

Epistemic Closure, Home Truths and Easy Philosophy*

There is a long-standing and well-known dispute in epistemology regarding whether knowledge is closed under known entailment. The general idea of such closure (“CLR”) is this:

CLR1: For all persons S, necessarily, if S knows both that p and that p entails q , then S knows that q .

In spite of the intuitiveness of CLR, there has been a decades-old stalemate regarding whether it is true, largely because some of the “Moorean” things we seem not only to know, but to know easily (like that I am sitting on a green chair)¹ seem clearly to entail “heavyweight” philosophical things. These deep, philosophical items (like that I am not actually lying in bed dreaming or a brain in a vat (“BIV”) being fed green-chair stimulation through a tube) would seem to be of a nature that, if we can know them at all, we surely cannot *easily* know them. There is general agreement on the sorts of things that may be easily known and the sorts of deep philosophical truths that surely may not be. But on one side of the CLR dispute we find those who believe that because certain things are both known and known to entail items that not only may not be easily known, but arguably may not be known at all, CLR must be false; while on the other side of this debate we find those who, believing CLR to be true, either infer that the two required instances of knowledge (that p and that p entails q) are not both simultaneous with and appropriately connected to each other, or take the standards for what is required to know p to be so different from those required to know q , that no reasonable CLR can bridge the equivocal instances of “knowledge.”² In this paper, I hope to demonstrate that two widely accepted facts

* Thanks are due to Larry Tapper, Jim Van Cleve, Ron Allen, J.P. de Mouy, and an extremely helpful and acute anonymous referee at this JOURNAL.

¹ I will sometimes call these “home truths.”

² The apparent failure to know that q in spite of knowledge both that p and that *if p then q* has been taken to be a function of the differing “relevant alternatives” to ostensibly known items in different contexts. See, for example, Fred Dretske, “Epistemic Operators,” this JOURNAL, lxvii, 24 (December 1970): 1007-1023 and Robert Nozick, *Philosophical Explanations* (Cambridge Mass.: Belknap Press, 1981), pp. 172-247). John Hawthorne has suggested that an apparent failure of closure might also result

about what we do and do not know—facts with which I believe any minimally acceptable understanding of the concept of knowledge must comport—are jointly inconsistent with the truth of CLR.

Before I divulge the two uncontroversial premises or data elements upon which I shall depend for making my argument against CLR, let me try to clear up a couple of anticipatory objections: first, that whether or not CLR holds is strictly (or at least largely) a function of how we define it; second, that whether or not CLR holds is strictly (or at least largely) a function of how we define (or explicate or “rationally reconstruct”) *knowledge*. For if one assumes that we may analyze or reconstruct CLR and knowledge any way we like, it would seem to follow that we could define one or another of these concepts in ways that could either require CLR’s truth or make it obviously false. I will consider these objections in turn.

Certainly there have been a wide variety of proposals regarding how one ought to define CLR.³ But the basic idea, which many find intuitive, should be clear. After all, if I know both that my chair is green and that anything that is green must be colored, it would seem that I must also know that my chair is colored. As most conceptions of knowledge require belief, however, and dispositions to belief can be idiosyncratic, it has been recognized that the concurrent knowledge that *p* and that *p* entails *q* might *not* result in the belief that *q*, even if it arguably ought to. So weaker versions of CLR, wherein belief has been strained out, have been created. This may be achieved by indicating that S *would* know *q* if he or she were to believe it. To wit:

CLR2: For all persons S, necessarily, if S knows both that *p* and that *p* entails *q*, then, if S were to believe that *q*, then S would know that *q*.

But now, since there is no telling *why* S happens to believe that *q* and might do so for silly reasons, additional amendments have been made to require S to believe that *q because of the*

from what he calls “metaphysical anxiety.” See his “The Case for Closure,” in Matthias Steup and Ernest Sosa, eds., *Contemporary Debates in Epistemology* (Oxford: Blackwell, 2005), pp. 26-42

³ For a hefty batch of them, see John A. Barker and Fred Adams, “Epistemic Closure and Skepticism,” *Logos and Episteme*, 1, 2 (2010): 221-46 and Steven Luper, “Epistemic Closure,” in Edward N. Zalta, ed., *The Stanford Encyclopedia of Philosophy* (Spring 2016 Edition),

<http://plato.stanford.edu/archives/spr2016/entries/closure-epistemic/>.

known antecedent and relevant entailment rather than for some other reason(s). Another way of achieving that goal is to shift the talk from knowledge to justification or warrant, since knowledge is commonly thought to be “built up” out of truth, belief, and either sufficient evidence, reasons, or the reliability of some source (call any of these “back-up” for short). The following version reflects this concern:

CLR3: For all persons S, necessarily, if S has sufficient back-up to know both that p and that p entails q , then S has sufficient back-up to know that q .

I note that all three versions above contain three items said to be known—or to be sufficiently “backed-up” to be known if both true and believed. As suggested above, the literature on CLR is full of descriptions of various sorts (and levels) of standards for “know” that depend on the situation of the ostensible knower and/or the use(s) to which the knowledge in question will be put.⁴ This has led a number of contextualists to argue that everyday Moorean

