

COMPARTMENTALIZED KNOWLEDGE

Forthcoming in *Philosophical Studies*

LEVI SPECTRE

The Open University of Israel

August 27, 2018

Abstract

This paper explores some consequences of David Lewis's (1996) understanding of how knowledge is compartmentalized. It argues, first, that he underestimates how badly it impacts his view. When knowledge is compartmentalized, it lacks at least one of two essential features of Lewis's account: (a) Elusiveness—familiar skeptical possibilities, when relevant, are incompatible with everyday knowledge. (b) Knowledge is a modality—when a thinker knows that p , there is no relevant possibility where p is false. Lewis proposes compartmentalized knowledge to keep treating knowledge as a modality while mitigating one of its unrealistic epistemological implications: In normal modal epistemic logic (and standard possible world semantics), a thinker always counts as knowing the strongest proposition that follows from the set of all the individual propositions that this thinker knows. Lewis's compartmentalization proposal is that thinkers merely know the conjunctions of propositions that are known in each of, but not across, their compartments. The irony is that in avoiding overblown knowledge the view now allows for thinkers to attend to skeptical error possibilities and yet knowledge is present. The account avoids inflation of knowledge in one sense only to acquire another type of knowledge the view denies subjects can have. This problem motivates an inspection of knowledge accounts whose intra-compartment closure principles are weaker than those that are valid in normal modal-logic. The conclusion is that some formulations of closure can avoid the challenge Lewis's view faces. Nevertheless, even these closure principles pose a barrier—perhaps an implausible barrier—for knowledge of ignorance. Even when the reasoning supporting the lack of knowledge is sound, a subject cannot always know she doesn't know. Interestingly, this obstacle is one that knowledge of knowledge doesn't seem to face.

* I presented several versions of this paper at Lund University; Stockholm University's 2016 international epistemology workshop; Israel's 2017 Philosophical Association conference; The Logic, Language, and Cognition Center at the Hebrew University; ECUP 9 at LMU Munich; University of Arizona; and University of Sheffield. I would like to thank everyone who attended. For discussions, comments, objections, and suggestions that significantly improved this paper I'm indebted to Stewart Cohen, Juan Comesaña, Julien Dutant, David Enoch, Paul Faulkner, Mikael Janvid, Moshe Halbertal, John Hawthorne, Maria Lasonen-Aarnio, Ofra Magidor, Ittay Nissan-Rozen, Erik J Olsson, Peter Pagin, Baron Reed, Daniel Rothschild, and Assaf Sharon.

I INTRODUCTION

To avoid internal conflict, to better utilize limited cognitive resources, or perhaps because of some evolutionary coincidence, our mental lives are to some extent compartmentalized. David Lewis (1996) appealed to this fact—specifically compartmentalized attention—to avoid an implausible consequence of his account: Just as in any normal modal logic, by knowing propositions a subject knows their conjunction. In addition to all the weaker propositions a subject counts as knowing, the strongest proposition that follows from the set of the known individual proposition—their conjunction—is something she automatically counts as knowing.

Lewis doesn't regard this consequence as it is often regarded when normal modal logic models are used, i.e., as an idealization, a simplification for modeling some specific epistemological property, as true for a limited number of propositions, or some such.¹ As his account unavoidably depends on normal everyday knowledge being a modality—roughly, that it depends on the relevance of what is possible for the speaker—Lewis cannot set this consequence aside. The motivation for this view, though, depends on what people ordinarily count as knowing. This, at the very least, does not include all the logical consequences of everything people know. The solution he proposes is Compartmentalized Knowledge (*Comp-K* for short). A thinker can know that p and that q and fail to know their conjunction ($p \wedge q$). A thinker's knowledge is closed only within and not across compartments, Lewis proposes. Two thinkers, then, can know the same individual propositions and yet because of the way they are compartmentalized, one of them knows their conjunction while the other does not. The latter's knowledge is such that at least one conjunct isn't in the same knowledge-compartment as the other

¹ For instance, Timothy Williamson (2013) models Gettier cases in (normal) modal logic and regards the above strong closure of known propositions as a “harmlessly simplifying idealization of the models.” (2013, 4). In Sharon and Spectre (2013b) we argue that normal modal logic cannot adequately capture Gettier cases within some of the internalist views that validate Gettier's assumptions. A salient problem, for instance, is suggested by Gettier's (1963, 121) condition that a justified belief can be false. Since Williamson's model has a non-reflexive doxastic accessibility relation, a justified belief can indeed be false. However, it will also be the case that justification is will be closed in the strong sense discussed in the main text. Though these two are not contradictory, it is difficult to reconcile the strong multi-premise justification closure with a view that accepts that false propositions are sometimes justified. Williamson (2015) resolves this issue by modeling Gettier cases with neighborhood rather than the standard possible world semantics. Multi-premise closure isn't valid in these models, let alone the normal modal logic type of closure from the main text that doesn't differentiate between known premises and known conclusions.

are. So we can say that in Lewis's epistemology a thinker's knowledge isn't closed in the strong modal logic sense, and yet, it is a modality. It still depends on the relevance of possibilities.

This paper argues that as ingenious as this scheme is, it cannot work. If knowledge is compartmentalized, it avoids the implication of normal modal logic only at a price that renders the account—by its own lights—unworkable. Comp-K conflicts with the most central aspect of his epistemology: Elusiveness. This is the idea that attention to (uneliminated) possibilities of error—especially possibilities familiar from skeptical arguments—makes knowledge that is otherwise present, evaporate.² Elusiveness is crucial for the balance Lewis aims at striking between the deathtrap of skepticism and the awkwardness of fallibilism (i.e., that statements of the form “He knows, yet he has not eliminated all possibilities of error.” sound wrong). Comp-K cannot achieve this balance—it isn't Elusive—because knowledge can be present even while attending to skeptical possibilities of error. Moreover, it has more bizarre consequences than those it is meant to avoid.

Since compartmentalized attention is a fact, not a theoretical proposal, the option of forgoing it and returning to Lewis's original compartmentalization-free view is off the table.³ So should the problems of Lewis's view concern other attention (or salience) dependent knowledge accounts? The paper argues that *some* realistic intra-compartment closure principles resolve Lewis's problem. These views will, however, need to get accustomed to compartmentalized subjects' knowledge being open (across compartments, that is). Moreover, on some plausible assumptions regarding the nature of Comp-K, the paper presents a problem for knowledge of ignorance (or lack of knowledge) even for the views proposing the weaker closure principles. It goes on, though, to resolve the puzzle by arguing that Comp-K imposes a special structural difficulty on knowledge of ignorance that knowledge of knowledge doesn't have to face.

The plan of the paper is as follows. §2 briefly comments on where the idea of com-

² More carefully, it is not that knowledge “goes away” or “evaporates.” Rather, in those contexts, “know” expresses a relation between a subject and a proposition that almost never holds.

