was at the party, but I don’t know whether he knows who was at the party’ sound plainly contradictory. We encounter the same patterns with ‘what’-, ‘which’-, ‘when’- and ‘where’-indirect questions.¹⁶

Substantially the same pattern is also displayed for ‘how’-indirect questions. Suppose that George’s partner Elizabeth did not dress at all the night of the party but instead went naked straight to bed and that all the relevant details of this circumstance are known to George. Consider then the ascription:

(6) George knows how Elizabeth dressed the night of the party

and the related knowledge-of-answer statements:

(6KA^E) George knows every correct answer to the question as to how Elizabeth dressed the night of the party

(6KA^S) George knows some correct answer to the question as to how Elizabeth dressed the night of the party.

Again, George does have an excellent answer to the question as to how Elizabeth dressed the night of the party—namely, ‘In no way’! Furthermore, he is also in a position to give an even more informative answer along the lines of ‘In no way—she didn’t dress at all, but had a shower and then went straight to bed’. Indeed, whatever the range of correct answers to the question as to how Elizabeth dressed the night of the party may exactly amount to in this case, the scenario can clearly be enriched (if it need be) in such a way as to have George in possession of all such answers. In such a situation, both (6KA^E) and (6KA^S) are true.

Moreover, not only are (6KA^E) and (6KA^S) true in the situation just described—in it, (6) itself is arguably true even if it is also in many ways just as infelicitous as (5) is in the situation where George knows that no one was at the party. Again, the semantic fact can be made to emerge in various ways. Firstly, and most simply, all the natural negations of (6)—such as ‘George does not know how Elizabeth dressed the night of the party’, ‘It is

¹⁶ ‘Why’-indirect questions do not seem to align in this respect (with respect to the fourth consideration, arguably not because knowledge of the answer to a ‘why’-question does not imply the truth of the relevant ‘know why’-ascription, but because there seems to be no answer to the question as to why *P* if it is not the case that *P*). An utterance of ‘George knows why the party was a success’ seems untrue if George knows that the party was not a success and hence that there is no reason in virtue of which the party was a success. No analogues of the considerations in favour of the truth of (5) with George knowing a negative answer seem to apply to ‘why’-indirect questions (see also fn 18). We wish to remain neutral here as to how to explain this peculiarity (maybe some version of the semantic-presupposition view to be mentioned later in the text could profitably be applied to this case)—as we proceed to observe, in this respect ‘how’ seems to follow ‘who’ rather than ‘why’. (Hence, some of the following claims in the text about ‘know wh’-ascriptions are to be understood as restricted to non-‘why’ cases.) Thanks to Jason Stanley for discussion on this issue.

not the case that George knows how Elizabeth dressed the night of the party’, ‘The following is not true: George knows how Elizabeth dressed the night of the party’—just sound plainly false. Secondly, the implication that Elizabeth dressed in some way the night of the party can be explicitly cancelled in such a way that (6) becomes even more clearly acceptable, as in: ‘George knows well how Elizabeth dressed the night of the party: he has seen with his own eyes that she didn’t dress at all!’, ‘I don’t know whether Elizabeth dressed in some way the night of the party, but ask George—he knows how she did’ and ‘No matter whether Elizabeth dresses red, or blue, or black, or doesn’t dress at all the night of the party: George always knows how she dresses’. Thirdly, (6) is straightforwardly entailed by some perfectly acceptable implicit comparison. Suppose for instance that we are comparing George, whose epistemic state is as just described, with James, who does not have the faintest idea about whether and what Elizabeth dressed the night of the party. Given such comparison, ‘George is the one who knows how Elizabeth dressed the night of the party’ is compelling. Fourthly, statements like ‘George knows the answer to the question as to how Elizabeth dressed the night of the party, but ‘George knows how Elizabeth dressed the night of the party’ is not true’, ‘George knows that Elizabeth dressed in no way the night of the party, but ‘George knows how Elizabeth dressed the night of the party’ is not true’ and ‘I know that George knows the answer to the question as to how Elizabeth dressed the night of the party, but I don’t know whether he knows how Elizabeth dressed the night of the party’ sound plainly contradictory.

An alternative view on the perceived infelicity of some ‘know wh’-ascriptions in situations where the subject possesses a correct negative answer would attribute to these ascriptions, and to the corresponding interrogations, a *false semantic presupposition* (such view has been explicitly taken for direct questions by Katz [1972] and Keenan & Hull [1973] and is actually briefly endorsed by Stanley & Williamson [2001], p. 421, fn 19 at least for the case of ‘know who’-ascriptions). Such an account could be supported by appealing to cases which present a more acute infelicity than those we have just discussed, like e.g. ‘George knows how Elizabeth passed the exam’ or ‘George knows which numbers are a counterexample to Fermat’s Last Theorem’ (‘how’ is not unique in exhibiting such cases). On this view, utterances of (5) and (6) in a situation where the subject possesses a correct negative answer would be relevantly akin to a present utterance of ‘George met the present king of France’ (if you don’t like this example, plug in your favourite case of semantic presupposition). However, it seems to us that these cases are clearly different from the point of view of intuitive acceptability. Even ‘George knows how Elizabeth passed the exam’ sounds much better than ‘George met the present king of France’ (with ‘George knows when the present king of France was born’, on the relevant narrow-scope reading of the definite description, lying in between).¹⁷ More importantly, an attribution of semantic presupposition would be very hard to motivate given that the same indirect questions clearly don’t carry any such presupposition when embedded under other verbs like ‘have a view about’, ‘bet on’, ‘depend on’ etc.¹⁸ In conclusion, given the various theoretical virtues of the kind of semantic theory of indirect questions we have been

¹⁷ Thanks to Herman Cappelen for suggesting to us to consider embeddings into indirect questions of expressions carrying semantic presuppositions.

¹⁸ Again, ‘why’-indirect questions behave differently here: for example, an utterance of ‘Whether Elizabeth went to the party depends on why the party was a success’ seems untrue if the party was not a success.

working with and given the overwhelming wealth of data we've just presented in favour of the truth of 'know wh'-ascriptions with known negative answers, we think it clear that the best explanation for the infelicities in question cannot be but pragmatic rather than semantic (the same conclusion is reached, among others, by Groenendijk & Stokhof [1982]; [1997]). It is an interesting question how exactly to characterise the pragmatic mechanisms involved—here, however, we shall not speculate further about this (see e.g. Karttunen & Peters [1976] for a specific proposal).

That being said, our upcoming argument is compatible with a version, which we take to be the least implausible, of the semantic-presupposition view. On the version of this view that we have in mind, direct and indirect questions could take on different semantic presuppositions depending on the conversational context. This would happen via the (direct or indirect) question's somehow contextually selecting a range of available answers. If one logically possible answer is not selected, its negation is semantically presupposed by the question. Thus, while in some contexts 'No one' would be a correct answer to the question 'Who was at the party?', in other contexts it would rather be a way of "rejecting the question" (analogously for indirect questions). If you should be attracted by such a view, just assume that the cases we are discussing are cases in which the question does not carry any relevant semantic presupposition.

With so much by way of stage setting, our argument from negativity should begin to emerge. For, while it might be granted that the previous considerations in favour of the truth of 'know wh'-ascriptions with known negative answers do not strictly speaking impugn the entailment from $(5KA^E)$ to $(5KO^E)$ (since, in the situation where George knows that no one was at the party, $(5KO^E)$ is vacuously true), they do show that the converse entailment fails, and, more importantly, that both $(5KA^E)$ and $(5KO^E)$ can be true in a situation where no one was at the party (with, on the one hand, the truth of $(5KA^E)$ requiring a substantial cognitive feat by George while, on the other hand, the truth of $(5KO^E)$ being vacuously guaranteed by the facts of the matter). Even more importantly, they also show that the entailment from $(5KA^S)$ to $(5KO^S)$ fails: while George has some correct answer to the question as to who was at the party ('No one'), there is no person p such that George knows that p was at the party. And, in such cases, we have argued that both theory and data couple the truth of (5) with the truth of $(5KA^S)$ rather than $(5KO^S)$.

Now, as we have already remarked in section 2, some 'know wh'-ascriptions lend themselves better to the 'every'-reading while some others lend themselves better to the 'some'-reading. Moreover, as Stanley & Williamson themselves seem to acknowledge, an analysis of 'know how'-ascriptions in terms of 'know wh'-ascriptions had better use the 'some'-reading on pains of leaving precious little knowledge-how around. If in order for Diego to know how to play football he has to know, of each and every way w in which he could play football, that he could play football in w , then, even taking into account possible reasonable restrictions on the domain of the universal quantifier, a standard for knowledge-how would have been set that even *el pibe de oro* could hardly ever satisfy. It is thus quite essential to the success of Stanley & Williamson's analysis that the 'some'-reading is adopted at least as default for 'know how'-ascriptions. This, however, makes the analysis particularly prone to the risk of the negativity glitch that has just been unearthed.