⁴ CLR could thus invite equivocation: one might even explicitly use two or three separate quantifiers to bind the variables for what might turn out to be several types of knowledge. Obviously, however, with as many as three different versions of “know” in our CLR, it may be quite difficult to determine the interaction between the principle and various skeptical claims—which may be relying on any one of the three—or even a fourth one! Furthermore, it is not only the “knowing”—or “competently deducing” (Hawthorne, *op cit.* at p. 29)—but “what is known” that may undergo subtle shifts as we move through the three instances of knowledge found in most CLR. Dretske, *op. cit.* at p. 1013 tells a story about his brother Harold, who, because of his (the brother’s) boorish behavior on a bus, has made it possible for a stranger to learn that Harold is not a nice person. So, this stranger now knows that Dretske’s brother is sometimes unpleasant on buses. She may know too, that if Dretske’s brother is ever unpleasant, then Dretske has a brother. But she nevertheless does not know that Dretske has a brother. Is this an example of a failure of epistemic closure? I think not. What is said to be known has shifted from an object of *de re* knowledge to an object of knowledge *de dicto*. The woman knows of the x that is Dretske’s brother that x is a boor, but does not know that the proposition *that Dretske’s brother is a boor* is true. Neither does it follow from what knowledge she does have that Dretske even has a brother. Since no one ought to expect CLR to underwrite referential fallacies of that sort, cases like this one should not be thought of as counter-examples to CLR.

statements require much less of knowers than do the anti-skeptical claims made by epistemologists.⁵

Although there have been many varying accounts of CLR, I hope the basic idea is clear: *when we know something, we also always know what we know to be entailed by that thing*. If we stay true to that basic idea, striving neither to enlarge nor constrict it inappropriately, I believe we may avoid any danger of making CLR false by definition. And indeed, the variations in CLR phraseology seem to have done little to settle the stalemate among (a) those who have inferred that CLR is false because they deny that we know apparently heavyweight conclusions in spite of our clear knowledge of the truth of the apparently lighter antecedents and conditionals; (b) those who believe the obviousness of both CLR and our ignorance of heavyweight metaphysics shows that in certain cases we simply cannot know both the antecedent and the conditional—at least not at the same time and in the same way; and (c) those who are comfortable with our knowledge not only of p and $p \rightarrow q$ in those contested cases, but also with what CLR entails for our epistemic relationship with the qs . Evidently, one’s choice of a particular CLR makes no important difference to the resolution of this stalemate, so for our purposes readers are free to use any of the versions above or any other version that may be preferred as we continue. Indeed, I think it makes sense to follow Nozick on this score when he cautioned that “It would be ill-advised to quibble over the details” of our presentation of CLR.⁶ We have only to ensure of any version that we choose that the basic idea expressed in CLR 1, CLR 2, and CLR 3 above is not lost.

The situation regarding the explication of *knowledge* is somewhat different from that surrounding the choice of a particular CLR, however. Certainly, with respect to the provision of necessary and sufficient conditions, the profusion produced by “knowledge” explicators has swamped that produced by CLR discussions.⁷ But, as indicated, because CLR is a technical principle, there are fewer constraints on how we may put it. Just as we are free to vary the

⁵ Among the numerous philosophers who have taken a contextual approach to objections to CLR are Stewart Cohen, Keith De Rose, David Lewis, and Michael Williams.

⁶ *Op. cit.* at p. 205.

⁷ In fact, I would hazard the proposal that no other term in English has received either the number of definitional offerings or the wealth of expository literature, although perhaps “love” or “goodness” would give “knowledge” a run for its money.

meanings of “know” within it, we may take “entails” to involve anything from material implication to necessary consequence. That is, we can make our CLR stronger or weaker, as we prefer, and then check to see if it seems true with respect to various types of knowledge. When we try to explicate knowledge itself, however, we find that there are additional constraints that we have no choice but to reflect, since “knowledge” is not a technical term, but a word in common use. If we violate those constraints, we will be analyzing something other than what is generally meant by the term.⁸

What are these limitations? As it turns out, the constraints I will focus on are precisely the two facts I will use to make my argument against CLR. First, like the vast majority of those who have ever written on epistemology, I will take as a fundamental starting point that that we often easily know such home truths as what our name is, that there are cows around here, that we once attended elementary school, that Mickey Lolich once pitched for the Detroit Tigers, and that we have two hands.⁹ Thomas Kelly¹⁰ has reminded us of David Lewis’s endorsement of this foothold in the latter’s “Elusive Knowledge”:

⁸ For a nice discussion of various types of knowledge, see Oswald Hanfling, *Philosophy and Ordinary Language: The Bent and Genius of our Tongue* (London: Routledge, 2000) at pp. 94-110. Naturally, we are here interested only in certain sub-types: we do not need to consider any such ostensible phenomena as “knowledge by acquaintance” or cases in which someone “just knows” something that subsequently turns out to be false. Compare also Richard Swinburne, “Review: Response to Warrant,” *Philosophy and Phenomenological Research*, LV, 2 (June 1995): 415-19 and Susan Haack, *Evidence and Inquiry: A Pragmatist Reconstruction of Epistemology* (New York, Prometheus Books, 2009); both of those philosophers consider attempts to provide necessary and sufficient conditions even for that particular sort of knowledge that is of interest here to be endeavors of utter futility.

⁹ If one agrees with Wittgenstein that such “hinge propositions” as that I have two hands strictly cannot be known because they cannot be doubted and it makes no sense to adduce additional evidence for them, then one may strike all such facts from my list. See his *On Certainty* (New York: Harper & Row, 1972). I here use “Moorean” to indicate every-day, common-sense “lightweight” material, but it may well be that some of the examples that Moore himself gave do not pass muster because of what we might call their “hinge weirdness.”