³ To see the difference I am speaking to here, think the role of attention in Lewis's view. It is possible to rethink the idea that what a knowledge ascriber is thinking of makes a difference to whether the attribution is true. Attention—the focus of awareness, may or may not play this theoretical role. It isn't up to the theoretician to say whether this attention is compartmentalized.

partmentalization comes from and about the (relatively sparse) philosophical literature. §3 is a full view of the essentials of Lewis’s knowledge account and two subsections spell out two central properties that it has: Elusiveness and Modal Closure. Lewis’s idea of how to incorporate Comp-K into his account and what problem he takes it to resolve is the subject of §4. It also presents a preface paradox type problem due to Hawthorne (2002) and comments on how this problem might not arise in Lewis’s Comp-K view. §5 is the problem for Elusiveness within Lewis’s Comp-K account. A possible response that indexes knowledge to compartments is considered and rejected in §6. §7 offers a more ready and plausible answer to the problem (that is incompatible with Lewis’s view). It requires rejecting the entire modal framework in favor of certain weaker operation dependent closure principles. In §8 another problem is considered even for the alternative view and a tentative reply is offered. It requires viewing knowledge of ignorance as somewhat—in a sense to be explained—hard to come by. It is there argued that this obstacle to knowledge of ignorance can very well be a structural issue.

2 COMPARTMENTALIZATION IN EPISTEMOLOGY

Compartmentalization is usually understood as a psychological mechanism that maintains (and may even improve on) normal mental functioning by keeping different types of tension from surfacing to awareness. Wikipedia defines compartmentalization as follows:

A subconscious psychological defense mechanism used to avoid cognitive dissonance, or the mental discomfort and anxiety caused by a person’s having conflicting values, cognitions, emotions, beliefs, etc. within themselves.

(Wikipedia: “[compartmentalization-psychological](#)”)

What is important in this description—assuming it is roughly true—is that compartmentalization allows conflicting mental states and processes to co-exist by inhibiting awareness of their coexistence, preventing direct interaction between them. Also important for present purposes is that there is some mechanism of separation between compartments. We can think of the compartments as mental states and activities belonging to two different people and compartmentalization as a whole as a way to hide the

existence of one from the other. To some extent, that is, sometimes the separation is only minimal as the cost of having them meet is low.

Compartmentalization in the epistemological literature mostly concerns doxastic (binary or degree) states that are in some kind of tension. Robert Stalnaker (1984) and David Lewis (1982)⁴, for instance, advance compartmentalized belief views (in an attempt) to dissolve problems of the logic of binary beliefs. Andy Egan (2008) appeals to compartmentalization as a way to resolve a problem he identifies between higher and lower level degrees of belief.

Here I will focus on a different type of compartmentalization that has rarely⁵ been explored—compartmentalized knowledge. This is simply the thesis that—roughly in the psychological sense characterized above—thinkers’ knowledge is fragmented. My point of departure will be Lewis’s (1996) knowledge account where due to the central role attention is given, compartmentalized knowledge plays a surprising role.

3 LEWIS’S EPISTEMOLOGY

The distinctive theoretical element of Lewis’s (1996) epistemology is his notion of Error Possibility (*EP*, for short). It plays the central role in his semantics of the verb “know” in an attempt to resolve the major paradoxes and puzzles of epistemology. The idea is that the word “know” expresses different relations in different contexts in which it is ascribed, allowing “all EPs” to refer to different domains. This in turn allows Lewis to address the main puzzle—the one he thinks is the key to resolving the others—of how to reconcile the following seemingly conflicting epistemological claims: Infallibilism—to know that *p* a subject must have evidence that eliminates all of its EPs, and yet, uneliminated EPs exist “everywhere”. Together these theses seem to entail skepticism—we have little if any knowledge if our evidence needs to eliminate any uneliminated EPs (let alone

⁴ Lewis (1996, 556) explicitly rejects the claim that knowledge entails belief. So belief compartmentalization doesn’t necessarily translate into Comp-K. There are though, questions about how the two papers—Lewis (1982) and (1996)—can be combined. For instance, how do beliefs that are part of an inconstancy fit into the Comp-K picture?

One consequence of belief compartmentalization that deserves mention, is that on Lewis’s view believing $\neg p$ doesn’t entail (even if knowledge does entail belief) that a subject doesn’t know that *p*. However, regardless of compartmentalization, it doesn’t seem like Lewis’s view rules this possibility out in the first place. See Lewis (1996, 556).

⁵ There are exceptions. Norby (2014) is a recent skeptic about belief compartmentalization. Greco (2015) discusses knowledge as well as belief.

all of them). However, it is a Mooreian fact that we have lots of everyday knowledge, Lewis says. So “[w]e are caught between the rock of fallibilism and the whirlpool of scepticism.” (Lewis 1996, 550).

The solution isn’t to deny the Mooreian fact or to find fault with his first desiderata (and become accustomed to fallibilism). Lewis’s major claim is that the way the truth value of sentences in which “know” figures is determined allows us to keep them both. “Know” is a three place relation between a subject, a proposition, and a set of EPs determined by the context of its use. In accordance with infallibilism, a speaker’s sentence of the form “S knows that *p*” will be true in a particular instance of its use just in case *no member* of the set of EPs is true in a world where S’s evidence is as it is in the actual one (or as in the world the knowledge ascription is about). In accordance with the Mooreian fact, the relation expressed by “know” often enough picks out sets of EPs that can all be eliminated. The sentence is false otherwise. His definition, then, is:

S knows that *P* iff *S*’s evidence eliminates every possibility in which not-*P*
– Psst! – except for those possibilities that we [the speakers] are properly
ignoring.

(Lewis 1996, 554)

So far we have seen the structural features of Lewis’s epistemology. But why should we believe that the expressed relation holds often enough? The content of his epistemology—the way by which he wants to avoid fallibilism on the one side, and skepticism on the other—is given by four rules. These rules determine which EPs are *not* ignored legitimately and hence are included in the set of EPs to be eliminated by the evidence (if the relation holds, that is): (1) *Actuality*: the actual world can never be properly ignored. (The actual world of the subject, not the world the speaker thinks the subject is in.) (2) *Belief*: “A possibility may not be properly ignored if the subject gives it, or ought to give it, a degree of belief that is sufficiently high, and high not just because the possibility in question is unspecific.” (Lewis 1996, 555). Note that this rule concerns the subject, not the ascriber. (3) *Resemblance*: “Suppose one possibility saliently resembles another. Then if one of them may not be properly ignored, neither may the other.” (Lewis 1996, 556). Finally we have the rule that is important here, the rule of *Attention* (4):

When we say that a possibility is properly ignored, we mean exactly that; we do not mean that it could have been properly ignored.

