

## Deductive Closure as a Sorites

1. When not wearing our philosophical hats, most of us would say that we know many things about what will happen in the future. If some non-philosopher were to call right know and ask me if I knew where I would be in a week, I would answer without hesitation that I know I will be in New Jersey. Yet this knowledge claim may seem overconfident in the face of the wide variety of ways in which things can go wrong. We do not have to invoke any fantastic science-fiction scenarios in order to raise doubts about knowledge of the future, as we would to raise doubts about my knowledge that I am sitting in my apartment typing on my computer. Every week a few apparently healthy people die suddenly, and there is a small chance that I will be one of those people. It seems as though this small chance prevents me from knowing where I will be next week. For if I can't rule out that I'll be six feet under next week, and I can't, I apparently must admit that there's some small chance that I won't be in New Jersey.<sup>1</sup>

One response is a fallibilist approach to knowledge, on which knowledge does not require ruling out every remote counter-possibility.<sup>2</sup> In this case, we might say that I can know that I'll be in New Jersey even though I can't rule out the remote possibility that I will die suddenly. This approach, however, seems to threaten our ability to deduce new knowledge from old knowledge. If I will be in New Jersey next week, then I will not be dead next week. So if I can know that I will be in New Jersey next week, then I should be able to deduce that I will not be dead next week; and this is precisely what I cannot

---

<sup>1</sup> See Vogel (1990) for many similar examples and an argument that they do not provide counterexamples to deductive closure.

<sup>2</sup> This approach is fallibilist in the sense of Cohen (1988), in which we can know things that are not in fact entailed by our evidence, rather than in that of Lewis (1996), in which we can say in one breath that someone knows and that she has not ruled out an alternate possibility.

know. This apparent failure of deductive closure is well known, and provides one motivation for contextualism about knowledge. According to (some) contextualist theories, in a context in which you are contemplating your whereabouts next week, I can know that I will be in New Jersey; in a context in which I am contemplating my lifespan, I cannot know that I will not die in the next week, nor that I will be in New Jersey next week. (And that is why I have to wait for my non-philosopher friend to call in order to claim knowledge; I must forget the sentences that I am typing now.) Deductive closure only seems to fail because the attempt to deduce that I will not die from my knowledge that I will be in New Jersey shifts the context, destroying my old knowledge.

Fallibilism can also lead to a less familiar apparent failure of deductive closure, as pointed out by Hawthorne (2002, 2004). According to fallibilism, we may know something even if we have not ruled out some counter-possibilities with a small but nonzero likelihood.<sup>3</sup> If a set of propositions  $P$  comprises many such fallible pieces of knowledge, the counter-possibilities to each piece may compound, so that it would be quite likely that some proposition or other in  $P$  is false. (Though in fact, on the supposition that each is a real piece of knowledge, each must be true; fallibilism does not allow that we may know falsehoods.) Then it seems implausible that we can know the proposition “Every proposition in  $P$  is false,” since this proposition itself is unlikely on our evidence; yet this proposition may be deduced from the propositions in  $P$ , each of which is *ex hypothesi* known. Contextualism will not solve this problem, as it aspires to

---

<sup>3</sup> Note that this goes beyond the ability to know even if you have not ruled out brain-in-the-vat scenarios and the like. Arguably there is no non-zero probability that we should assign to brain-in-the-vat scenarios; so the probability of error need not compound across a large set of propositions known by that standard. In addition, if the only possibility of error rests on a brain-in-the-vat scenario, then there is no additional risk in conjoining a set of propositions known by this standard; if one is true (and you are not a brain in a vat) then all are true. So the proponent of unqualified deductive closure need not be disturbed by Hawthorne’s argument as applied to brain-in-a-vat scenarios.

solve the other alleged failure of deductive closure. Contextualist theories do not yield an obvious account of how contexts might shift in the course of this deduction, so that the attempt to draw the conclusion destroys the knowledge of some of the premises.

Nor, I will argue, should they. The fault lies not in the fallibilist ascription of knowledge of the premises (and denial of knowledge of the conclusion), but in the version of deductive closure that is used in this argument. Deductive closure is a sorites premise, which seems intuitively obvious but which leads to absurdities if it is accepted without qualification. Indeed, this result should be less surprising than it is. Deductive closure's intuitive appeal comes from the way that it seems to embody our practices of deriving new knowledge from old knowledge; but if we examine those practices more closely, we will see that they do not support the unqualified use of deductive closure in many-premise arguments. The contextualist attempt to account for apparent failures of single-premise closure can then stand or fall on its own merits (which I will not attempt to evaluate); the reason that apparent failures of multi-premise closure cannot be accounted for in the same way is that in those cases multi-premise closure really does fail.

2. I will borrow most of John Hawthorne's excellent exposition of how deductive closure makes trouble for contextualist fallibilism (Hawthorne 2002, 2004). As Hawthorne points out, there are two different kinds of deductive closure, one more unshakable than the other. The unshakable version involves only one premise:

*Single-Premise Closure (SPC)*. Necessarily, if S knows  $p$ , competently deduces  $q$ , and thereby comes to believe  $q$ , while retaining knowledge of  $p$  throughout, then S knows  $q$  (Hawthorne 2004, p. 34).