¹⁰ Thomas Kelly, “Moorean Facts and Belief Revision, or Can the Skeptic Win?” *Philosophical Perspectives*, xix, 1 (December 2005): 179-209

We know a lot . . . We have all sorts of everyday knowledge, and we have it in abundance. To doubt that would be absurd . . . It is a Moorean fact that we know a lot. It is one of those things that we know better than we know the premises of any philosophical argument to the contrary.¹¹

I take it that no philosopher need apologize for this opening move, since it has been made by epistemologists of nearly every stripe¹² at least since the time of Hume. It may be that a long road must be travelled before a child may be said to have mastered Moorean information, but most will acknowledge that once such mastery has been accomplished, learning new lightweight facts becomes easy—as “easy” is to be understood here. At any rate, enrollees in many schools—externalists, evidentialists, contextualists, Wittgensteinians, coherentists, foundationalists, (and various combinations and hybrids of them)—have largely been on board with people easily knowing vast numbers of ordinary things at most times of their adult lives. Indeed, all but (the paradoxical and likely reclusive) skeptics have generally agreed that we may relatively easily—at least compared to metaphysical propositions—know our own names, when we are sitting down, or that we once regularly attended elementary school. No doubt, various philosophers may require that we put these knowledge claims very particularly: they may insist, for example, that we may not be considering any epistemological issues at the time we make our lightweight claims, or they may take such apparently easily won knowledge to be limited to “internal” assertions rather than extend to “metaphysical realities.” But all except the skeptic will concede that if we are sufficiently careful, we can produce an indefinite number of statements regarding which there is certainly a consensus that we may know them without understanding any ontology or epistemology. Indeed, many philosophers will insist that anyone who denies we can ever know home truths must be using “know” in a highly unorthodox manner. In any case, I

¹¹ David Lewis, “Elusive Knowledge,” *Australasian Journal of Philosophy*, LXXIV, 4, (December, 1996): 549-67 at p. 549.

¹² Kelly, *op. cit.*, provides a list of philosophers who have accepted the primacy of Moorean home truths that includes Armstrong, Fine, Lycan, Harman, Pollack, Wittgenstein and Wright. The list could be extended indefinitely, I think—especially if we look back as far as Reid. As will become clear below, I share Kelly’s view that we may glean important epistemological information possibly unobtainable elsewhere from what would seem to be strictly metaphilosophical material.

will accept as a raw (and not itself particularly heavyweight) datum—and as my first constraint on explications of “knowledge”—that people easily know lots of lightweight propositions. This or that example of ostensible knowledge of a Moorean fact might turn out to be mistaken, but if nobody has ever known anything like their own name or that they are now sitting on a chair, discussions of knowledge will be at a standstill, indeed evidently pointless.

The second premise/constraint that I will rely upon involves a claim of metaphilosophy that has come to be almost as widely accepted in epistemological circles as that we sometimes know that we are sitting down. Indeed, it could be argued that a recognition of the distinction between what we are terming “home” and “heavyweight” facts partly depends upon this very recognition. It is that deep philosophical propositions are *not* the sorts of items that may be known easily; they must be significantly “harder won” than knowledge of Moorean items generally can be. Note that I say “harder won” rather than simply “hard won” because it may be that *every* sort of knowledge must be hard won in some sense. Perhaps knowing that we are sitting can only be the result of years of language learning, the mastering of proprioception, and so on. Indeed, it may well be that reaching knowledge of simple home truths has required millions of years of evolution. I take no position on that. I restrict myself here to a metaphilosophical assertion that implies that controversial philosophical theses, like those proclaiming the falsity of skepticism, cannot be easy to prove or otherwise come to know. We cannot, as we might with respect to a question about the presence of a chair, just point, or sit, or show a picture.¹³

¹³ Of course, Moorean items need not always be easily known. E.g., it might be quite difficult to come to know that a green chair is before one in a completely dark room, and it may always be relatively difficult for a blind person to know this. When such knowledge is difficult, however, it is not thereby made heavyweight. All heavyweight knowledge is uniquely (perhaps even *impossibly*) difficult, and none is Moorean, but not all knowledge of Moorean material need be easy. Thus, one reason that might be given for some proposition—whether intuitively heavyweight or not—to fail to be easy knowledge for some person S, is that S came by it through a demanding or arduous process. And a reason that an apparently easily reached conclusion might nevertheless be heavy is that at least one premise used to derive it is as well. So, for example, there is no incongruity for *There are physical objects* to be heavyweight even though someone’s deduction of it (say, from *Either there are physical objects or I am*

Naturally, this has all been noticed before. Consider Jonathan Vogel's famous story of Roxanne's easy inference that her gas gauge is reliable.¹⁴ According to Vogel's parable, Roxanne notes at t1 (when the tank is full) that the pointer is on "F," and infers that the tank is full; notes at t2 (when the tank is half-full) that the pointer is on "1/2," and infers that the tank is half-full; and so on. Vogel claims that if reliabilism is true, then Roxanne, without ever touching a dipstick, trying to put gas in a tank when the pointer is at F, asking an expert, or having any idea whether others are either driving her car, occasionally putting gas in it, or both, can infer by induction that her gas gauge is reliable.

Unsurprisingly, this consideration has wider than automotive application. William Alston had previously noted¹⁵ that if at t1, S forms the perceptual belief that p1, and p1 is true; and at t2, S forms the perceptual belief that p2, and p2 is true; and so on, then, if perception happens to be trustworthy, a simple reliabilist conception of knowledge seems to allow S to infer that sense perception is a dependable source of knowledge. Furthermore, epistemologists have pointed out that this is not a problem for reliabilism only. It is an acknowledgement of the problematic aspect of assuming reliability in order to demonstrate it, something that involves the nature of epistemic reliance generally. Indeed, Stewart Cohen has argued that any epistemology that does not approve this principle:

KR: A potential knowledge source K can yield knowledge for S, only if S knows K is reliable.

standing up; and *I am not standing up*) has been rudimentary. Since the conclusion is metaphysical and the inference is undemanding, we simply infer that at least one of the premises must be heavyweight.