(Lewis 1996, 559)

This means that if a p -EP is being attended to by us (the ascribers), the subject's evidence must eliminate it if the "S knows that p " can be true. Otherwise the Error Possibility in question is, in modal logic terminology, *accessible*, i.e., the sentence is false. It doesn't matter how farfetched this EP is or how many other EPs are ignored or eliminated.

One last ingredient of the theory—how EPs are individuated: "A possibility will be specific enough if it cannot be split into subcases in such a way that anything we have said about possibilities, or anything we are going to say before we are done, applies to some subcases and not to others." (Lewis 1996, 552).

These, then, are the essential workings of Lewis's epistemology.⁶ I turn now to two of its features.

3.1 Elusiveness

How does Lewis avoid the whirlpool of skepticism? The answer is by now familiar but it will be worthwhile to see how central this idea is to his project. The solution is given by the property that the Lewis's paper is named after. I will refer to it as *Knowledge Elusiveness* (or *Elusivness*, for short). What we properly ignore in one setting may not be properly ignored (or not ignored at all) in another. Keeping the subject, the proposition, the time, and the sentence fixed, what is the difference between contexts of knowledge attribution? The subject's evidence and hence the EPs that can be eliminated, the beliefs of the subject (including what the subject ought to believe) regarding EPs, the world of the subject, all remain constant across the contexts. What changes is the salience of resemblances between EPs, but more often (and relatedly), what is entirely different across contexts is the EPs that speakers pay attention to. And this is where Lewis believes the truth in skepticism resides. When skeptics get us to think about skeptical EPs our knowledge attributions to ourselves and to others will be false. The expressed relation in "S knows that q " is too hard to satisfy, specifically when q is itself a skeptical hypothesis:

⁶ Lewis also has "permissive rules" that tell us when we may properly ignore an EP. But if we may ignore an EP but are not, that would not matter to whether the subject knows that p . *Attention* says S will not know that p if we are attending to an EP that isn't eliminated by her evidence regardless of whether we could have properly ignored it (legitimately). See quote in the main text.

Do some epistemology. Let your fantasies rip. Find uneliminated possibilities of error everywhere. Now that you are attending to them, just as I told you to, you are no longer ignoring them, properly or otherwise. So you have landed in a context with an enormously rich domain of potential counter-examples to ascriptions of knowledge. In such an extraordinary context, with such a rich domain, it never can happen (well, hardly ever) that an ascription of knowledge is true. Not an ascription of knowledge to yourself ... and not an ascription of knowledge to others. ... That is how epistemology destroys knowledge ... That is how knowledge is elusive. Examine it, and straight-way it vanishes.

(Lewis 1996, 559-60)

To focus the argument here, it is good to have some claim about what the Elusiveness thesis is: In a context in which an uneliminated q -EP is attended to, “S knows that q ” is false.^{7,8} I will take a standard skeptical EP as a familiar example of Elusiveness but many others would do just as well. If the following is false, so is Elusiveness:

(BIV-Elusiveness) If a subject is attending to the possibility that she is a brain in a vat, the sentence “I know I’m not a brain in a vat” is false.⁹

The idea is that skeptics win the battles in which they participate because all they need to do is raise EPs to awareness. In most battles, however, they don’t participate. That is how everyday knowledge—the Mooreian fact that we know better than any philosophical premise—is preserved. Elusiveness, through shifts of awareness, is the mechanism by

⁷ Another way of stating this is by defining a (minimally reflexive) accessibility relation R between worlds in the set of worlds W . For this purpose, propositions can be coarse-grained sets of worlds in W . Knowledge that p in the $\langle W, R \rangle$ frame can be $K(p) = \{w : wRx \models x \in p\}$. If $p \vdash \neg q$ then given the *Attention* rule, if S is attending to q then either S ’s evidence eliminates it, or there is at least one q -world x' such that wRx' . In which case since $(wRx' \wedge x' \notin p)$, $\neg K(p)$ in w . For more formal and detailed 5-tuple frame definition of Lewis’s accessibility relation, see Salow (2016, 1576) and references there to rival Lewisian knowledge definitions.

⁸ This isn’t quite right because attending doesn’t always make an uneliminated EP accessible (in the sense of the previous note). But the cases are odd or at least extraordinary. Moreover, it’s hard to see how they could work in first person knowledge attributions. See (Lewis 1996, 556) about the possibly a well trained dog shot a gun and hence the court doesn’t know it’s was the suspect who did it. This amusing example is presented in relation to the Belief rule, I don’t know how if it can sidestep the rule of Attention.

⁹ Since it is almost always the case that the subject/speaker is attending to what she is saying, we could also use the claim: “I know I’m not a brain in a vat.” is false whenever uttered.

which he accommodates the multitude of uneliminated EPs and knowledge without rejecting infallibilism.¹⁰

3.2 Modal Closure

There are some properties that Lewis’s epistemology exhibits that come as a byproduct of the framework in which he presents it. A notable problematical byproduct is known as the (Logical) *Omniscience problem*. Other properties are more essential to his view. The property I am calling *Modal Closure*, both follows from the framework—from propositions being sets of possible worlds—but also from treating knowledge as a type of local necessity, as a modality concerning all the EPs. That is, the relation will hold or not dependent on the accessibility or inaccessibility of sets of EPs (depending on the expressed relation). To see this consider whether by knowing that p and knowing that q , a subject will know that $p \wedge q$. By Lewis’s definition, knowing p means that every p -EP has either properly been ignored or eliminated by the subject’s evidence. Likewise regarding q ’s EPs.¹¹ Hence, there are no $p \wedge q$ EPs that are not properly ignored in the context of the ascription or not eliminated by the subject’s evidence. So she knows $p \wedge q$ if she knows that p and knows that q . Likewise with the weaker principle of strict equivalence. These principles follow from Lewis’s definition regardless of whether they are presented within a normal modal logic framework. They are, we might say, a byproduct of his brand of infallibilism, not the framework (but see below).

Unlike the Omniscience problem—the problem that a subject knows all logical and mathematical truths—Modal Closure, then, is essential to Lewis’s view.¹² He is well

¹⁰ Lewis uses the rule of resemblance in his attempt to resolve the “static” problems, why there isn’t knowledge in the Gettier cases, the fake barn case, the available but not yet perceived evidence case, and lottery cases. There is room, though, for shifts that are due to salience changes. See Hawthorne (2002) for a Lewisian solution to the preface paradox that hinges on salience in accordance with the rule of Resemblance. The point is, skepticism is mostly a matter of attention shifting and the solutions to other puzzles and paradoxes mostly employ other prohibitive rules.

¹¹ Note that the rule (2), *Belief*, cannot introduce a new EP because Lewis explicitly rules out any case that accumulates a high enough degree of belief by being less specific. See his individuation of EP criteria in (1996, 552) and Hawthorne (2002, 243).