Indeed, the analysis is in effect fatally affected by the negativity glitch. For concreteness, suppose that George knows full well about all the relevant details concerning the impossibility of squaring the circle. Consider then the ascription:

(7) George knows how to square the circle,

the related knowledge-of-answer statements:

(7KA^E) George knows every correct answer to the question as to how to square the circle

(7KA^S) George knows some correct answer to the question as to how to square the circle

and the corresponding knowledge-of-objects statements:

(7KO^E) For every way w in which he could square the circle, George knows that he could square the circle in w

(7KO^S) For some way w in which he could square the circle, George knows that he could square the circle in w .

If we were to treat (7) like a ‘knowledge wh’-ascription, then, as we have argued (and assuming as default the ‘some’-reading), (7) would have (7KA^S) as its truth condition. Unfortunately, given the facts of the matter, George does have an excellent answer to the question as to how to square the circle—namely, ‘In no way’! Furthermore, he is also in a position to give an even more informative answer along the lines of ‘In no way, since, as von Lindemann has proved, π is transcendental’. Indeed, whatever the range of correct answers to the question as to how to square the circle may exactly amount to in this case, the scenario can clearly be enriched (if it need be) in such a way as to have George in possession of all such answers. In such a situation, (7KA^S) (and (7KA^E)) is true (with (7KO^E)’s being vacuously true and (7KO^S)’s being false), and so, if semantically on a par with other ‘know wh’-ascriptions, (7) would have to be true as well. Unfortunately, in no way can (7) possibly be true—George does not know how to square the circle!¹⁹

It should be clear that the disaster is not avoided by switching, somehow arbitrarily, from the ‘some’-reading to the ‘every’-reading: given that, whatever the range of correct

¹⁹ Jonathan Schaffer has suggested to us that the question ‘How could you square the circle?’ could in some contexts carry the false semantic presupposition that there is a way to square the circle: hence, in such contexts, George could not know the answer to the question (since, in such contexts, every admissible answer would be positive). As we have already discussed in the text, we are rather sceptical about the prospects of ever counting such presuppositions as semantic. Be that as it may, even on this view there will also be contexts in which the question does not carry the false semantic presupposition (for example, if someone were offered a multiple choice of answers including ‘In no way. It is impossible to square the circle’). We think it clear that in these contexts, while it is true that the subject knows the answer to the question, attributing to the subject knowledge of how to square the circle would still be absurd.

answers to the question as to how to square the circle may exactly amount to in this case, the scenario can clearly be enriched (if it need be) in such a way as to have George in possession of all such answers, $(7KA^E)$ will be true in some extension of the scenario—thus forcing Stanley & Williamson’s analysis to admit the truth of (7)—while, in such an extension, George will still not know how to square the circle. True, $(7KO^S)$ is false in all these problematic situations, since it requires the existence of a way in which George could square the circle but no such way exists. Could Stanley & Williamson then simply by-pass the intuitive analytical step constituted by $(7KA^S)$ and rather opt for a straightforward analysis of (7) in terms of $(7KO^S)$? Indeed they could. But, no matter what the independent merits of such analysis might be, this move would come at a very high dialectical price in our context, for it would completely undercut Stanley & Williamson’s only argument in favour of their analysis—namely, its conformity with contemporary semantic theory (and with general data concerning ‘know wh’-ascriptions). Knowledge-how is not knowledge of a correct answer.

Finally, in contrast to knowledge-that and knowledge-wh, it is interesting to note that an analogous requirement that some way exist is present for a wide variety of *abilities*, like e.g. those variously denoted by ‘can’: if George can square the circle, then there is some way of squaring the circle. The existence requirement is thus an important aspect in which, against the intellectualism encapsulated in (R^{SW}) , knowledge-how falls in the realm of practice rather than in the realm of theory.

4. SINGULARITY

A prominent feature of propositional attitudes like knowledge-that, belief-that, desire-that etc. is their capacity for what have come to be known as *singular contents*.²⁰ As is well-known, an exhaustive specification of what these exactly amount to is a matter of great controversy in the philosophy of mind and language, but a minimal, fairly neutral and perforce sketchy characterization will suffice for our purposes. Let us say that a propositional attitude has a singular content iff it is *directly about some particular objects*. Of course, the crucial work in this characterization is being done by the phrase ‘directly about’, which, as it stands, is not much more than a place-holder for a fuller articulation of the relevant intentional relation.²¹ Still, we take it that the philosophically trained reader

²⁰ We’ll be using ‘content’ and its like interchangeably with ‘proposition’ and its like.

²¹ We stress that our understanding of direct aboutness is deliberately minimal (and hence hopefully sufficiently neutral). In particular, we don’t wish to presuppose any strict correlation between direct aboutness at the level of thought and what has been called ‘*direct reference*’ at the level of language (that is, very roughly, the property that an expression might have of picking out its referent without the mediation of a condition that the referent satisfies, see Kaplan [1989]). Direct aboutness is thus compatible with the object’s being thought about with the mediation of a (possibly *sui generis*) condition that the object satisfies (like for example a “singular sense”, see McDowell [1977]). Nor do we wish to equate without further ado propositional attitudes towards singular contents with so-called “*de re*” attitudes: for example, a typical Italian peasant’s belief that Bangkok is an Asian city, expressed by using the proper name ‘Bangkok’, is directly about Bangkok, but it’s not clear whether the peasant is thereby *believing of Bangkok* that it is an

has an intuitive understanding of and good judgement about such relation, and that these will suffice in what follows. To calibrate intuitions, on such understanding, a piece of knowledge most accurately expressed by an utterance of the sentence ‘Carla is gorgeous’ (where ‘Carla’ is hereafter used as an ordinary proper name referring to the wife of the French President in 2008)—i.e. the piece of knowledge that Carla is gorgeous, which is typically arrived at by perceiving Carla and its relevant physical features—is singular. It is directly about a particular individual, Carla Bruni. On the other hand, a piece of knowledge most accurately expressed by an utterance of the sentence ‘The wife of the French President in 2008 is gorgeous’ (where the definite description ‘The wife of the French President in 2008’ is used “attributively”, in Donnellan [1966]’s sense)—i.e. the piece of knowledge that the wife of the French President in 2008 is gorgeous, which is typically arrived at by hearing or reading that the same person who is the wife of the French President in 2008 is also a gorgeous person—is not singular. Even though there is a particular individual, Carla Bruni, that uniquely satisfies the description, that piece of knowledge is not directly about her, consisting rather in the general information that whoever is the wife of the French President in 2008 is also gorgeous: such information is so much about Carla Bruni as it is about Juliette Binoche.

Arguably, propositional attitudes with singular contents present spectacular failures of closure principles under *known* logical consequence. Talking with you about your new acquaintance Carla, who—unbeknownst to him—is in fact Carla Bruni, Elia might know from your testimony that, if Carla is gorgeous, her husband Nicholas will be jealous of her. Elia might also know by looking at pictures in newspapers and magazines that Carla is gorgeous, but these two pieces of knowledge (plus his knowledge that the proposition that Carla is gorgeous and the proposition that, if Carla is gorgeous, her husband Nicholas will be jealous of her entail the proposition that Nicholas will be jealous of Carla) will not put him in a position to know²² that Nicholas is jealous of Carla by drawing the relevant *modus ponens* inference. This is so, we may suppose, even though Elia is an acute logician. Thus, arguably, the following principle of *closure of knowledge-that under known logical consequence* fails:

(CKT^K) If *s* knows that P_0 , that P_1 , that $P_2 \dots$ and *s* knows that $\langle P_0 \rangle$, $\langle P_1 \rangle$, $\langle P_2 \rangle \dots$ entail $\langle Q \rangle$,²³ then *s* is in a position to know that Q

once we allow singular contents as values of ‘ P_0 ’, ‘ P_1 ’, ‘ P_2 ’ ...²⁴

Asian city (or, possibly equivalently, whether the peasant is thereby *believing Bangkok* to be an Asian city). *De re* thought might have stronger access requirements than merely singular thought.

²² Throughout, by ‘being in a position to know that P ’ we’ll mean ‘being in such an epistemic position that only a straightforward reflection is still needed for one to come to know that P —that is, if anything is still needed at all’ (see Williamson [2000], p. 95). Admittedly, such a notion remains very vague, but will do for our purposes.

²³ We use ‘ $\langle \varphi \rangle$ ’ to denote the proposition denoted by the phrase ‘the proposition that φ ’.