¹⁴ "Reliabilism Leveled," this JOURNAL, xcvii, 11 (November 2000): 602-23.

¹⁵ In his "Epistemic Circularity," *Philosophy and Phenomenological Research*, XLVII, 1 (September 1986): 1-30 at p. 9.

will thereby suffer from the problem of Too Easy Heavyweight Knowledge.¹⁶ And, as indicated above, without the concern that there is a crucial distinction between such home truths as “This looks red” and such heavyweight claims as “Vision is an acceptable basis for (or method of achieving) knowledge,” it is hard to see why the debate over CLR ever got off the ground, let alone has had the staying power it has displayed.

Whether or not **KR** is a correct diagnosis of the failure of epistemic reliability, it is insufficiently general to apply to what I take to be importantly similar cases that are not explicitly about knowledge or warrant at all. For our recognition that something is awry seems to me exactly the same in cases that are not-specifically epistemological. When someone infers that *There is free will* from *John could have gone to the game if he had wanted to* or that *Universals exist* from *Sarah and Shelley share the property of being suspicious of others*, there is the same cognitive dissonance as is engendered by the Roxanne case. It is thus not only heavyweight epistemology that cannot be learned or proved easily. Just as waving one’s hands about or kicking a rock has not succeeded in providing a refutation of skepticism, displaying two red balls cannot refute nominalism. Rather, *all* such “proofs” have seemed to miss the point completely. It may be, as the contextualists urge, that the standards for philosophical knowledge are higher or that when the subject turns to metaphysical matters, what we once knew we somehow immediately know no longer. Whatever the cause, *something* seems to prevent heavyweight knowledge from being purchased cheaply. But it need not be, as one might infer from the Alston/Vogel/Cohen attack on epistemic reliabilism, anything involving *reliability* at all.

Roxanne’s story is nevertheless helpful when trying to determine just what it is that makes one example of knowledge easy and another difficult. Attributions of weight involve both the “how” of acquisition and the “what” of the prize. Knowledge that a gauge pointer is on $\frac{1}{2}$ may be easily obtained in the sense that even a child can manage it. It is Moorean in nature and thus may be deemed easy just in case it is obtained with little effort. But if Roxanne were to tell

¹⁶ See his “Basic Knowledge and the Problem of Easy Knowledge,” *Philosophy and Phenomenological Research*, LXV, 65 (September 2002): 309-29. James Van Cleve explains how KR creates problems for coherentism and “Two-Levels” conceptions in his “Is Knowledge Easy--Or Impossible? Externalism as the Only Alternative to Skepticism,” in Stephen Luper, ed., *The Skeptics: Contemporary Essays*. (Hampshire, UK: Ashgate, 2003), pp, 45-59.

us that she had easily and quickly learned that we are not BIVs, we would doubt her: it offends our constraint (and our intuitions) for such deep, controversial information to be easily acquired. We would insist that if the proposition in question were indeed easily derived from one or more others, at least one of the premises used must itself be heavyweight.

I believe I have now answered both anticipatory objections: First, we need not make CLR false by fiat, simply because nothing prevents the choice of a version that does not do that. Second, while a conception of knowledge *could* affect the truth-value of our chosen CLR, the constraints that we have settled on—(i) the commonplace existence of easily known Moorean truths, and (ii) the unavoidable difficulty connected with coming to know heavyweight truths—do not beg the question. Neither constraint seems to prevent epistemological disputants—whether internalists, externalists, evidentialists, tracking-theorists, foundationalists or coherentists—from retaining their favorite positions regarding the truth or falsity of CLR. Indeed, I take neither constraint to itself be a heavyweight epistemological or metaphysical claim, but consider each to be a rather commonplace starting point for discussion. The objection that CLR simply *must* be false if it can easily take us from easily known Moorean premises to a heavyweight proposition, ignores that we have to this point taken no position on either the weight of the internal conditional said to connect *external physicality* with such properties as *chairhood*, or on the ease with which the consequent may (given CLR) be known. Thus, there are several avenues by which CLR could be true without requiring either that we cannot easily know there are chairs or that we often easily know there are external physical objects.

In spite of my conviction that my second metaphilosophical constraint above is intuitive and relatively lightweight, it is clear that I ought to address some apparently clashing provisions found in the “Easy Ontology” theory. Amie Thomasson, perhaps the most prominent recent exponent of this position, has written,

[N]ot only can we know the answers to ontological questions, but that they are too easily resolved to be the subjects for extended and deep disputes....I hold that a great many ontological questions can be answered by making trivial inferences from uncontroversial premises. For example, we can make a trivial inference from ‘There are two cups on the table’ to ‘The number of cups on the table is two’, and so to ‘There is a number’; or from ‘The shirt is red’ to ‘The shirt has the property

of redness', and so to 'There is a property', thereby settling the questions of the existence of numbers and properties.¹⁷

In response to Thomasson's take on the matter, Stephen Yablo has written:

Whether or not P is a moot question [depends on whether] (a) P is ordinarily presupposed, not asserted, and (b) P is highly extricable from the claims that presuppose it. The mootness of ontological questions is the special case of this where P is an existence-claim....But how is this sort of question to be decided in general? I would not want to be the one who sifts through the proposed reductions to determine which run counter to common sense. A good part of philosophy consists of squabbling about precisely this —about which hypotheses are revisionary/eliminative and which the ordinary speaker will, or ought to, take in stride.¹⁸

J.W.N. Watkins pronounced a similar verdict on the paradigm case argument when he wrote of one of its purveyors:

Either the connotation of any expression to which he applies the argument is determined by the chief items among its denotation or its connotation is prior to its denotation....If he opts for the first alternative, he will transform those factual assertions he is most anxious to assert into tautologies, and you cannot kill metaphysics with this sort of paper shot.¹⁹

¹⁷ Amie Thomasson, "Quizzical Ontology and Easy Ontology," this JOURNAL, CXI, 9/10 (September/October 2014) 502-528 at pp. 502-3.