¹² Unfortunately things are not as neatly separable as my comments so far suggest. There are mixed cases that involve both properties. Cases, for instance, that are part deductive and part modal framework driven, for instance that a priori equivalences are indistinguishable regardless of whether the subject knows of these equivalences. Though we often know well enough what to treat seriously and what to hang on the framework, this isn’t always so. Not too much depends on this, but it looks like Lewis regards the Omniscience Problem as something he can think of in terms of idealization of rationality. See quotes below where he says that he cannot treat the Modal Closure issue in this way.

aware of this:

If we analyse knowledge as a modality, as we have done, we cannot escape the conclusion that knowledge is closed under (strict) implication.

(Lewis 1996, 563)

No knowledge of the implication itself is needed. As before, this follows both from the possible worlds framework but also from—as Lewis says—his analysis of knowledge. Like with Elusiveness, it is useful to have a thesis that follows from the view to which it will be easy to refer. The following is one that will soon be important.

(Modal-CL) A subject *S* knows the conjunction of all the (individual) propositions that *S* knows. (More generally: For all *S*, any accessible EP of a premise of a deductively valid argument is also an EP of its conclusion. If all are eliminated from the premises by *S*'s evidence, *S* knows the conclusion.¹³)

To see how problematical this is, consider John Hawthorne's challenge to Lewis's view (2002):

I write down a manuscript consisting of a very long sequence of sentences, $s^1 - s^n$, sincerely saying to myself at each juncture, 'I know that what I'm writing down is true.' Let us label these silent soliloquies $k^1 - k^n$. Let us presume that each sentence has, individually, high epistemic credentials, and that, as it turns out, each sentence is true. Let us suppose, indeed, that by ordinary standards each individual sentence would be reckoned to express a piece of knowledge. . . . I say, on inductive grounds, 'It is very likely that I will make (will have just made) a mistake. After all, every manuscript that any of my highly intelligent friends has ever written has contained a mistake.' Once again, by ordinary standards, this would appear to express a piece of knowledge. Even if, as chance would have it, the sequence s^1 to s^n is mistake free, the claim about likelihood would appear to be both true and known. Assuming closure, it seems that I both know that a long conjunction is true and also that it is likely to be false.

¹³ Lewis notes that this result also follows in any normal modal logic in footnote 21 of (Lewis 1996).

(Hawthorne 2002, 249-50)

Hawthorne's conclusion is bad enough, but the situation for Lewis's view is even worse if you assume that *p is likely* is a weakening of *p* (so that if you know the later you know the former). On Modal-CL it would then follow that "I know that the conjunction *q* is both highly likely and also unlikely." What went wrong?

The problem seems to be that the EPs for each of the individual propositions has—by *Belief* (Lewis's rule of)—a low credence, low enough for the rule not to render any of them *illegitimately* ignored. So by Modal-CL, the conjunction is known. However, the number of EPs over the entire conjunction is vast and various. We know this because we know that it's very hard to write an error free book (this is expressed in the argument the inductive knowledge of the low likelihood of an error free book). Hawthorne proposes to change *Belief* (2002, 245), but as we will see in the next section, this amendment may not be necessary.

4 COMPARTMENTED MODAL CLOSURE

Hawthorne has shown that Lewis's view has potential difficulties with differences between propositions and their conjunction, differences that the view doesn't seem capable of capturing. The framework doesn't even seem to have the resources to differentiate between them. Ascribing knowledge of conjuncts is ascribing knowledge of their conjunction. The relation "know" can only hold for both or none. So it seems that there is no way to give up Modal-CL without giving up the whole framework.

In fact, however, Lewis rejects Modal-CL without giving up—or so it first seems—his view. Let's first be positive: Without Modal-CL, Hawthorne's problem doesn't necessarily arise. Here is Lewis's idea:

Suppose two or more premises jointly imply a conclusion. Might not someone who is compartmentalized in his thinking – as we all are? – know each of the premises but fail to bring them together in a single compartment? Then might he not fail to know the conclusion? Yes; and I would not like to plead idealization-of-rationality as an excuse for ignoring such cases. But I suggest that we might take not the whole compartmentalized thinker, but rather

each of his several overlapping compartments, as our ‘subjects’. That would be the obvious remedy if his compartmentalization amounted to a case of multiple personality disorder; but maybe it is right for milder cases as well.

(Lewis 1996, 565)

Lewis, then, is suggesting a replacement of Modal-CL with Compartmented Modal Closure (*CompModal-CL* for short):

CompModal-CL: A subject *S* is guaranteed to know any proposition equivalent to, or logically weaker than, the conjunction of all the propositions that *S* knows *within any one* of *S*’s compartments.

Compare, then, *S* and *S*’. They may both know all of s^1 to s^n , but only *S*’, say, knows their conjunction. Or maybe it’s impossible to know their conjunction because this number of propositions can’t all be held in one knowledge compartment. This is not quite enough to claim that CompModal-CL is a solution to Hawthorne’s challenge, but it’s a start. To have a full resolution there would have to be some substantive connection between the likelihood of writing a mistake free book that has n sentences and the possibility of all of the propositions that s^1 to s^n express in one knowledge compartment. Perhaps it is possible, but only if the subject has high enough support for each of the n propositions so that the book does have high enough likelihood of being mistake-free (or enough so that it is false that one could learn by induction that the book is probably mistake free).¹⁴

It isn’t clear—at least not to me—if Lewis has a deep reason for replacing Modal-CL with CompModal-CL. Perhaps he realizes that knowing the conjunction of everything one knows isn’t plausible on the face of things, perhaps he thought Modal-CL would lead to the kind of preface problems Hawthorne points out, or maybe he just saw that compartmentalized attention—which we all have to some extent—inevitably leads to Modal-CL’s failure. Whatever his reasons may be, I will argue that CompModal-CL conflicts with Elusiveness (and in particular BIV-Elusiveness) that we saw is at the center of his account. At this point it should be noted, though, that preserving the modal nature of his account—the knowledge relation to EPs that are all eliminated—without endorsing

¹⁴ There are psychological facts about compartmentalization that need to be explored before we can be satisfied with an answer to the preface cases along the lines suggested in the main text. Or so it seems to me, at least.

the full implications of Modal-CL is ingenious. It is unfortunate, then, that it cannot work. This may lead some to reevaluate Elusiveness (but see §8).

For later use it is worth presenting the counter-case to Modal-CL that Lewis proposes in the following way:

Table 1: Compartmented Modal Closure

| <u>Comp. 1</u> | <u>Comp. 2</u> |
|--------------------------|----------------|
| $K(p)$ | $K(q)$ |
| <u>Subject Knowledge</u> | |
| $K(p)$ | |
| $K(q)$ | |
| $\neg K(p \wedge q)$ | |

5 COMPARTMENTALIZED ELUSIVENESS KNOWLEDGE?

Lewis notices that compartmentalization has implications further than rendering Modal-CL invalid.