²⁴ The example might be resisted if one adopts an extremely fine-grained individuation of singular contents, such that the content expressed by an utterance of the sentence ‘Carla is gorgeous’ in the context of our conversation about your new acquaintance is different from the content expressed by an utterance of the very same sentence in the context of my reading and thinking about celebrities (even though the name ‘Carla’ refers to the very same person in both contexts). Some Fregean approaches to the famous Paderewski case of

Equally interesting failures of closure principles under *obvious* logical consequence are present once one pays due attention to a prominent feature of singular contents, namely that one cannot entertain any propositional attitude towards them simply on the basis of, roughly speaking, one's general linguistic and conceptual competence: in addition, one needs to be in *some non-trivial kind of cognitive relation* to the objects that they are directly about. For example, in order for one to entertain the singular content that Carla is gorgeous, it is not sufficient that one be able in some sense to single out Carla Bruni, for instance by means of the definite description 'the wife of the French President in 2008'—one needs to be in some non-trivial kind of cognitive relation to Carla Bruni, such as the one afforded by the competent use of the name 'Carla' as ultimately grounded in Carla Bruni herself, or the one afforded by seeing Carla Bruni on the street, or the one afforded by Nicholas' talk about his wife etc.²⁵ Consider a Scotsman and an Inuit, Calum and Kirima respectively, such that Calum does not stand in the relevant cognitive relation to Kirima, and hence cannot entertain singular contents about Kirima. Informally speaking, Calum has never heard about Kirima. Suppose further that Calum, an impeccable logician, knows that everything is self-identical. The proposition that everything is self-identical entails, we may assume, the proposition that Kirima is self-identical, indeed obviously so by the fundamental rule of inference of universal instantiation.²⁶ Yet, Calum is not in a position to know that Kirima is self-identical by drawing the relevant universal-instantiation inference, since he lacks the ability to entertain the content that Kirima is self-identical. Thus, arguably, the following principle of *closure of knowledge-that under obvious logical consequence* fails:

Kripke [1979] do adopt such an individuation. An assessment of the ultimate tenability of this extreme measure lies beyond the scope of this paper and is quite unnecessary for the argument of this section.

²⁵ There are at least two factors on which one's ability to entertain a singular content about Carla Bruni (or anything else!) depends. One factor concerns the *fine-grainedness* of singular contents: the more fine-grained one individuates the content that Carla is gorgeous (contrasting it for example with the content that that woman (pointing at Carla Bruni on the street) is gorgeous), the more stringent the requirements will be for entertaining the content that Carla is gorgeous. Another factor concerns the *strength of the cognitive relation to the relevant objects*: the stronger the cognitive relation to Carla Bruni has to be (from merely having a description or name for Carla to being closely acquainted with her), the more stringent the requirements will be for entertaining the content that Carla is gorgeous. The exact nature and possible context dependence of the relevant non-trivial kind of cognitive relation need not concern us here. It will suffice for the argument in this section to go through that the cognitive relation be of a non-trivial kind or other, so that a subject with a general linguistic and conceptual competence may still fail to be able to entertain a certain range of singular contents (such as those involving Carla Bruni).

²⁶ A lot of distracting noise is caused here by the fact that the (suitably regimented version of the) argument form 'Everything is *F*. Therefore, *a* is *F*', even though valid in *classical* first-order logic, is arguably strictly speaking invalid (consider the instance 'Everything exists. Therefore, Vulcan exists'—setting aside Meinongian views, the premise is true while the conclusion is false). A *free* logic of some sort or other should be used in order to detect invalidities generated in natural language by the use of what arguably are non-referring proper names (and other singular terms). However, as these problems are clearly orthogonal to our concerns in this paper, we simply wish to ignore existence issues, assuming that every singular term has a referent. For those who remain suspicious, we should like to point out that a more clumsy example would anyway go through with no assumption about existence: the argument 'Snow is white. Therefore, either snow is white or Kirima is fishing' is valid in all standard (bivalent) free logics, and obviously so, and its premise is known by Calum, yet he is not in a position to know its conclusion.

(CKT⁰) If s knows that P_0 , that P_1 , that $P_2\dots$ and it is obvious that $\langle P_0 \rangle$, $\langle P_1 \rangle$, $\langle P_2 \rangle\dots$ entail $\langle Q \rangle$, then s is in a position to know that Q

once we allow singular contents as values of ‘ Q ’.²⁷

The stage is now ready for lodging our argument from singularity. In short, the trouble with (R^{SW}) is that it entails that knowledge-how should suffer from the same closure failures as knowledge-that does (in particular, from failures of (CKT⁰)), whereas there seems to be all evidence to the contrary. We show this by focussing on a peculiarly strong and yet quite widespread closure feature of knowledge-how, involving indefinite descriptions used with universal (or generic) force. Go back to Calum and Kirima. Enrich the scenario by adding that, alas, Calum is also quite a proficient and experienced serial killer, very well acquainted with all sorts of murdering methods, ranging from strangling to shooting. In such a situation, the sentence:

(8) Calum knows how to kill a man

is undoubtedly true. Moreover, the sentence:

(9) Calum knows how to kill Kirima

is also arguably true, and *just for the same reason*: the general competence acquired by Calum with the relevant murdering methods, if sufficient for giving him knowledge how to kill a man in general, is certainly also sufficient for giving him knowledge how to kill Kirima in particular.

Of course, on the by far most natural reading,²⁸ (8) would still be true even in a situation where an individual exists who, while belonging to the species *Homo Sapiens*, is for some reason or other invulnerable by all, as it were, common or garden murdering methods mastered by Calum—even though having generalising force, the indefinite description ‘a man’ as occurring in (8) has *generic* rather than *universal* force, just as ‘A penguin is black and white’ is not false in a situation where, in addition to the standard wealth of black and white penguins, an extraordinary genetic mutation has caused one single penguin to be red. Thus, in order for the inference from (8) to (9) to be truth preserving, we should assume that Kirima is a typical specimen of *Homo Sapiens*. That assumption is unproblematic and will be tacitly understood as holding in what follows. Under these conditions, the truth preserving nature of the inference from (8) to (9) can be even better appreciated by further enriching the Calum/Kirima scenario with the additional assumption that a species exists all of whose specimens are indeed invulnerable by all murdering methods mastered by Calum. Let us pick one particular such specimen and call it ‘Athantos’. In such a situation, the sentence:

²⁷ Note that it is singular contents as *premises* that create trouble for (CKT^K), while it is singular contents as *conclusions* that create trouble for (CKT⁰).

²⁸ We are inclined to believe that the most natural reading is in the case also the only admissible reading (so that on no reading would (8) be false in the situation to be described in the text), but wish to remain officially neutral on this issue as nothing in our argument hinges on it.

- (10) Calum knows how to kill Kirima, but does not know how to kill Athanatos

is clearly true, but its truth evidently requires the truth of (9).²⁹

Extrapolating from this particular case, we can conclude that, quite generally, knowledge-how obeys the following principle of *closure of knowledge-how under indefinite instantiation*:

- (CKH^{II}) If s knows how to $A_{\text{an } F}$ and d is a typical³⁰ F , then s knows how to $A_{d/\text{an } F}$,³¹

where the designated occurrences (see fn 31) of the instance of ‘an F ’ are “safe”. Interesting and challenging as a fully precise characterisation of the operative notion of safety of a designated occurrence may be, for our purposes we can leave such notion at an intuitive level; suffice it to say that the restriction to safe designated occurrences is needed in order to rule out blatant falsehoods arising from complex embeddings of the instance of ‘an F ’ in the instance of ‘ A ’, such as ‘If s knows how to make Calum believe that a man is approaching, then s knows how to make Calum believe that Kirima is approaching’.

On the other hand, (R^{SW}) offers the following truth condition for (8) in terms of knowledge-that:

- (8^{SW}) For some way w in which he could kill a man, Calum knows that he could kill a man in w ,

²⁹ Since, under the present assumptions, there is nothing peculiar about Kirima (as far as killing goes) and the standard grounds for asserting (9)—the general competence acquired by Calum with the relevant murdering methods—are just as good grounds for asserting (8) (which, given (CKH^{II}), is strictly stronger than (9)), an assertion of (9) would usually violate the plausible conversational maxim of asserting the stronger. Hence, an assertion of (9) would usually implicate that there is something peculiar about Kirima (as far as killing goes). We submit that whatever oddity an assertion of (9) might be felt to have is due to this fact (as can be tested by cancelling the implication). Thanks to Felix Mühlhölzer for discussion of these issues.

³⁰ To stress, the qualification ‘typical’ is supposed to take care of cases in which the object has some special feature that prevents application of the knowledge-how to it, thus constituting a counterexample to an unqualified version of the principle (for example, (8) would still be true even if a man existed who is invulnerable to all murdering methods mastered by Calum). As Dylan Dodd pointed out to us, the specification could perhaps still be deemed insufficient, as an object might be a typical F in most respects, while being still peculiar in very few others, where, again, such a “local” peculiarity happens to prevent application of the knowledge-how (for example, the man invulnerable to all murdering methods mastered by Calum may well be typical in most other respects relevant to humanity). Such far-fetched cases could be ruled out e.g. by relativising typicality to respects and consequently weakening (CKH^{II}) with the qualification that d is a typical F in the relevant respects. The details of how this can best be done are however clearly inessential to our main line of argument (which ultimately only requires the truth of (9) against the falsity of (9^{SW})).