¹⁸ Stephen Yablo, "Carnap's Paradox and Easy Ontology," this JOURNAL, CXI, 9/10 (September/October 2014): 470-501 at pp. 490-493.

¹⁹ "Farewell to the Paradigm-Case Argument," *Analysis*, XVIII, 2 (December 1957): 25-33 at p. 33.

It is not my intention to come down in favor of the Yablo/Watkins side here, except to the extent that it supports the claim that heavyweight philosophy cannot be easy. Whether statements like “There are numbers” or “There are physical objects” may be made in an “internal” fashion is itself a deep philosophical question. If it is resolved along Thomasson’s neo-Carnapian line, what *seem* like quintessential metaphysical existence statements may sometimes be uncontroversially asserted, and we may thus engage in what she calls “easy ontology.” But surely they *can* and often *have* also been made to express heavyweight truths.²⁰ I believe, however, that, whether or not I am right about this, her theory is both important and controversial, and that this is so regardless of the specific details of the neo-Carnapian position she adumbrates.²¹

One must be careful here not to infer too much from philosophical controversy. Suppose, for example, that Rosalyn and Rollie both take note of the longstanding intractability of various philosophical disputes. Rosalyn deduces from this ancient controversy that Thomasson’s work on easy ontology must be wrong, while Rollie takes from it that, since traditional metaphysics seems to lack an agreed-upon decision procedure, Thomasson’s ordinary language approach is likely to be correct. Both seem to me to be inferring too much from the strictly metaphilosophical information they are relying on. All of the following seem to me clearly epistemically possible in light of that evidence: that, to date, all ontologists have been wrong; that a third of them have been right and the rest wrong, and that only Carnap and Thomasson have been correct. We cannot infer the reasonability or unreasonability of any particular

²⁰ For an excellent discussion of this matter, see Everett W. Hall, *Philosophical Systems: A Categorical Analysis* (Chicago: University of Chicago Press, 1960) at pp. 12-25.

²¹ Daniel Z. Korman, “Easy Ontology without Deflationary Metaontology,” *Philosophy and Phenomenological Research*, (Forthcoming) has argued against Thomasson’s “deflationary metaontology,” but his complaint is substantive in a way that my claim is not intended to be. He writes, for example, that “prolonged debates about existence questions “have a point....[namely] to assess strategies for reinstating or circumventing easy routes to knowledge once putative defeaters...have rendered them suspect.” I myself make no claim here that ontological disputes “have a point.” I push only what I take to be the much more obvious fact that these disagreements are stubborn: they have existed for a long time and show no signs of slackening.

ontology from the intractability of the debate alone; all we may properly infer from it is that the questions are essentially both difficult and contentious.

I expect that by this point many of my readers will have become impatient. The dissonance resulting from the combination of the easy knowledge of home truths and the unavoidable difficulty in knowing what those truths seem to entail are both known phenomena. And the role of CLR in this dispute is also common knowledge in epistemology. We have gained some ground, perhaps, but it is well-worn. At the start we had twice the following positions to choose from (assuming either the truth or the falsity of CLR):

(A) We know both that there are chairs (“ p ”) and that p entails that there are physical objects (“ q ”).

(B) We know p , but not that p entails q .

(C) We do not know p , but we do know that p entails q .

(D) We know neither that p nor that p entails q .

If we grant easy knowledge of Moorean facts, we knock out both versions (CLR true & CLR false) of (C) and (D). If we also insist there can be no easy knowledge of heavyweight truths, the truth-value of the CLR false version of (A) will be partly a function of whether the conditional is itself heavyweight (and how it has come to be known), while the CLR true version will depend both upon those matters and upon whether CLR here must “pass along” any ease with which the knowledge of its constituents has been gained.²² Thus, we are left with four perennially conflicting possibilities: both the CLR true and CLR false versions of (A) and (B). Is the condition hopeless? Is it impossible, even for those who accept the truth of Moorean claims and the impossibility of easy philosophy—to decide whether or not CLR is true?

²² According to the principle of the closure of easy knowledge (“ECLR”), if p and $p \rightarrow q$ are both known easily, then q must (or at least always *may*) also be easily known. As we shall see, I take the question of whether the conditional here may be known easily to be largely a function of the truth-value of a particular psychological thesis to be discussed below. But in any case, ECLR seems to me no less controversial than CLR, since some will be bound to claim that the conditional (like the antecedent) may be easily known, while agreeing with us that the entailed consequent can never be. For an approach to these issues that involves limitation by knowledge “source,” see Tim Black, “Solving the Problem of Easy Knowledge,” *Philosophical Quarterly*, LVIII, 233 (October 2008): 597-617.