A compartmentalized thinker who indulges in epistemology can destroy his knowledge, yet retain it as well. Imagine two epistemologists on a bushwalk. As they walk, they talk. They mention all manner of far-fetched possibilities of error. By attending to these normally ignored possibilities they destroy the knowledge they normally possess. Yet all the while they know where they are and where they are going! How so? The compartment in charge of philosophical talk attends to far-fetched possibilities of error. The compartment in charge of navigation does not. One compartment loses its knowledge, the other retains its knowledge. And what does the entire compartmentalized thinker know? Not an altogether felicitous question. But if we need an answer, I suppose the best thing to say is that S knows that P iff any one of S 's compartments knows that P . Then we can say what we

would offhand want to say: yes, our philosophical bushwalkers still know their whereabouts.

(Lewis 1996, 565)

One thing to note is that when Lewis asks what the compartmentalized thinker knows, he somewhat reluctantly and casually suggests a criterion that I'll label Compartmentalized Subject Knowledge (or *Subject-K* for short) connecting compartment to subject level knowledge:

(Subject-K): S knows that p iff any one of S's compartments knows that p .

For Lewis, despite his casual phrasing, Subject-K is at this point already unquestionable. His point in replacing Modal-CL with CompModal-CL was based on the claim that a subject knows p and knows that q while not knowing their conjunction. He is already committed to the idea that compartments collect to the subject level (Subject-K). Otherwise, it's hard to see how we would have a counterexample to Modal-CL with only two compartments. In other words, \neg Modal-CL by way of CompModal-CL implies Subject-K.¹⁵

The major thing to note, however, is that CompModal-CL (and if we want to add it as another premise Subject-K) is a rich source of BIV-Elusiveness counterexamples. The epistemological bushwalker's self-ascribed claim that they don't know that they are not BIVs—which should be true according to BIV-Elusiveness—is false. By attending to the far-fetched possibility that they are BIVs, they do destroy the knowledge that they have in their philosophical compartment. However, they know all along where they are and where they are going, because this knowledge is retained in another compartment that doesn't attend to the possibility that they are BIVs. (Different partially overlapping compartments are, we have said, like different subjects with different sets of EPs.) Now if their bushwalking compartment is modally closed, i.e., if CompModal-CL holds, they know in that compartment that they are not-BIVs. Vat bound brains can't walk anywhere

¹⁵ CompModal-CL doesn't itself entail Subject-K, it only does so assuming that it can provide a counterexample to Modal-CL with only two compartments. To see this suppose S has three compartments A, B and C. p could be known in A and B, and q known in B and C. The criterion for subject knowledge could be: S knows that p (simpliciter), iff p is known in the majority of S's compartments. We would then have the same counterexample to Modal-CL that Lewis gives but with no commitment to Subject-K. Nevertheless, the problems I will discuss below will also hold with this and similar compartment-to-subject-knowledge criteria. (But see the section on indexing §6.)

(at least not in the sense that the philosopher-bushwalkers are walking) and moreover they would not know where they are. But then this knowledge collects to the subject level in accordance with CompModal-CL (or at least Subject-K). Hence, as subjects they do know that they are not BIVs. So, their claim that they don't know they are not BIVs is false after all contra BIV-Elusiveness .

Where as before K is the knowledge operator, b is the proposition *I am bushwalking*, v refers to the proposition that they (each referring to oneself) are legless BIVs and l is *I have legs* (and assuming that $b \vdash l$ and $l \vdash \neg v$) the diagram may help clarify their epistemic state (the argument, though, has already been stated above):

Table 2: Modaly-Closed Compartmentalized Knowledge

| Epistemological Comp. | Bushwalking Comp. |
|-----------------------|-------------------|
| $\neg K(\neg v)$ | $K(b)$ |
| $\neg K(l)$ | $K(l)$ |
| $\neg K(b)$ | $K(\neg v)$ |
| Subject Knowledge | |
| $K(b)$ | |
| $K(l)$ | |
| $K(\neg v)$ | |

The lesson is that a subject doesn't know that p (at least in the self-ascribing situation)¹⁶ only if she doesn't know in some compartment propositions that entail that p . A sentence like "I don't know that I'm not a BIV" will in many cases be false. How extensive this phenomenon is, is hard to say. It depends on how frequently, and in what way, we are compartmentalized. Nevertheless, it seems clear that just by attending to far-fetched

¹⁶ It isn't easy to apply the case Lewis presents to third person knowledge ascriptions. Here is one stab at how it is supposed to work: Presumably, the compartmentalization that is relevant to the determination of the expressed relation is the attributer's (and maybe also the third party listeners). What seems the most pertinent here is the operative compartment—the compartment from which the sentence is uttered (or more generally the compartment that guides the organisms behavior). Since this will determine the truth value of "S knows that p ", the way the subject is compartmentalized is immaterial (leaving to one side the rule of Belief). She doesn't have the kind of relation that the philosophical compartment expresses. However, this should also be the case for everyday bushwalking-like compartments. And if so, it isn't clear how supposing "S knows that p " and "S knows that q " are both true, "S doesn't know that $p \wedge q$ " can be true. It seems—in line with the argument in the main text—that Lewis can have CompModal-CL or Elusiveness but not both.

possibilities, at least in the self-ascribing case, we do not destroy our knowledge. Other conditions must hold. Elusiveness is false and so we are back to Lewis's original problem of how to avoid both Scylla and Charybdis. It seems that the view has landed on the side of anti-skeptical *fallibilism*: "I know I'm not a brain in a vat though I haven't eliminated this possibility" is true whenever I retain some everyday logically related knowledge in some compartment.¹⁷

Let us look at another example. Suppose I am thinking of farfetched EPs but I have knowledge in some other compartment that later I need to pick my kinds up from school, which entails, say, driving and holding on to the steering-wheel. In this case, I know I have hands, I have kids, a car etc. But I do not know things that do not follow a priori from my knowledge in that compartment. I don't know who won the last presidential election, that Paris is in France and so forth. Unless, that is, those things too are known in some further compartment. Either way, this clearly can't be right, or at least, knowledge is not Elusive in Lewis's sense. I might have landed in a situation with a very rich domain of EPs, but it isn't true that any self ascribed knowledge will be false as Lewis claims (see above).

6 INDEXING

I have argued that Lewis's attempt to weaken his epistemic modal-closure (Modal-CL) view and replace it with modally closed compartments (CompModal-CL), fails (assuming the elusiveness is non-negotiable). Technically, there is a way out: Replace Subject-K—the principle that collects knowledge in any compartment to the subject level—with a principle that says a subject knows what is known in *all* of S's compartments. (Other principles that collect knowledge to the subject level, e.g., if *p* is known in most compartments then S knows that *p*, are prone to a similar, albeit more complicated,

¹⁷ John Hawthorne has suggested to me a way to avoid this problem. Suppose our mental life is similar to a Hydra with many heads each calling attention (in a non-technical sense) to itself. Imagine further that somehow, "I" manages to pick one head/compartment. If this were the case, it seems Lewis could claim that "I know I'm a brain in a vat" is false when referring to the epistemological compartment. So if the referring method picks out the right compartment, BIV-Elusiveness could still be true. We agree, however, that on this suggestion there is no counter example to Modal-Closure. We also agree that third person attributions are hard to make sense of here.

problem to the one we've seen.)¹⁸ The best version of this is, I think, an indexing view. The idea is to index knowledge to compartments and attribute unindexed subject level knowledge when all compartments know.