³¹ ‘ A ’, ‘an F ’ and ‘ d ’ are schematic, respectively, for an infinitive verb phrase, an indefinite description and a singular term. ‘ $A_{\text{an } F}$ ’ is schematic for an infinitive verb phrase with designated occurrences of the indefinite description instantiating ‘an F ’. ‘ $A_{d/\text{an } F}$ ’ is schematic for the result of substituting occurrences of an instance of ‘ d ’ for the designated occurrences of the instance of ‘an F ’ in the instance of ‘ $A_{\text{an } F}$ ’.

and the following truth condition for (9) in terms of knowledge-that:

- (9^{SW}) For some way w in which he could kill a man, Calum knows that he could kill Kirima in w .

For the very same reasons as those that determine the failure of (CKT^O), (8^{SW}) cannot however possibly entail (9^{SW}). Letting ‘Murdery’ denote a suitable witness for (8^{SW}), Calum’s knowledge of the general content that he could kill a man in Murdery by no means guarantees knowledge of the singular content that he could kill Kirima in Murdery, given that, in our situation, Calum fails to be in the relevant non-trivial kind of cognitive relation to Kirima, and hence is not even in a position to entertain—let alone know—the singular content about Kirima that Calum could kill her in Murdery. Indeed, these considerations show not only that (8^{SW}) fails to entail (9^{SW})—they also show that, in our situation, the latter is false. The former circumstance refutes (R^{SW}), since, according to that principle, (8^{SW}) should entail (9^{SW}) (given that, as we have argued, (8) entails (9)). The latter circumstance also refutes (R^{SW}), since, according to that principle, (9^{SW}) should be true (given that, as we have argued, (9) is true). Notice that, for the very same reasons, (8^{SW}) fails to entail (9^{SW}) even if it is assumed that ‘Calum could kill a man in Murdery’ in some sense actually “entails” ‘Calum could kill Kirima in Murdery’:³² the general failure of (CKT^O) shows that even an entailment from ‘Calum could kill a man in Murdery’ to ‘Calum could kill Kirima in Murdery’ (indeed, even an obvious entailment from the former to the latter) will not suffice to generate an entailment from (8^{SW}) to (9^{SW}).

We have developed our argument from singularity exploiting knowledge-how’s closure under indefinite instantiation. However, with some more ingenuity, the argument could have been centred upon a closure principle under obvious logical consequence for knowledge-how. We briefly sketch how such a modified argument would go. In order to state the knowledge-how analogue of (CKT^O), we would ideally need a definition of what it is for the semantic values (modal possibility-like properties) of a set of ‘to’-infinitive verb phrases to entail the semantic value of another ‘to’-infinitive verb phrase (see section 2 for the modal interpretation of ‘to’-infinitive verb phrases in the constructions of interest here). We take it that such a notion of entailment is naturally modelled on the notion of entailment among the semantic values of sentences (propositions) and that, for our present purposes, we have enough an intuitive grasp of it as to obviate the lack of a precise definition: ‘to eat’ entails ‘to eat or drink’, ‘not to shout’ entails ‘not to shout and play football’, ‘to shut the door’ entail ‘to shut something’ etc. (but ‘to sit’ and ‘to stand’ do not entail ‘to sit and stand’, just like the proposition that it is possible that Calum sits and the proposition that it is possible that Calum stands do not entail the proposition that it is possible that Calum sits and stands). Availing ourselves of this new notion of entailment,

³² Such an assumption would hold in effect that instantiation of a safe occurrence of an indefinite description with a term denoting a typical specimen is indeed an entailment. In so doing, the assumption would go against contemporary orthodoxy in at least two ways: in maintaining that entailment sometimes depends on the semantic values of occurrences of non-logical expressions (such as ‘Kirima’), and in maintaining that entailment sometimes does not require necessary truth preservation (since Kirima might well have been an atypical Inuit).

we can then state the following principle of *closure of knowledge-how under obvious logical consequence*:

(CKH⁰) If s knows how to A_0 , to A_1 , to A_2 ... and it is obvious that $\langle \text{to } A_0 \rangle$, $\langle \text{to } A_1 \rangle$, $\langle \text{to } A_2 \rangle$... entail $\langle \text{to } B \rangle$, then s is in a position to know how to³³ B .³⁴

(CKH⁰) is arguably correct, for reasons similar to those for which (CKH^{II}) is correct. Rather than expanding on a defence of (CKH⁰), we note that it justifies the inference from (8) to:

(11) Calum knows how to kill a man or please Kirima,

while the truth condition in terms of knowledge-that offered by (R^{SW}) for (8), namely (8^{SW}), fails to entail the truth condition in terms of knowledge-that offered by (R^{SW}) for (11), namely:

(11^{SW}) For some way w in which he could kill a man or please Kirima, Calum knows that he could kill a man or please Kirima in w .

Moreover, (11) is true in the Calum/Kirima scenario while (11^{SW}) is false.

It is important to note that failure of knowledge-that to be closed under indefinite instantiation is mirrored by an analogous failure of knowledge-wh to be so closed. We may further enrich the Calum/Kirima scenario by adding that, having diligently followed the geography class while at school, Calum knows where an Inuit lives. However, and again for the very same reasons as those that determine the failure of (CKT⁰), this piece of knowledge-wh on Calum's part by no means guarantees the other piece of knowledge-wh on his part which would consist in knowing where Kirima lives: clearly, in our situation, failing to be in the relevant non-trivial kind of cognitive relation to Kirima, Calum does not know where Kirima lives (indeed, failing to be in the relevant non-trivial kind of cognitive relation to Kirima, Calum is not even in a position to entertain any hypothesis as to where Kirima lives).³⁵ Similar examples also show that knowledge-wh is not closed under obvious logical consequence (presupposing a suitable notion of entailment for the semantic values of 'wh'-clauses): informed by a knowledgeable and highly reliable source that Hamish is the one and only one person that likes everything, Calum knows who likes everything, but does not know who likes Kirima.³⁶

³³ We write 'in a position to know how to' merely to respect uniformity with (CKT⁰). Admittedly, contrary to knowledge-that and knowledge-wh, we don't have a clear idea of what the distinction between knowing how and being in a position to know how could possibly amount to (in the relevant sense of 'being in a position to know', see fn 22). We take this to be yet another symptom of the difference between knowledge-how on the one hand and knowledge-that and knowledge-wh on the other hand.

³⁴ We use ' $\langle \text{to } \Phi \rangle$ ' to denote the property denoted by the phrase 'the property of possibly Φ ing'.

³⁵ A slight variation on the example suffices to drive the point home also for infinitive knowledge-wh (as opposed to finite knowledge-wh): Calum knows where to find an Inuit, but does not know where to find Kirima.

³⁶ Again, a slight variation on the example suffices to drive the point home also for infinitive knowledge-wh (as opposed to finite knowledge-wh): informed by the same knowledgeable and highly reliable source that

Finally, in contrast to their observed failures for knowledge-that and knowledge-wh, it is again interesting to note that analogous closure principles under indefinite instantiation and closure under obvious logical consequence do hold for a wide variety of *abilities*, like e.g. those variously denoted by ‘can’: if Calum can kill a man, then he can kill Kirima and, furthermore, he can kill a man or please Kirima (just like if a poisonous mushroom can kill a man, then it can kill Kirima and, furthermore, it can kill a man or please Kirima). The validity of (CKH^{II}) and (CKH^O) is thus another important aspect in which, against the intellectualism encapsulated in (R^{SW}), knowledge-how falls in the realm of practice rather than in the realm of theory.

5. INFERENCE

If knowledge-how were a species of knowledge-that, it would be such as to put a subject in a position also to know that certain other things are the case, just as more usual knowledge-that does—it would serve as basis for *acquiring, via* deductive and non-deductive inference, *further knowledge-that*. The undeniable role of knowledge-that as a basis for acquiring further knowledge-that *via* inference is one of the main factors that makes closure principles for knowledge-that so appealing. Even if, as shown in section 4, the most simple-minded formulations of closure principles for knowledge-that are subject to straightforward counterexamples, we’ll assume that a more sophisticated, satisfactory formulation of a closure principle of knowledge-that under logical consequence can be achieved. We won’t attempt to give such a formulation here, but we’ll simply assume that it exists and call it ‘*Closure*’. We stress that, while Closure has just been introduced in a such a way as not to be subject to the cheap counterexamples of section 4, it too may well turn out to be in the end untenable *in its full generality* (recall e.g. the very deep and structural problems raised by preface-paradox situations). Still, something like Closure does have to hold *in a sufficiently wide range of circumstances* if deduction is to preserve its traditional role in the acquisition of new knowledge-that, and this indisputable restricted validity of Closure is all we’ll need to assume in the following.³⁷

(R^{SW}) thus predicts that knowledge-how should equally well serve as a basis for acquiring further knowledge-that. However, this prediction is problematic in a number of ways. Consider first the following conjunction:

(AC^{KT}) John does not know that he is not a brain in a vat (BIV), but he knows that he has hands.