I believe there is a way. The first part of the trick is to specify precisely a certain thesis that would make CLR undeniably problematic—even to contextualists. I will call this (somewhat Thomassonian) thesis “categorialism.” Put simply, categorialism is the general idea that some (sub-categorial) properties are such that, necessarily, for all persons S, if S ascribes²³ one of them to something, then S must simultaneously and in every context ascribe another (categorial) property to that thing. It is an “infusion” principle: some properties are necessarily inseparable from—because “telescoped” into—others. The truth or falsity of all such doctrines must be left open here, not only because of the danger of begging the question with respect to CLR, but also because it is a controversial (and heavyweight) matter whether or not the axioms of conceptual/categorial systems always infuse many of our everyday general terms and predicates. I take it that a determination of the truth-value of the thesis would involve an understanding of the process of language acquisition, the nature of assertion and presupposition, the role of concepts in thinking and communication, and countless other extremely perplexing questions.²⁴

There is more than one variant of the “conceptual infusion” position, however, so to carry out the strategy referred to above, we must be more precise than we have been to this point. For one thing, there may be a large number of conceptual “universes of discourse.” Perhaps basketball makes up such a universe, and nobody can assert that anyone has “dished out an assist” without ascribing a large number of sport-related properties. On my limited understanding of the ancient history of the philosophical concept of Category, the particular brand of big “C” Categorialism I wish to utilize follows Aristotle in involving only the broadest, most basic

²³ Here, and in the discussion that follows, I use “ascribe” and “assert” such a way that one cannot be properly said to ascribe a property or assert that it is exemplified without representing it in some manner.

²⁴ I have considered some of these in my “Metaphysical Realism and the Various Cognitive Predicaments of Everett W. Hall,” in Walter Horn, ed., *The Roots of Representationism: An Introduction to Everett Hall* (Saarbrücken, Germany: LAP Lambert, 2013), pp. 259-97. See also Hall, *op.cit.*; Hilary Putnam, *The Many Faces of Realism: The Paul Carus Lectures, Series 16* (Chicago: Open Court, 1988), particularly the material on Carnap and the Polish Logician; and Terry Horgan and Mark Timmons “Conceptual Relativity and Metaphysical Realism,” *Philosophical Issues*, XII, 1 (October 2002): 74-96. Of course, there are countless classic works on presupposition and conceptual schemes by Russell, Strawson, Wittgenstein, Quine, Davidson, Sellars, Fodor, and many others.

categories of thought, and is thus a position that is more clearly apposite here, where we are specifically concerned with knowledge of heavyweight propositions, than would be a small “c” categorialism that takes such concepts as banking or basketball to be telescoped into various subsidiary items.²⁵ In fact, I want to push the idea of Category to its extreme by defining “Categorialism” so narrowly as to limit claimed infusion to a single category, that of *external physical object*. I thus take an assumption of the truth of “Categorialism” to entail just this: where G is *being external and physical* and F is any such garden-variety property as *being a chair*, necessarily, for all persons S, if S ascribes F to something, then S must also ascribe G to that thing.

The relevance of Categorialism to the discussion of CLR should be obvious. Although we may arguably contemplate zebras without representing (consciously or unconsciously) vats, demons, or disguised mules, according to the Categorialist, we cannot contemplate zebras without in some sense representing the allegedly infused property of being an external physical object. But while Categorialism is, thus, a psychological thesis, to be useful here it must wield a sort of epistemic power.²⁶ I therefore include in my definition of the doctrine the following implications that it has for both knowledge and the ease of its attainment:

- (1) In all instances in which it is easily known that something is F, it will be easily known to a competent language user/reasoner that, necessarily, whatever F is G; and
- (2) In all instances in which it is easily known that something is F, if (i) it is easily known that whatever is F must be G, and (ii) CLR is true, then it will be easily known (to any competent cogitator) that there are G things.

Thus, if someone knows that she is sitting, Categorialism has the power to confer easiness upon any ensuing knowledge she obtains that physical objects exist. Perhaps this aspect

²⁵ In addition to this large “C” small “c” distinction, the Carnapian position, wherein the connections are thought to be “analytic,” may differ significantly from a Categorialist position according to which the categorial “infusions” result from (and in) a somewhat different sort of relation between the concepts. I will make no attempt to settle that matter, but simply require any version of Categorialism satisfying my definition to have the specified implications for easiness of knowledge.

²⁶ As Thomasson’s version does. See her *Ordinary Objects* (Oxford, UK: Oxford University Press, 2007) at p. 30.

of the position will seem more plausible if one considers that if it is false, even if we have easy knowledge that there are chairs, it may be possible for chairs to be coherently believed not to be physical,²⁷ a situation which would alone seem to make any knowledge of the $p \rightarrow q$ in the middle of CLR more difficult to come by. Of course, if Categorialism is true the unavoidably joint ascription of the relevant properties naturally suggests that knowing certain types of facts is a rather modest accomplishment. This, again, is psychology, but it is worth remembering here that, in spite of being paired with *knowledge* throughout this and other works ostensibly focusing on epistemological matters, *easiness* and *difficulty* are themselves psychological properties.

As I have claimed above that Categorialism is controversial, it might be objected that if that thesis is itself difficult to arrive at, easy knowledge cannot result from it. But it is important to recognize that Categorialism could be true without anybody believing it—never mind knowing it. Its operation may in this way be covert. Its truth would simply allow for the possible easy knowledge both of such conditional propositions as *If there are chairs, there must be physical objects* and, if CLR is true, also of such consequents as *Therefore, there are physical objects*.

Of course, for the existence of physical objects to be *easily* known, it must be known *simpliciter*, and Categorialism cannot quite guarantee that. Just as CLR is not a thesis about conceptual infusion, Categorialism is not a doctrine about (or that can alone promise) an external world.²⁸ It is, as a Spinozist might put it, a thesis regarding the connection of ideas only. It

²⁷ Coherent only in the sense that it seems only to be Categorialism that is responsible for the requirement that representations of F and G are always inextricably entwined, even when the representation of F is a function of a belief that some F item is *not* G. It might still be the case that any such belief is necessarily false—and so incoherent in that sense. I take it that the Categorialist (and any Wittgensteinian allies) may hold that one cannot “strip” the external physical object presuppositions out of terms like “chair” without rendering them meaningless, and that if (*per impossibile*) one *could* perform a successful “categoriotomy,” entailments to any alleged heavyweight propositions would not only be difficult, there is no reason to suppose that they would hold at all—with or without the assistance of CLR. As indicated above, such issues are deep and difficult.