There is evidence that we do something similar to indexing when we say things like “as a scientist, I believe. . .”, “I believe, as an observant Jew that. . .” or “as a judge, I think. . .” so we do sometimes relativize beliefs to roles that are similar to compartments.^{19,20} However, when it comes to knowledge, these “as a” claims are mostly made in order to emphasize that the subject has access to some special evidence or expertise. The “as a” belief statements themselves (on the assumption that knowledge entails belief) don't seem very promising for the purposes of indexing knowledge. They are sometimes meant to qualify beliefs or maybe explain away apparent inconsistencies. For instance, “I believe that God exists as a religious person but not as a scientist”, or “I believe in impartiality as a judge, not as a parent.”²¹ This type of indexing or relativizing isn't common when knowledge is concerned . But even if I am wrong about this, it isn't central or robust enough for the kind of theory Lewis is proposing. A theory, that is, that is designed to account for all knowledge attributions.²²

¹⁸ David Enoch suggested that there may be other principles that determine the relation between compartment and subject knowledge. For instance, a subject knows whatever is know in a relevant compartment. Like with the Hydra case of the previous footnote, I don't deny that this is possible. I just don't see—at least not without complete *ad hocery* what could motivate CompModal-CL without losing Elusiveness. In this case, Modal-CL is retained, it seems. I don't want to deny the possibility of some combination of CompModal-CL with Elusiveness. Nor can I. What I do want to emphasize is that the epistemology as it stands makes incompatible commitments and that the shape it can take by rejecting one (or more) of them depends heavily on the principle governing the relation between compartment and subject knowledge. It seems to me that the psychology of attention and compartmentalization can contribute here.

¹⁹ But there is also a dissimilarity here because “as a” beliefs are also expressions of awareness of other (at least potential) roles. This is not typical of (and not always possible in) compartmentalization cases.

²⁰ Moshe Halbertal has proposed the following type of case: A Captan of sinking ship is also the father of two of the young passengers. As a farther he will do anything to save his kids while as a Captan he must be impartial. We are not concerned with this kind of compartmentalization (assuming that it is). But it is useful to see that this kind of devision need not be unconscious.

²¹ In several contexts “belief that” combines well with “as a” but so would “belief in.” E.g. “As a reporter I believe in free speech, but not as a conservative.” In other contexts “as a” combines rather badly with “believe that.” In reference to the past (assuming there isn't some backstory) “as a X I believe that” sound bad and is better when it comes to the future: “As a reporter I believe that it will rain tomorrow, but not as a reader.” “As a reporter I believe that it rained yesterday...” This might be because it is harder to imagine how the role is relevant. But it is not clear to me that this is the explanation.

²²I do not want to deny that there are other expression types that mark compartments (or index to compartments). I just don't know of any. Is it possible that the indexing is subconscious? Here too, psychological data may be useful.

There is a more direction problem here for Lewis. Even if some kind of indexing Lewisian theory is possible, it won't be able to do what Lewis sets out to do by replacing Modal-CL with CompModal-CL. A subject would not know p and knows q but not their conjunction. Rather, the compartments would know but not the subject. It would be similar to claiming that two different subjects know different things. Or if the subject does know each proposition, the subject also knows their conjunction. Nevertheless, I don't think I can close the door on this type of view (nor do I want to). I'll leave it for further development. The main issue is how to connect the indexed knowledge to subject knowledge and to third person knowledge attributions.

7 CLOSURE

If attention is compartmentalized in a minimal way we have (with Lewis) assumed, how do those who are not committed to Lewis's theory—particularly to treating knowledge as a modality—account for Comp-K? This is an important question because at least attention compartmentalization, as noted earlier, is not a theoretical hypothesis. If it is anything, it is a psychological fact. The remainder of this paper will detail some of the implications of compartmentalization. I start with ways of avoiding the situation Lewis's view finds itself in.

One obvious idea is that since the problem for Lewis's view—the one he appealed to Comp-K to resolve—is his overly strong closure principle, weakening will resolve any inconsistencies or tensions with Comp-K. I will argue that we need to proceed with care here, but first I want to say something about the epistemological views that are vulnerable to problems concerning closure and compartmentalization.

Two views (or frameworks) I will focus on are contextualism and Impurism (i.e. Subject Sensitive Invariantism). The reason for this is that both views often incorporate shifts from knowledge to ignorance as a function of attention or salience. Compartmentalized subjects might know in one compartment something that entails knowledge that the view says they can't have because of what they are attending to in another.²³ So

²³ Let me say again that Comp-K can be true even if attention (or salience) has no impact on what is known or unknown. Comp-K can result from belief compartmentalization together with the view that knowledge entails belief. For instance, assuming I believe that p in compartment A and $\neg q$ in compartment B while $p \vdash q$. I might then know that p in A but not in B. So though I cannot develop these ideas here,

a potential problem with Comp-K is that the wrong closure principle might return the same problem Lewis’s view faces. Again, this problem can arise at least for views that assume knowledge is sometimes attention dependent.

To see the problem more clearly suppose that a subject, due to high stakes, say, doesn’t know that q in Compartment-1. And assume the following Closed Compartmentalized Knowledge principle:

(ClosedComp) Necessarily, if S knows that p entails q and in that compartment S knows p , then S knows that q in that compartment.

In this case, though compartment-1 doesn’t know q , the subject might. This will be the case if p and $p \vdash q$ are known in compartment-2. This subject, then, does count as knowing that q due to his knowledge in compartment-2.²⁴ This structure will run against cases that Impurists and contextualists have presented in support of these views.

Suppose, e.g., that a subject knows he won’t be able to afford a Safari trip (Hawthorne 2003). If he knows the relevant background facts, he knows that if the ticket in his pocket to the state lottery will win he will be able to afford going on Safari. Now it seems, though, that we will have difficulty with saying that this subject doesn’t know that his ticket won’t win the lottery. Why? The structure is the following:

Table 3: ClosedComp

compartmentalization might have very wide epistemological implications.