Hamish is the one and only one person to whom everything should be mentioned, Calum knows whom to mention everything, but he does not know whom to mention Kirima.

³⁷ We’ll continue to talk of “closure” even though the issues involved with the acquisition of knowledge-that *via* deduction cry out for drawing the usual distinction between closure and *transmission* (see e.g. Wright [2000]) and for focussing on the latter rather than the former. Fortunately, the scope of our argument is so restricted as to allow us to disregard the otherwise important differences between these two epistemic properties.

DeRose [1995] dubbed this ‘the abominable conjunction’, pointing out that it constitutes an unacceptable result of certain denials of Closure (see e.g. Nozick [1981], pp. 167–288). Whatever your ultimate view on Closure’s validity in such cases is, we think it’s clear that (AC^{KT}) at least *sounds* very bad. This judgement is shared by many epistemologists. For instance, Stanley [2005], pp. 67–69 objects to epistemological contextualism on the grounds that conjunctions like (AC^{KT}) should sometimes be acceptable if ‘know’ were a genuine context dependent expression, but they never seem to be such. Substituting a ‘know wh’-ascription for the second conjunct would not at all make the conjunction more acceptable, as witnessed by:

(AC^{KW}) John does not know that he is not a BIV, but he knows where his friend lives.

By contrast, consider:

(AC^{KH}) John does not know that he is not a BIV, but he knows how to swim.

The striking fact is that (AC^{KH}) , far from being abominable, does not sound bad at all. To be clear, it does not “sound bad” not in the sense that it is not the case that, in most usual situations where the second conjunct is true, the first one is not, for, presumably, most usual situations where John knows how to swim are also situations where John knows that he has hands. Rather, (AC^{KH}) does not “sound bad” in the sense that it does not strike us as presenting an *internal contradiction*, or even a tension, between its two conjuncts—what, quite to the contrary, both (AC^{KT}) and (AC^{KW}) do. One appealing explanation of this disanalogy is, of course, that while knowledge-wh at least involves relevant knowledge-that, and therefore similarly licences knowledge-that extending inferences, knowledge-how does not involve any relevant knowledge-that, and therefore there is nothing relevant one could infer from it. Such an explanation is however clearly unavailable to a defender of (R^{SW}) .

Still, of course, the defender of (R^{SW}) can look for alternative explanations. As we see things, there are in particular two other explanatory strategies that enjoy some initial plausibility. According to (R^{SW}) , John knows how to swim iff, for some w in which John could swim, John knows that he could swim in w . Let’s call ‘Swimmy’ the particular way that witnesses the (alleged) truth of the *analysans*. The first explanatory strategy consists then in maintaining that the belief that John could swim in Swimmy is typically *implicit*, and therefore not typically available as a starting point for inferences. This would however be a very weak reply. Firstly, the intuition about (AC^{KH}) would not change at all if it were specified that John is a swimming teacher (we are assuming that swimming teacher John has an explicit belief that he could swim in Swimmy if anyone has). Secondly, the pattern of intuitions relating to (AC^{KT}) , (AC^{KW}) and (AC^{KH}) would not change if we were to substitute uniformly ‘is in a position to know’ and ‘is not in a position to know’ for ‘knows’ and ‘does not know’ respectively.

A further thought in this connection is that the defender of (R^{SW}) could here appeal again to practical modes of presentation. Since, according to (R^{SW}) , the proposition that

John could swim in Swimmy is known by John under a practical mode of presentation, it may be argued that John is somehow blind to inferences involving it essentially together with other propositions that he knows under a *different* mode of presentation (see section 4 for an illustration of failure of a closure principle in such cases). If this is supposed to be a fully general principle, however, it would seem *ad hoc*: sometimes, we are not blind in this way to inferences involving propositions that we know under different modes of presentations. For example, one of the authors can effortlessly infer from the proposition that *Daniele* is not in London (presented under a third-personal mode of presentation) to the (same) proposition that *he himself* is not in London (presented under a first-personal mode of presentation) and *vice versa*. It would thus be utterly mysterious why, in contrast, propositions known under a practical mode of presentation are never similarly available to the speaker's cognitive system in *any* (AC^{KH})-like case, and it would cast doubt on them being known at all in anything like the standard sense of 'know'. Of course, as we've already seen in section 4, in some cases a difference in modes of presentation does prevent the inference from being available. So, for example, the author just mentioned might be an amnesiac and fail to know that he himself is Daniele, in which case he will neither be inclined to infer the proposition that he himself is not in London from the proposition that Daniele is not in London nor *vice versa*. Typically, we submit, the unavailability of a certain inference will be easily explained by failure on the part of the subject to know, or even believe, an identity claim (such as the claim expressed by 'I am Daniele'). However, no reason seems to be forthcoming as to why a similar failure should occur in *all* (AC^{KH})-like cases. We will come back to issues revolving around different modes of presentation later in this section.

Meanwhile, the more promising explanatory strategy consists in maintaining that the proposition that John knows when he knows how to swim, namely the proposition expressed by:

(12) John could swim in Swimmy,

does not after all entail—not even together with trivial propositions such as the proposition that BIV don't have hands—the proposition that John is not a BIV. This might be because (12) is *modal*. One interpretation of its truth condition that would block the entailment is given by the metaphysical-possibility claim:

(12^{MP}) It is metaphysically possible that John swims in Swimmy.

Of course, (12^{MP}) would be true even if John were a BIV, because there surely is a metaphysically possible world—for instance, any world that is similar to (what we take to be) the actual one—in which Swimmy, which we are imagining to be one of the standard ways in which human beings could swim in normal conditions, is a way for John to swim. However, to suppose that, according to the defender of (R^{SW}), (12^{MP}) exhausts the content of John's knowledge-how leads to quite unacceptable results. For suppose that there is a way (let's call it 'Swimmy*') that is not a way in which John could swim, but is nevertheless such that it is metaphysically possible that John (given appropriate changes in human biology, in the environment, in the laws of nature etc.) swims in it. Suppose also

that John practised swimming in a virtual-reality machine specifically devised so as to make Swimmy* a good way of virtually swimming. Suppose finally that John is well aware of all this and does not know of any other possible way of swimming. Then, presumably, John knows (under a practical mode of presentation) that it is metaphysically possible that he swims in Swimmy*, even though he does not know how to swim, which would be very much against the reductionist *spirit* of (R^{SW}).³⁸

The defender of (R^{SW}) could improve by offering a truth condition for (12) intermediate in strength between a simple metaphysical-possibility claim and a simple claim about the actual world, for example by going *counterfactual*, as in:

(12^C) If John were to try to swim in Swimmy, he would succeed.

This kind of interpretation certainly seems to work for some cases: for example, in order for John to know how to drive his car, it will not be required that he is in a position to know that the car has not been stolen, or destroyed, since the time he parked it this morning. Indeed, the car might have been stolen, or destroyed, and it would still be true that he knows how to drive it. On the present proposal, this is so because, for some way *w* in which he could drive his car, he knows that, if he were to try to drive his car in *w*, he would succeed. In effect, to help ourselves to a very influential semantic analysis of counterfactual talk, in the closest possible worlds in which he tries, he does succeed. Still, the counterfactual reading is not generally adequate. Consider again the case of poor Diego whose legs have been chopped off (see fn 12): it is arguably false that, for some way *w* in which he could play football, if he were to try to play football in *w*, he would succeed; yet, as Stanley & Williamson concur, it is still plausibly true that he knows how to play football.

It is therefore not clear to us how exactly (12) should be analysed. However, any reading that, as it seems unavoidable in order to avoid (12^{MP})-like troubles, is strong enough as to be incompatible with the environment's having certain unfavourable properties, will not help with our problem. To see that, think of a possible world, call it 'NSW' (for 'No Swimmy World'), in which trying to swim in Swimmy would result in certain drowning: this is so because, while human beings are relevantly similar in NSW, water, or whatever "waterish" substance exists in NSW, behaves very differently from the way water behaves in the actual world. Let us then consider the possibility that John is in NSW. Given the lesson of (12^{MP}) about the required strength of (12), we take it that such a possibility is incompatible with (12) (the reader who thinks otherwise is welcome to substitute in the following whatever possibility she does think is incompatible with (12)). Let us also suppose that, being aware of the incompatibility, John is considering this possibility. Consider the conjunction:

(AC^{KH+}) John is not in a position to know he is not in NSW, but he knows how to swim.

³⁸ If not the letter, since, as we've stated it, in order for it to be satisfied the *analysans* of (R^{SW}) would also require that Swimmy* is a way in which John could swim, which it is not.