²⁸ See, however, my “A New Proof of the Physical World,” *Philosophy and Phenomenological Research*, XLIV, 4 (June 1984): 531-537, which uses a related psychological principle to attempt a demonstration of the necessary instantiation of physicality.

cannot by itself produce knowledge—easy or difficult—that there are physical objects. It provides no more than the possibility of ease of the attainment of metaphysical knowledge when all other conditions of knowledge have been met. In spite of such frailty, however, the effect its truth would have on at least a century of literature on the natures of analyticity, aprioricity, understanding, ascription, representation, presupposition, and assertion ensures that Categorialism is itself a heavyweight doctrine.

As indicated, “conceptual infusion” cannot alone entail heavyweight knowledge of any non-conditional item (easy or not) when we know we are sitting—only (what at least non-Categorialists will deem to be) heavyweight *ascription*—that which (all should agree, because the definition requires it) has the power make for easy knowledge in any context of the conditional: an assertion of the connection between such concepts as *being a chair* and *physicality*. According to the Categorialist, it may be that both this conditional and the conclusion that there are physical objects are rendered lightweight by the infusion of physicality, while, pursuant to our second constraint, at least the knowledge of the consequent can never be easy. My definition of “Categorialism” should not be understood to determine—one way or the other—what propositions its proponents will likely take to be heavyweight. Rather, the doctrine should be taken to mean precisely that which would make the first premise of both the informal and formal arguments below true—and nothing more.

With this understanding of “Categorialism” added to our premises regarding the ubiquity of easy knowledge of home truths and the impossibility of easy knowledge of heavyweight philosophical claims, I need make only one additional observation before setting forth my proof that CLR is false. That is that a demonstration utilizing only a couple of instances of *modus ponens* and one of *disjunctive syllogism*—could result in the easy acquisition of knowledge of the conclusion of that derivation.²⁹ And so to work:

²⁹ I do not here assume the truth of any ECLR principle. I merely take the view that, from such propositions as are known easily, we *might* reach an easily known conclusion via an argument as simple as mine: that is, the method would not alone necessarily generate difficulty. I believe that neither my constraints on knowledge nor any assessment I have made or will make here regarding the weight of this or that proposition should itself be considered heavyweight material. Though they are metaphilosophical—and so not Moorean—I take it that no assertion of mine on these matters here would,

- (i) If both Categorialism and CLR are correct, whenever we know such things as that we are sitting on a green chair, a competent reasoner can easily know that we are in the presence of an external physical object.³⁰
- (ii) We sometimes know such things as that we are sitting on a green chair.
- (iii) Therefore, if CLR is true then either Categorialism is false or we can sometimes be in a position that allows us to easily know such (deep philosophical) propositions as that there are external physical objects.
- (iv) But no one may ever be in a position that allows one to easily know such (deep philosophical) propositions as that there are external physical objects.
- (v) Therefore, if CLR is true, Categorialism is false and may be easily known to be so (simply by following steps (i)-(iv) above).
- (vi) But no one may ever be in a position that allows her to easily know any such deep philosophical propositions as that Categorialism is false.
- (vii) Therefore, CLR is false.

To make this argument intuitive, I have put it a bit loosely. It can be fleshed out by schematizing it as follows:

P = Categorialism is true.

Q = CLR is true.³¹

R = There is easy knowledge of some Moorean facts.

S = There can be easy knowledge of some heavyweight facts.

if true, put some ancient, apparently intractable metaphysical controversy to final rest. When, as in Roxanne's case, an apparently heavyweight proposition seems to follow quite simply from a few known lightweight facts, I take the appropriate response to be to conclude that the lightweight propositions and logical principles utilized cannot have actually been sufficient to prove such a deep conclusion. Something crucial must have been left out.

³⁰ As suggested above, on the Categorialist view, achieving such a "position" is no great accomplishment. Roughly, according to Thomasson's *Ordinary Objects*, *op. cit.* at p. 30, one need only be a competent reasoner/language user and understand the relevant terms.

³¹ Again, I encourage the reader to use whatever variant of CLR seems most satisfying.

T = This argument is sufficiently simple that its structure alone could not prevent any of its conclusions from being easily known by a competent, alert reasoner/language user who understands it.

U= There can be easy knowledge of the falsity of Categorialism.

1. $(P \bullet Q) \rightarrow (R \rightarrow S)$ (premise)
2. R (premise)
3. $\therefore Q \rightarrow (\sim P \vee S)$ (from 1 and 2)
4. $\sim S$ (premise)
5. $\therefore Q \rightarrow \sim P$ (from 3 and 4)
6. T (premise)³²
7. $\therefore Q \rightarrow U$ (from 5 and 6)
8. $U \rightarrow S$ (premise)
9. $\therefore \sim U$ (from 4 and 8)
10. $\therefore \sim Q$ (from 8 and 9)

³² As can be seen, T is not about the contents of any propositions used in this argument (the “what”), but solely relates to the complexity of the argument form (the “how”). That is because the argument requires that there be no *methodological* impediment to the easy acquisition of its conclusions. If any of those conclusions are nevertheless apparently heavyweight, either some premise must be as well, or, pursuant to our second constraint, such conclusion cannot actually be easily known. As explained above, besides simplicity of structure, the proof also relies on the definition of “Categorialism” (rather than any principle of ECLR) for the transmission of easiness. With respect to the “what”—the contents of each premise and conclusion—we must each make our own determination regarding whether it is the sort of item that is simply too heavy to be easily acquired, regardless of the simplicity of the derivation. I repeat that neither my constraints nor my weight assignments seem to me to be heavyweight material, but, naturally, readers will have their own opinions. I note that in the passage from “Elusive Knowledge” I quoted above, David Lewis seems to have agreed with me, since he claimed that the arguably metaphilosophical “we know a lot” is itself “a Moorean fact.” It may help in making evaluations of philosophical weight to consider whether the “proof” of the falsity of Categorialism from the material here utilized resembles something commonly associated with Roxanne’s “proof” of the reliability of her gas gauge and with Moore’s proof of the external world: the production of what I have called “cognitive dissonance.”