²⁴ Jason Stanley’s view (2005, 18-9) is vulnerable to this problem in cases where p and $(p \vdash q)$ are known as a conjunction (i.e., $K((p \vdash q) \wedge p)$). Of course there is no problem stating the case in this way, but for clarity I will use the two-premise version. Stanley’s reason for focusing on the more restrictive principle is orthogonal to accepting his “single-premise closure” and how it differs from ClosedComp. So the problem concerns not merely possible but actual views.

| | |
|---|---|
| <div style="text-align: center;"> <u>Comp.-1</u> $\neg K(q)$ </div> | <div style="text-align: center;"> <u>Comp.-2</u> $K(p)$ $K(p \vdash q)$ $K(q)$ </div> |
| <u>Subject Knowledge</u> $K(p)$ $K(p \vdash q)$ $K(q)$ | |

Cases that fit this mold quite naturally are familiar to epistemologists: Our subject, we have said, knows that he will not be able to afford the Safari trip in the safari-trip-compartment. He knows by closure in that compartment that he will not win the lottery. He therefore, as a subject (by Subject-K), knows that he won't win. In another compartment where he's concerned with the question of whether to buy a lottery ticket, he doesn't know that he will lose because he is attending, say, to the possibility that he will win (not on Safari prices).²⁵ So though in his lottery compartment he doesn't have the relevant knowledge, that isn't sufficient for him as subject to buy a ticket. Because as subject, he knows the ticket will lose. Closure problems of this type are easily generalized to question the solutions to problems that are resolved by knowledge shifts in accordance with attention or salience changes.

In general, then, a combination of ClosedComp and Subject-K can make trouble for many of the contextualist and Impurist responses to skepticism, lottery puzzles, and other closure related problems.²⁶

Should we get used to these kind of subject-compartment knowledge states? Should we deny closure altogether? Should we find and motivate some principle other than Subject-K? Though there may be other reasons for taking these options seriously,²⁷ there is a less radical measure that can resolve problems of this general type. The idea is not to require knowledge of entailment, but to require a logical operation. One good

²⁵ Naturally, it is possible to plug in other knowledge obstructing attention or salience conditions.

²⁶ I also note that even if there is a solution to this problem and the subject doesn't in fact know that he will lose the lottery, the advocates of these views still need to get used to claims like the following: *He knows he will not be able to afford a Safari trip but he doesn't know that he will not win this money in the lottery.*

²⁷ For the arguments I have in mind here, see Sharon and Spectre (2013a, 2017) and Spectre (2009).

formulation of closure that can avoid our problem is found in Williamson (2000) and Hawthorne (2003). This is the Multi Premise Closure (MPC) with compartmentalization added in:

MPC: Necessarily, if S knows p^1, \dots, p^n in one compartment, competently deduces q , and thereby comes to believe q while retaining knowledge of p^1, \dots, p^n throughout in the same compartment, then S knows q .

How does this formulation resolve the problem above? By inferring that q from p and thereby coming to believe that q , it would not be presumptuous to say that if p is, say “I cannot afford to go on Safari” and q is “I will not win the lottery” that the lottery compartment is no longer separated from the Safari compartment. Or at least that is the idea. So on this example by Hawthorne, plausibly, the subject doesn’t retain his knowledge that he will be unable to afford a safari trip in the safari-trip-compartment. This is because it seems right to say that by inferring that q and coming to believe it, he no longer knows that p (that he will be unable to afford the Safari trip).

Nevertheless, even if this resolves this problem, I am not altogether convinced that compartmentalized MPC can escape preface problem like the type Hawthorne levels against Lewis’s view (see above). That will depend on whether claims made in the book can all be known in one compartment, while in a second compartment there is knowledge of how likely it is that books contain mistakes. It seems to me that some such case can be constructed. If so, it is open to proponents of the relevant view to adopt the single-premise closure version of MPC, i.e. SPC. But I will not further speculate on the matter here. I think that a lot here depends on psychological facts that I (at least) do not know. Moreover, there is a different type of problem that I want to focus on.

8 KNOWLEDGE OF IGNORANCE

Suppose a subject knows that compartmentalized MPC (or SPC) is valid and knows that she doesn’t know that q . The following is a problem that could arise where we first spell out what follows from knowing that closure of this type holds. (\vdash is knowledge of an intermediate principle between MPC and SPC, parentheses are used to remove ambiguities, “ \vdash ” refers not only to the logical entailment but also represents the inferring

process from premise to logical conclusion, and “|=” refers to a priori (or axiomatic-like) truth that within compartments.):²⁸

1. $K(K(p \vdash q) \models Kp \supset Kq)$
2. $K(K(p \vdash q) \models \neg Kq \supset \neg Kp)$
3. $K(K(p \vdash q) \wedge \neg Kq \models \neg Kp)$
4. $K(K(p \vdash q) \wedge \neg Kq) \models K\neg Kp$
5. $KK(p \vdash q) \wedge K\neg Kq \models K\neg Kp$

But the subject does know that she doesn’t know that q and since it’s true *a priori* we can assume that she knows that p entails q (and can know that she knows this). She would now be in a position to know that she doesn’t know that p . The problem is, there is nothing stopping her—nothing we’ve said about compartments is stopping her, in any case—from knowing that p in some other compartment. But if so, we now have a contradiction. Assuming knowledge is factive and assuming Subject-K, that is. The reason is that knowledge collects to the subject level even when it is knowledge of ignorance. Here again the diagram may be useful.

Table 4: Knowledge of ignorance problem

| | |
|--|---|
| <u>Comp.-1</u> $K\neg K(q)$ | <u>Comp.-2</u> $K(p)$ |
| $5. KK(p \vdash q) \wedge K\neg K(q) \models K\neg K(p)$ | |
| <u>$KK(p \vdash q)$</u> $K\neg K(p)$ | <u>Subject Knowledge</u> $K(p)$ $K\neg K(p)$ \perp |
| | |

²⁸ The principle itself, even without knowledge of its truth, is weaker than the K-axiom $K(p \supset q) \supset (Kp \supset Kq)$ that with the inferences and assuming knowledge is retained is equivalent to MPC. It is stronger, though than SPC.

Now we did assume that the subject knows that MPC (or a weaker principle, to be exact) is valid and that is quite a heavy assumption. But it's not clear that we couldn't directly make an a priori type inference and come to know—to take Stewart Cohen's (1999) example—that since I don't know that my plane to New York won't crash, I don't know that my plane will be landing on the way to NY in Chicago. So it seems that even without knowing that MPC holds, I can come to know that I don't know my plane will land in Chicago. Or so it seems, at least. Assume I do. The question is whether what I know about my knowledge (or lack of knowledge) now prevents me from knowing something about whether the plane will stop on the way to New York in Chicago (by looking it up in the itinerary) in another compartment. We have already seen that I can't know this because of the contradiction. The question is why couldn't I know that I don't know.²⁹ Even if I don't yet know because, say, I haven't made any inferences and have not come to believe that I don't know that p , we need to say something about my being in a position to know. We need an answer, that is, to what goes wrong with my a priori reasoning that I don't know (and so know that I don't know).³⁰

There are several ways to resolve this inconsistency. One is to simply claim that one cannot both know that p in one compartment (say that one has hands) while knowing in another compartment that one doesn't know p (by deducing it from the knowledge of lack of knowledge that one is not a BIV). Unlike the earlier example where the very same proposition was at issue, however, it doesn't seem implausible—or at least not as implausible—that a subject remains compartmentalized in the way that puts her in a position to make the relevant inferences. Moreover, it doesn't seem plausible that this subject will lose her knowledge that she doesn't know by making the required inferences.