We think that (AC^{KH+}) does not sound bad (we'll discuss shortly one possible source of distraction about this). If this is so, the initial disanalogy between (AC^{KT}) and (AC^{KW}) on the one hand and (AC^{KH}) on the other remains ultimately unexplained from the point of view of a defender of (R^{SW}) .

There is, admittedly, an important complication in the case just discussed, pursuing which will both strengthen our intuitive judgement about (AC^{KH+}) and lead us to another related argument. In a case of the kind just described, John will typically *know that he knows* how to swim. He also knows that, if NSW is realised, he could not swim in Swimmy. Typically, he will also realise that if he could not swim in Swimmy, he does not know how to swim. Therefore, by Closure, if he knows that he knows how to swim, he should be in a position to know that he is not in NSW. Now, any residual doubts against the coherence of (AC^{KH+}) , we submit, are due to this further typical and yet distracting feature of the case (that is, that John knows that he knows how to swim). To illustrate the point, let us consider a different case, in which the second-order knowledge is absent.³⁹

Mary, the physicist, believes that she does not know how to swim. As a matter of fact, she was actually taught how to swim when she was 3-year old, but she has somehow lost her memory of this learning experience. She has not however lost her competence: if she were to be thrown in water, she would instinctively perform the right movements. Indeed, she sometimes dreams of being in water and swimming by performing those movements. In this situation, it seems true that she knows how to swim. According to (R^{SW}) , Mary knows how to swim iff, for some w in which Mary could swim, Mary knows that she could swim in w . Let's call 'Floaty' the particular way that witnesses the (alleged) truth of the *analysans*, presumably the way involving the movements that Mary instinctively makes. Let it also be understood that, while both denoting Floaty, 'Floaty^P' is associated with the relevant practical mode of presentation, whereas 'Floaty^T' is associated with a theoretical physical description fully grasped by Mary. A few times, we may suppose, Mary has deliberately tried out Floaty's movements out of sheer curiosity—while walking around, or perhaps even in the air (while she was suspended, on purpose, on the roof of her laboratory).

Despite being in many other respects a very good and well-informed physicist, let us also suppose, Mary does not know much about water: she knows of course that water is H_2O , but she does not know enough about the properties of H_2O to know whether it permits swimming in Floaty^T (we may suppose that Mary has spent her recent years locked in her laboratory, that most of her knowledge dates from this period and that she hasn't conducted any experiment concerning water: for all Mary presently knows, water is a sticky treacherous liquid in which any terrestrial animal would drown). Given her general good knowledge of physics, Mary does however know the proposition expressed by the following sentence to be true:

(WN) If I could swim in Floaty, then water is Non-Sticky,

³⁹ We should note that, assuming (R^{SW}) , the following case would also therefore constitute a plausible counterexample to the "KK-principle" (roughly, if one knows that P , one is in a position to know that one knows that P).

where ‘Non-Sticky’ stands for the negation of a disjunction of properties that Mary knows are incompatible with any form of natation and floating. It seems clear that, while knowing the proposition expressed by (WN), Mary—who doesn’t know very much about the pattern of instantiation of properties in the world outside the lab, and in particular about which properties are co-instantiated with water—is in no position to infer the proposition expressed by its consequent, which represents a very substantial piece of knowledge about water. That water is Non-Sticky just isn’t the kind of thing that Mary can come to know simply by deduction without conducting further experiments. Given Closure (and Mary’s excellent mastery of *modus ponens*), that in turn seems to entail that she does not know the proposition expressed by (WN)’s antecedent to be true. But she knows how to swim, and hence (R^{SW}) is false.

Here, the intellectualist is likely to appeal again to *differences* in modes of presentation: indeed, it might be said, Mary does not even believe that she could swim in Floaty^T (as opposed to believing that she could swim in Floaty^P), and so it’s not at all surprising that she is not in a position to infer that water is Non-Sticky. She believes that, if she could swim in Floaty^T, water is Non-Sticky, but she believes that she could swim in Floaty^P. She fails to recognise that the latter proposition is the antecedent of the former. Now, granting for the sake of argument that the notion of belief under a practical mode of presentation can be made good sense of, we think that this reply faces the following problem. We observed above that, in general, failure of availability of inferences involving propositions known under different modes of presentation is explained by failure of recognising, under these different modes of presentation, some components of the propositions as being the same. But there is no very plausible candidate for such a failure of recognition in our example. The most plausible one would be Floaty itself. We specified however that Mary has tried out Floaty’s movements (which presumably gives her access to Floaty as Floaty^P) and that she can also give an accurate description of the movements in a physical vocabulary (which presumably gives her access to Floaty as Floaty^T); we can now add that she does recognise the identity between Floaty^P and Floaty^T. Indeed, she might accept the proposition expressed by (WN) while thinking about Floaty as Floaty^P: she might be trying out Floaty’s movements and then realise that that way (Floaty^P) is just the such-and-such physical way (Floaty^T), and so is such that, if she could swim in it, water is Non-Sticky. Therefore, the reply is *ad hoc*: there is no plausible reason why Mary should fail to recognise that the proposition that she could swim in Floaty^P is the antecedent of the proposition that, if she could swim in Floaty^T, water is Non-Sticky.

It might be thought that we’ve been overlooking a further feature of Mary’s belief: its being only implicit. However, all Closure presumably requires is that Mary *be in a position to* know that water is Non-Sticky. If she does have knowledge that she could swim in Floaty, it seems that this requires that, if she were competently to infer that water is Non-Sticky from her (implicit) belief that she could swim in Floaty, she would come to know this. But that counterfactual is clearly false. That water is Non-Sticky just isn’t the kind of thing that Mary can come to know simply by deduction without conducting further experiments. This is another way to see that, if knowledge-how were a species of

knowledge-that, it would be mysteriously barren—incapable of serving as basis for acquiring further knowledge-that *via* inference.

Finally, in contrast to the knowledge-that and knowledge-wh, it is once again interesting to note that an analogous inferential isolation does exist for a wide variety of *abilities*, like e.g. those variously denoted by ‘can’: even if John can jump over the hedge in your garden, this does not put him in a position to know that there is a hedge in your garden. Inferential isolation is thus yet another important aspect in which, against the intellectualism encapsulated in (R^{SW}), knowledge-how falls in the realm of practice rather than in the realm of theory.

6. GRADABILITY

As has been exposed in section 2, it is crucial for Stanley & Williamson’s argumentative strategy for showing that *knowledge-how* is a species of knowledge-that that the intermediary step that *knowledge-wh* is a species of knowledge-that be secured. All of our objections so far have been directed at the move from this assumption to the conclusion that knowledge-how is also a species of knowledge-that. Indeed, we have been at pains to emphasise common aspects with respect to which knowledge-that and knowledge-wh behave differently from knowledge-how. In closing, we wish however to discuss briefly some evidence that casts doubt on the assumption that knowledge-wh is a species of knowledge-that. Although this evidence falls short of providing a conclusive argument against the reducibility of knowledge-wh to knowledge-that, we think that it does at least give some reason to think that the reduction is not as obvious as it is usually thought to be, and that further arguments are needed in order to establish it.

Stanley [2005] considers briefly a possible objection to the thesis that knowledge-that is not *gradable*, based on the fact that knowledge-wh and knowledge-how seem to be clearly gradable. The same data, holding fixed the assumption that knowledge-that is not gradable, could be used to question the assumption that knowledge-wh is reducible to knowledge-that. We think that Stanley [2005] provides a good case for the former assumption, and we therefore take it for granted. Here are a few examples of felicitous grading that Stanley [2005], p. 42 considers:

- (13a) John knows how to swim well
- (13b) John knows how to ride a bicycle better than Mary does
- (13c) Hannah knows where Texas is better than John does.

Stanley’s own explanation of the felicity of (13a)–(13c) is that what is being compared (in the case of (13a), evaluated) is the quality of the answers that the subjects have to the indirect questions.

Let us first consider grading of knowledge-wh. Surely, we do use sentences like (13c) in the way Stanley suggests, e.g. in a situation where Hannah can say that Texas borders on Mexico plus the other US states of Louisiana, Arkansas, Oklahoma and New Mexico, while John can only say that Texas is somewhere in the US. However, it seems that an utterance of (13c) would also be felicitous in a situation where Hannah and John have the same answer to the question as to where Texas is (say, the answer that Texas is in the US), but Hannah knows that answer on very good grounds (say, she has been told so by a world expert of US geography and has double-checked this with another ten authorities), while John only knows the answer on worse but still knowledge-conferring grounds (say, he has read so on Wikipedia). Note that this creates a clear disanalogy between knowledge-that and knowledge-wh. In the latter situation, while an utterance of (13c) would be felicitous, an utterance of the following sentence would not be:

(13c^{KT}) Hannah knows that Texas is in the US better than John does.

Here is one more example that clearly confirms the pattern:

(13d) Hannah knows who will come to the party better than John does.