I have already defended both the truth and my weight assessments of the underived premises, but there is one possible objection I want to address straightaway. That is the claim that any such proof of the falsity of CLR is itself an example of the acquisition of easy knowledge of a heavyweight proposition. That seems clearly wrong to me. The claim that many of our homely general terms and predicates are shot-through with a Categorical axiom—where such an infusion would make it impossible for a competent reasoner to ascribe, for example, the property *being a chair* without knowing that all chairs must be external and physical—is a far-reaching cognitive theory, affecting nearly every aspect of philosophy, from ethics to philosophy of mind.³³ And, naturally, the denial of the truth of such a contention is also deep and weighty. CLR, on the other hand, while not a home truth, resembles my premises in also not being particularly deep—even if there is no consensus on its truth. Asserting the factuality of CLR is thus akin to asserting that there are no moose in Italy.³⁴ It is exotic rather than Moorean information, in the sense of possibly being useful to crossword puzzlers (or analytic philosophers), but it is not philosophically deep or central to any philosophical concern.³⁵ The controversy surrounding CLR has always been entirely derivative: a function of claims about how we should understand *knowledge* and what can be entailed by home truths.

I will conclude by noting the similarities and differences between what Luper has called “The Argument [against CLR] from Not Easily Knowable Propositions”³⁶ and what I am up to

³³ As noted above, it has even been suggested that the assumption of a somewhat similar (though this time necessary) psychological principle might allow one to construct a proof of the existence of an external physical plenum. See my “A New Proof for the Physical World,” *op. cit.*

³⁴ Or perhaps this is closer to CLR: If erasable bond is a type of stationery and a good stationery store carries all types of stationery, then a good stationery store carries erasable bond. Support for such a principle might engender argument, but it could hardly be considered heavyweight material. See my “Metaphysical Realism and the Various Cognitive Predicaments of Everett W. Hall, *op. cit.*, at pp. 269-284.

³⁵ The above argument shows that if CLR holds, we must deny Categoricalism. But it is important to note that the falsity of CLR does not imply either the truth or the falsity of Categoricalism.

³⁶ Steven Luper, “Epistemic Closure,” in Edward N. Zalta, ed., *The Stanford Encyclopedia of Philosophy* (Spring 2016 Edition), <http://plato.stanford.edu/archives/spr2016/entries/closure-epistemic/>

here. The earlier argument, which Luper attributes to Dretske and (and “possibly Nozick”), may be put quite simply:

NEKP: I know both that I am sitting in a green chair, and that that knowledge entails that I am not a BIV. If CLR is true, I must know that I am not a BIV. But that I am not a BIV is not the sort of thing one can know from only banalities like that I am sitting on a green chair and what such knowledge entails. Therefore, CLR is false.

The similarity is clear: both NEKP and my argument utilize the claim that it must be false that we can infer heavyweight truths from trivialities. But NEKP is question-begging. The skeptic will simply counter that no one *can* know she is sitting in a green chair without thereby knowing that she is not a BIV—so if we cannot easily know the latter, we cannot easily know the former. Thus, we have a clear *modus ponens/modus tollens* impasse. My argument is different: it avoids this stalemate by explicitly not deciding between Categorialism (a heavyweight doctrine which, in conjunction with CLR, *would* make too-easy metaphysical knowledge possible) and non-Categorialism, which, arguably, would leave the banality of home truths pleasantly independent of philosophical assertions, even if CLR were true.³⁷

The seminal Dretskean arguments against CLR have not universally convinced because, without assuming Categorialism, the inference from green chairs to “no-deception” utilizes arguably non-essentially involved concepts. That is, it seems clear that one may ascribe chairhood without having any concurrent representations regarding BIV simulations, so one arguably may assert the existence of chairs without simultaneously asserting the absence of all heavyweight deception. And CLR supporters have argued that in those instances in which the principle seems to come out false, it is only because the cogitator in question does not simultaneously know both p and $p \rightarrow q$. But according to the Categorialist, whether one realizes it or not, no one *can* assert that a chair is present without simultaneously representing and

³⁷ Obviously, if the argument goes through against traditional CLR, it can also be used to prove the falsity of any analogous principle that uses Crispin Wright’s “transmission of warrant.” See his “(Anti-) Sceptics Simple and Subtle: G. E. Moore and John McDowell,” *Philosophy and Phenomenological Research*, lxxv, 2 (September 2002): 330-348.

ascribing external physicality to the chair in all contexts. Achieving knowledge of the conditional is made trivial. Add Categorialism to CLR, and easy knowledge of hoary metaphysical doctrines becomes almost unavoidable.

Now, of course, the Categorialist may be completely wrong about all of this. Philosophy being difficult, it is not at all easy to be sure. But I have shown that if CLR were true, Categorialism could be demonstrated to be false, a derivation which would trivially resolve— instantly and for all time—numerous age-old disputes at the very core of philosophical discourse. Regrettably (to me, anyhow), the accomplishment of any such titanic feat is far beyond my abilities. We may therefore safely conclude that knowledge, given any sound conception, is not closed under known entailment.