A second idea (we have considered) is that one cannot know that MPC is true and deny the a priori reasoning can lead to knowledge of lack of knowledge directly. Clearly this won't strike proponents of closure (both MPC and since we can make the same argument with one premise, also SPC advocates) as the right answer. But should open

²⁹ I am assuming that there is no problem in a compartment that has nothing to do with skeptical scenarios or whatever—there is no problem for me coming to know that my plane is stopping in Chicago. After all, if I didn't reason a priori in the first compartment we would not think of denying that I can know this in the second.

³⁰ Much of what follows I owe to Daniel Rothschild.

knowledge advocates be pleased with the potential to derive this inconsistency? It isn't clear at this stage. One reason is that we might not need closure at all to generate the knowledge of ignorance problem. An open knowledge contextualist would have to explain why one is unable to know that one doesn't know that p while having first order knowledge that p in another compartment. The answer here is simple, there can be no direct path to inconsistency because knowledge is factive. The indirect path to inconsistency goes through closure (or at least knowledge of closure). No matter how things turn out, a subject won't be able to know that she doesn't know that p by inferring it from knowledge of lack of knowledge that q . Nevertheless, the reason for this could be an inability to know that one doesn't know q in the first place, not an inability to transform this knowledge by inference into knowledge that p . So initially it isn't clear that the open knowledge advocate has any argument (from knowledge of ignorance) to support her view.

It should be clear, though, that these are not the only two options here. Subject-K is also a premise of the argument. In particular we may try and reconsider indexing to compartments as a way to replace Subject-K. Yet there is, I think, a better structural reason that explains how the knowledge of ignorance problem is avoided (or at least explained). It is simply the idea that compartmentalized subject have an obstacle in knowing that they do not know. The reason is implicit in Subject-K. In fact Subject-K says that to know that you don't know, you must know that for every compartment C , you don't know in C . And it seems odd—especially if as we have assumed, compartments are at least somewhat isolated from one another—that we could survey our compartments in the way required for knowledge of ignorance while retaining our Comp-K. In other words, even if we could survey compartments in the required way, it might no longer be clear that we would or could remain compartmentalized. (In this sense, Subject-K imposes a Lewisian type compartment elusiveness.)

In contrast to knowing you don't know, knowing that you know is much easier (on the Subject-K model). We can call this thesis an anti-Socratic thesis.³¹ All you need

³¹ Socrates (or Plato's Socrates) probably never states the paradoxical claim that in popular culture is often attributed to him: that the only thing he knows is that he knows nothing. The name, then, shouldn't be taken too seriously.

to know—and this could be achieved within a single compartment—is that you have knowledge in *a compartment*. An analogy may help here. Knowing that no one in a group knows that p , requires (besides it being true) that we check with every member of the group. To have knowledge that there is knowledge that p requires finding a member that knows. On Subject-K, once we find him or her, we have knowledge that the group has knowledge that p . So it is like comparing knowing that no one knows to knowing that someone knows. The former is harder knowledge to come by. Supposing the group members cannot communicate—as it would be if the group members were compartments of the same subject—knowing that no one knows might be very difficult.

There are surely propositions that we know no compartment can know. Trivial cases are propositions that we know are false. Less trivial but still unhelpful for present purposes are propositions that we know no compartment believes. And there surely are other general considerations that may lead us to believe that we could not know even if we did believe: *The number of stars is odd*, for instance. So it should be clear that the asymmetry between knowledge of ignorance and knowledge of knowledge only concerns limited cases. The odd thing is, though, that I can have a perfectly good reason to believe, say, that I do not know that I will lose the lottery, but this belief that I do not know will not be knowledge if I am compartmentalized. So if this seems implausible, maybe the open knowledge advocate has an advantage here after all. On the open compartmentalized knowledge theory there is no way to argue back from knowledge of ignorance of lottery propositions to everyday financial ignorance about the near future (not by closure, in any case). So by sound reasoning, the open knowledge advocate can know she doesn't know without risking contradiction with any everyday knowledge. Be that as it may, nothing analogous—not due to compartmentalization (even with closure), in any event—complicates knowledge of knowledge.

REFERENCES

- Cohen, S. (1999). Contextualism, skepticism, and the structure of reasons. *Philosophical Perspectives*, 13(S13), 57–89.
- Egan, A. (2008). Seeing and believing: Perception, belief formation and the divided mind. *Philosophical Studies*, 140(1), 47–63.
- Gettier, E. L. (1963). Is justified true belief knowledge? *Analysis*, 23(6), 121–123.
- Greco, D. (2015). Iteration and fragmentation. *Philosophy and Phenomenological Research*, 91(3), 656–673.
- Hawthorne, J. (2002). Lewis, the lottery and the preface. *Analysis*, 62(3), 242–251.
- Hawthorne, J. (2003). *Knowledge and Lotteries*. Oxford University Press.
- Lewis, D. (1982). Logic for equivocators. *Noûs*, 16(3), 431–441.
- Lewis, D. (1996). Elusive knowledge. *Australasian Journal of Philosophy*, 74(4), 549–567.
- Norby, A. (2014). Against fragmentation. *Thought: A Journal of Philosophy*, 3(1), 30–38.
- Salow, B. (2016). Lewis on iterated knowledge. *Philosophical Studies*, 173(6), 1571–1590.
- Sharon, A. & Spectre, L. (2013a). Epistemic closure under deductive inference: What is it and can we afford it? *Synthese*, 190(14), 2731–2748.
- Sharon, A. & Spectre, L. (2013b). Models of internalism. *MS*.
- Sharon, A. & Spectre, L. (2017). Evidence and the openness of knowledge. *Philosophical Studies*, 174(4), 1001–1037.
- Spectre, L. (2009). *Knowledge Closure and Knowledge Openness: A Study of Epistemic Closure Principles*. PhD thesis, Department of Philosophy, Stockholm University.
- Stalnaker, R. (1984). *Inquiry*. MIT.
- Stanley, J. (2005). *Knowledge and Practical Interests*. Oxford University Press.

Williamson, T. (2000). *Knowledge and its Limits*. Oxford University Press.

Williamson, T. (2013). Gettier cases in epistemic logic. *Inquiry*, 56(1), 1–14.

Williamson, T. (2015). A note on Gettier cases in epistemic logic. *Philosophical Studies*, 172(1), 129–140.