An utterance of (13d) can be felicitous in a situation where Hannah has a more accurate answer to the question as to who will come to the party (say, the answer that at least George, Bill and Ronnie will come to the party, which is more accurate than John's answer that at least George will come to the party). However, an utterance of (13d) can also be felicitous in a situation where Hannah and John have the same list of names of people coming to the party (say, they both had it from Mary, a trustworthy source), but, while John only read the list, Hannah actually called each person on the day of the party to double-check.

The evidence indicates that Stanley's explanation of the gradability of 'know wh'-ascriptions is (partially) inadequate. What is being compared can be both the quality of *the answers* and the quality of *the epistemic relations that the subjects have to those answers*. Similar, if slightly more unusual, examples can be given when what is being compared are not different answers that different subjects have to the same question, but different answers that the same subject has to different questions, or even different answers that different subjects have to different questions.

One of the most frequent arguments against knowledge-that and knowledge-wh expressing different relations is that they seem to fail some familiar tests for ambiguity, such as disallowing zeugmas where a single occurrence is used in the same sentence in the two senses and being expressed by different words in some other languages. However, one who wants to defend the non-reducibility of knowledge-wh to knowledge-that needn't posit any strong form of ambiguity here: clearly, the senses would have to be systematically related. It would therefore be a case of *polysemy* rather than *homonymy*;⁴⁰

⁴⁰ Some theorists wish to count both homonymy and polysemy as (different) forms of ambiguity, others prefer to reserve the word 'ambiguity' for the former phenomenon (for a survey, see Ravin and Leacocke

the tests mentioned before, however, are notoriously not very effective in detecting polysemy—they systematically yield false negatives. Take the case of ‘book’: an utterance of ‘That thought-provoking book has been written by Gilbert Ryle in 1949 and has dust on its cover’ is typically felicitous, and there are no languages that we know of which reserve two different words for book types and book tokens. To take a case more related to our main topic, consider the sentence ‘Mary is justified both in believing that $1 + 1 = 2$ and in eating healthy food’. The prevalent view, we take it, is that the kinds of correctness that the two conjuncts are concerned with are fundamentally different: epistemic justification in one case and practical justification in the other one. Maybe the prevalent view is not correct; but it is certainly not refuted by the felicity of the sentence just quoted, nor by the fact that, as far as we know, there are no languages that reserve two different words for epistemic and practical justification.

7. CONCLUSION

Our aim in this paper has not been to give a positive account of, or positive arguments for, a view on knowledge-how on which it does not turn out to be a species of knowledge-that. Rather, we have considered a number of objections to the opposite view, intellectualism, and in particular to the specific version of the view which we judge most interesting (and which is most prominent in the contemporary literature)—namely, the one ingeniously developed and defended by Stanley & Williamson [2001]. Stanley & Williamson’s argument for intellectualism is semantic: it draws on the alleged uniformity of ‘know how’- and ‘know wh’-ascriptions, both of which are best treated, they claim, along the lines of standard semantics for indirect questions (with a little help from practical modes of presentation).

We believe that we have provided good evidence to think that this assimilation is problematic. To sum up, the three main reasons to this effect are the following. Firstly, in contrast to ‘know wh’-ascriptions, ‘know how’-ascriptions with known negative answers are false: for instance, while ‘No one’ can be a correct answer to the question as to who was at the party and ‘In no way’ is a correct answer to the question as to how one could square the circle, possession of the former answer provides sufficient grounds for attributing to someone knowledge of who was at the party, but possession of the latter answer provides no grounds at all for attributing to someone knowledge of how to square the circle. Secondly, knowledge-how obeys closure principles whose counterparts clearly fail for knowledge-that and knowledge-wh when the contents involved are singular: for instance, Calum’s knowledge that he could kill a man in a certain way w does not give him knowledge that he could kill Kirima in w , but Calum’s knowledge of how to kill a man does give him knowledge of how to kill Kirima. Thirdly, knowledge-how seems, in complete disanalogy with knowledge-that and knowledge-wh, to be inferentially isolated from further knowledge-that: for instance, however the content of the knowledge-that which allegedly constitutes knowledge of how to swim is specified, it seems that standard

[2000]). ‘Ambiguous’ is ambiguous. The choice is however clearly irrelevant for our purposes: whether or not it is a form of ambiguity, polysemy is not usually detected by the proposed tests.

principles of closure for knowledge-that would have to fail for it. There is precious little, arguably, that you can infer from your knowledge-how (although you can clearly infer a lot from your knowledge that you have knowledge-how).

Finally, we've indicated some evidence which casts doubt even on the reducibility of knowledge-wh to knowledge-that. It's a well-known fact that knowledge-wh enjoys some forms of gradability, while knowledge-that plausibly does not. Stanley tries to explain away this semantic evidence arguing that the gradability of knowledge-wh only concerns the quality of the answer that the subject has to the indirect question rather than the quality of her epistemic relation to the answer. However, it seems that felicitous grading persists in cases in which two subjects can only give exactly the same answer to a certain question, which provides good evidence that the gradability of knowledge-wh can also concern the quality of the epistemic relation that the subject bears to the answer.

Even though we haven't had space to develop the point at length here, we take our arguments—pointing at crucial semantic and epistemic disanalogies between knowledge-how and knowledge-that—not only to cast into serious doubt the adequacy of Stanley & Williamson's particular analysis, but also to present a quite general challenge to intellectualist views.

Unifying all forms of human cognition is a grandiose task; our suspicion however is that there is a humble and too simple answer to the question of how to accomplish it—pretty much the same one as the answer to the question of how to square the circle.

REFERENCES

- Boër, S. & Lycan, W. [1986], *Knowing Who*, MIT Press, Cambridge MA.
- DeRose, K. [1992], 'Contextualism and Knowledge Attributions', *Philosophy and Phenomenological Research* **52**, pp. 913–929.
- DeRose, K. [1995], 'Solving the Sceptical Puzzle', *The Philosophical Review* **104**, pp. 1–52.
- Donnellan, K. [1966], 'Reference and Definite Descriptions', *The Philosophical Review* **75**, pp. 281–304.
- Ginet, C. [1975], *Knowledge, Perception, and Memory*, Reidel, Dordrecht.
- Goldman, A. [1976], 'Discrimination and Perceptual Knowledge', *The Journal of Philosophy* **73**, pp. 771–791.
- Groenendijk, J. & Stokhof, M. [1982], 'Semantic Analysis of *WH*-Complements', *Linguistics and Philosophy* **5**, pp. 175–233.

- Groenendijk, J. & Stokhof, M. [1997], 'Questions', in van Benthem, J. & ter Meulen, A. (eds), *Handbook of Logic and Language*, MIT Press, Cambridge MA, pp. 1055–1124.
- Hintikka, J. [1976], *The Semantics of Questions and the Questions of Semantics*, North-Holland, Amsterdam.
- Kaplan [1989], 'Demonstratives', in Almog, J., Perry, J., Wettstein, H. (eds), *Themes from Kaplan*, Oxford University Press, Oxford, pp. 481–563.
- Karttunen, L. [1977], 'Syntax and Semantics of Questions', *Linguistics and Philosophy* **1**, pp. 3–44.
- Karttunen, L. & Peters, S. [1976], 'What Indirect Questions Conventionally Implicate', in Mufwene, S., Walker, C., Steever, S. (eds), *Papers from the Twelfth Regional Meeting*, Chicago Linguistic Society, Chicago, pp. 351–368.
- Katz, J. [1972], *Semantic Theory*, Harper & Row, New York.
- Keenan, E. & Hull, R. [1973], 'The Logical Presuppositions of Questions and Answers', in Petöfi, J. & Franck, D. (eds), *Präsuppositionen in Philosophie und Linguistik*, Athäneum, Frankfurt a.M., pp. 441–466.
- Kripke, S. [1979], 'A Puzzle about Belief', in Margalit, A. (ed.), *Meaning and Use*, Reidel, Dordrecht, pp. 239–283.
- Lewis, D. [1979], 'Attitudes *De Dicto* and *De Se*', *The Philosophical Review* **88**, pp. 513–543.
- McDowell, J. [1977], 'On the Sense and Reference of a Proper Name', *Mind* **76**, pp. 159–185.
- Nozick, R. [1981], *Philosophical Explanations*, Harvard University Press, Cambridge MA.
- Perry, J. [1979], 'The Problem of the Essential Indexical', *Nous* **13**, pp. 3–21.
- Ravin, Y. & Leacock, C. [2000], 'Polysemy: An Overview', in Ravin, Y. & Leacock, C. (eds), *Polysemy. Theoretical and Computational Approaches*, Oxford University Press, Oxford, pp. 1–29.
- Ryle, G. [1949], *The Concept of Mind*, Chicago University Press, Chicago.
- Ryle, G. [1971], 'Knowing How and Knowing That', in Ryle, G., *Collected Papers II*, Hutchinson, London, pp. 212–225.
- Russell, B. [1912], *The Problems of Philosophy*, Williams and Norgate, London.

Schaffer, J. [2007], ‘Knowing the Answer’, *Philosophy and Phenomenological Research* **75**, pp. 383–403.

Stanley, J. [2005], *Knowledge and Practical Interests*, Oxford University Press, Oxford.

Stanley, J. & Williamson, T. [2001], ‘Knowing-How’, *The Journal of Philosophy* **98**, pp. 411–444.

Williamson, T. [2000], *Knowledge and its Limits*, Oxford University Press, Oxford.

Wright, C. [2000], ‘Cogency and Question-Begging: Some Reflections on McKinsey’s Paradox and Putnam’s Proof’, *Philosophical Issues* **10**, pp. 140–163.

Zardini, E. [2007], ‘Knowledge-How, True Indexical Belief, and Action’, ms.

APPENDIX: KNOWLEDGE–HOW AND CONTEXT DEPENDENCE ON NON-TRADITIONAL FACTORS

Many contemporary epistemological discussions are shot through with examples that *seem* to demonstrate the context dependence⁴¹ of ‘know that’-ascriptions on *non-traditional factors* such as *salience of counterpossibilities* and *practical stakes*. In this appendix, we aim to show that, while ‘know how’-ascriptions too are arguably sensitive to some non-traditional factors or others, they do not seem to be sensitive to the non-traditional factors on which ‘know that’-ascriptions do seem to be sensitive, at least not to the same extent. To wit, the intuitions that arise in certain kinds of cases for ‘know that’-ascriptions do not arise in analogous kinds of cases for ‘know how’-ascriptions. Independently of the hard question as to whether such intuitions should be trusted in the case of ‘know that’-ascriptions and hence lead to a revision of the standard invariantist picture of such ascriptions, this asymmetry in intuitions will suffice to constitute yet another piece of evidence against (R^{SW}).

We start by observing that ‘know how’-ascriptions seem to be dependent on context: an utterance of ‘John knows how to play chess’, if John knows the rules and nothing more, would seem true when uttered in a context where we are only looking for someone who can offer a first introduction to the game to a six-year old child, while it would seem false when uttered in a context where we are looking for someone whom Kasparov can train with. Indeed, the example shows that ‘know how’-ascriptions are sensitive to factors (such as the purposes of the situation at hand) that have not been traditionally thought to be relevant to epistemic assessment. Of course, this kind of *prima facie* context dependence can be explained in different ways, just as the abundant literature of the last decade has shown for the *prima facie* context dependence of ‘know that’-

⁴¹ To be clear, by ‘context dependence of a sentence φ ’ we simply mean that different utterances of φ in different contexts can have different truth values.

ascriptions. The point we are interested in here is whether the *prima facie* context dependence of ‘know how’-ascriptions is sensitive to the *same factors* and to the *same extent* as the *prima facie* context dependence of ‘know that’-ascriptions is. Such factors can for our purposes be usefully summarised as consisting in either the salience of counterpossibilities or in practical stakes.⁴²

Both kinds of factors can then be tested for context dependence by changing their value for either the ascriber of ‘know’ or for the subject to which ‘know’ is ascribed. In the space of this appendix, we limit ourselves to offering a couple of what we believe are simple and yet telling examples to test the alleged context dependence of ‘know how’-ascriptions on practical stakes, in particular cases where the ascriber coincides with the subject. We believe that the *insensitivity* uncovered can be shown to hold quite generally both for ascriber and subject, and that it can also be shown to extend to the salience of counterpossibilities (again, for both ascriber and subject), but will not attempt to generalise these results here—the relative insensitivity of knowledge-how (against the seeming sensitivity of knowledge-that on the same range of cases) will be sufficient to provide the final piece of evidence against (R^{SW}).

Consider first two variant cases of the much discussed bank example (taken with only very few alterations from Stanley [2005], pp. 4–5—the bank example goes back to DeRose [1992]):

High Stakes. Hannah and her wife Sarah are driving home on a Friday afternoon. They plan to stop at the bank on the way home to deposit their paychecks. Since they have an impending bill coming due, and very little in their account, it is very important that they deposit their paychecks by Saturday. But as they drive past the bank, they notice that the lines inside are very long, as they often are on Friday afternoons. Hannah notes that she was at the bank two weeks before on a Saturday morning, and it was open. But, as Sarah points out, banks do change their hours. Hannah says, ‘I guess you are right, I don’t know that the bank will be open tomorrow’.

[...]

Ignorant High Stakes. Hannah and her wife Sarah are driving home on a Friday afternoon. They plan to stop at the bank on the way home to deposit their paychecks. Since they have an impending bill coming due, and very little in their account, it is very important that they deposit their paychecks by Saturday. But neither Hannah nor Sarah is aware of the impending bill, nor of the paucity of available funds. Looking at the lines, Hannah says to Sarah, ‘I know the bank will be open tomorrow, since I was there just two weeks ago on Saturday morning. So we can deposit our paychecks tomorrow morning’.

The intuition we have according to Stanley is that Hannah’s utterance of ‘I don’t know that the bank will be open tomorrow’ in *High Stakes* is true, while her utterance of ‘I know that

⁴² To stress, what makes the alleged examples of context dependence so interesting is at least in part the fact that such factors—as against many other factors, such as e.g. time—have not been recognised by traditional epistemological theorising as influencing the truth values of ‘know that’-ascriptions.

the bank will be open tomorrow' in *Ignorant High Stakes* is false. Taking these intuitions to be solid, at least for the sake of the argument, we wish to see whether and how they change when the object of the attribution is changed from knowledge-that to knowledge-wh or knowledge-how.

The cases are very easy to adapt for knowledge-wh: just imagine that what Hannah says is rather 'I don't know [I know] when the bank will be open tomorrow'. It seems clear to us that in this case there is no difference in intuitions: the context-dependence intuition persists that Hannah would speak truly in (suitably modified) *High Stakes* and falsely in (suitably modified) *Ignorant High Stakes*. Knowledge-wh perfectly aligns with knowledge-that.

However, we find that the pattern drastically changes once we move to knowledge-how. Let us consider the two following modified cases:

High Stakes Knowing-How. Hannah and her wife Sarah are driving home on a Friday afternoon. They plan to stop at the bank on the way home to deposit their paychecks. Since they have an impending bill coming due, and very little in their account, it is very important that they deposit their paychecks by Saturday. But as they drive past the bank, they notice that the lines inside are very long, as they often are on Friday afternoons. Hannah and Sarah know full well that the bank is going to be open on the Saturday. However, since they moved quite recently to the area, there might be a worry about getting the way from their new place to the bank wrong on Saturday morning. Of course, on Saturday morning they would get to the bank sooner or later, but, if they were to get lost while driving there, they would probably miss the bank's opening hours. Hence, being there, they might take the opportunity to deposit the paychecks now. As it happens, however, Hannah actually did manage very well to drive from their new place to the bank two weeks earlier. Sarah nevertheless points out that people sometimes misremember routes and end up getting the wrong way. Hannah replies, 'I guess you are right, I don't know how to get to the bank'.

Ignorant High Stakes Knowing-How. Hannah and her wife Sarah are driving home on a Friday afternoon. They plan to stop at the bank on the way home to deposit their paychecks. Since they have an impending bill coming due, and very little in their account, it is very important that they deposit their paychecks by Saturday. But neither Hannah nor Sarah is aware of the impending bill, nor of the paucity of available funds. Hannah and Sarah know full well that the bank is going to be open on the Saturday. Looking at the lines, Hannah says to Sarah, 'I'll come to deposit the paychecks tomorrow. I know how to get to the bank from our new place, I've done that two weeks ago'.

Keeping in mind that, as in the original case it was supposed to be true that the bank was open on Saturday morning, in this case Hannah would drive to the bank without making errors on Saturday morning, it seems clear to us that the intuitions that are now elicited are exactly the opposite to those elicited for knowledge-that and knowledge-wh: Hannah would speak falsely in *High Stakes Knowing-How* and truly in *Ignorant High Stakes Knowing-How*. What intuitions track for knowledge-how, in general, seems to be simply the presence of some sort of competence. What counts as a high enough level of proficiency to possess the competence, in turn, can depend on the context (as for the

competence in playing chess), but not to the same extent as the correctness of 'know that'-ascriptions seemingly does.

The same conclusion could be reached by looking at the following further data. Consider the retrospective assessment we would obtain in *High Stakes* if Hannah and Sarah decided to take the risk, and found, as we know, that the bank is open on that Saturday morning: it does not seem that Hannah could then truly say, 'What I said yesterday turned out to be true: I knew that the bank was going to be open!' (analogously for knowledge-wh). Even if Hannah's belief turned out to be true, this does not seem to show she knew all along that the bank would be open on Saturday. By contrast, in *High Stakes Knowing-How*, having successfully driven to the bank on the Saturday following the Friday of the conversation, it would seem that Hannah could then truly say, 'What I said yesterday turned out to be true: I knew how to get to the bank!'.