The less unshakable version can involve multiple premises:

*Multi-Premise Closure (MPC)*. Necessarily, if S knows  $p_1, \dots, p_n$ , competently deduces  $q$ , and thereby comes to believe  $q$ , while retaining knowledge of  $p_1, \dots, p_n$  throughout, then S knows  $q$  (Hawthorne 2004, p. 33).

The first modification I would like to make to these principles involves what counts as a premise. The inference from “I will be in New Jersey next week” to “I will not be dead next week,” considered as a deduction, is enthymematic; it requires the additional premise “If I am dead next week, I will not be in New Jersey next week.” This additional premise, however, is virtually certain; if it has any empirical content whatsoever, that content concerns only the general fact that when one dies, one ceases to be anywhere.<sup>4</sup> In considering deductive closure, I will not count such virtually certain premises as premises. So the inference from “I will be in New Jersey next week” to “I will not be dead next week” will count as an instance of SPC. (It would make little difference if we counted it as a two-premise argument, since two-premise arguments will turn out to be acceptable in almost every case.)

Counting premises matters not just for SPC but also for the version of MPC I will consider. When we draw inferences from a set of premises, we do not usually lay out a

---

<sup>4</sup> I am supposing that for one’s corpse to be in a cemetery or morgue in New Jersey is not for oneself to be in New Jersey.

thousand premises, say “Ah,  $q$  follows from these,” and come to believe  $q$ . Indeed, it is dubious whether such a deduction would count as competent, since people with our limited capacities would be likely to make mistakes in performing such thousand-premise leaps. What we do instead is to gather a few premises together, consider what follows from them, take this intermediate conclusion as one of a few premises in the next deduction, consider what follows from them, etc. Hence, I will consider not MPC but *Few-Premise Closure*:

*Few-Premise Closure (FPC)*. Necessarily, if S knows  $p_1, \dots, p_n$ , competently deduces  $q$ , and thereby comes to believe  $q$ , while retaining knowledge of  $p_1, \dots, p_n$  throughout, then S knows  $q$ , if  $n$  is no higher than 5 or so.

Obviously, any conclusion that could be obtained by MPC could be obtained by repeated applications of FPC. But the repetition will matter; for I will be arguing that FPC is not true without qualification, but is a sorites premise. So repeated application of FPC can cause a lot of trouble, even if a single application of FPC can cause only a little.

As Hawthorne observes, SPC is tremendously intuitive. It may seem as though we might want to deny SPC in order to (for instance) “align itself with our instinctive verdicts about what we can and cannot know by perception” (Hawthorne 2004, p. 46).<sup>5</sup> Thus it seems that I can know by perception that I have hands, but I cannot know just by perception that I am not a handless brain in a vat, even though the latter follows from the former. Yet, as Hawthorne argues, an attempt to salvage these intuitive judgments by

---

<sup>5</sup> Hawthorne is here discussing the account of Dretske (2000), but he remarks that other accounts that deny SPC fail in the same way.

denying SPC will prohibit all sorts of inferences that do intuitively yield knowledge.<sup>6</sup> I agree with Hawthorne's treatment of SPC and will not address it further here.

Hawthorne observes also that "MPC seems tremendously intuitive: the idea that one can add to what one knows by deduction from what one knows has a powerful grip on us regardless of whether the deduction proceeds on one premise or many" (Hawthorne 2004, p. 46). (I will address his defense of MPC after presenting my case against it.) Stated thus, MPC does seem tremendously intuitive. At least, given (as discussed above) that most knowledge-preserving deductions proceed a few premises at a time, FPC seems tremendously intuitive, and MPC follows from it. Yet I will argue that FPC is like sorites premises such as "Any two hues that are visually indistinguishable are the same color." Such premises are also tremendously intuitive, and are safe to apply a few times. Nevertheless, they cannot be held true without qualification. Similarly, applying FPC a few times will never take us from premises we know to conclusions we don't know, but applying it many times sometimes will.

Indeed, a close examination of how we argue from old knowledge to new will make unqualified MPC seem less intuitive, in exactly the way that we would expect if FPC were a sorites premise. We are simply not inclined to employ FPC repeatedly without safeguarding our arguments so that they do not gradually take us from knowledge to not-knowledge. Thus, I will argue, the best prospect for preserving our ordinary knowledge attributions is to treat FPC as a sorites premise.

---

<sup>6</sup> See Hawthorne (2004), pp. 36-46.

3. Let us begin by reviewing Hawthorne's discussion of how MPC causes trouble for the contextualist. The contextualist's account of apparent failures of SPC is familiar (and already sketched above). A contextualist may hold that it can truly be said that I know that

(1) My feckless friend Bill will never be rich

while respecting the intuition that it cannot truly be said that I know that

(2) Bill's ticket will not win the lottery tomorrow.<sup>7</sup>

The idea is (roughly) that knowledge that  $p$  requires ruling out all alternatives to  $p$  that are relevant in the context of the ascription. In a context in which knowledge of (1) is being ascribed, the possibility that Bill's ticket is drawn is not relevant, and so I can be said to know (1). I could also be said to know (2) in that context, if saying that I knew (2) did not change the context and thus which alternatives are relevant. In fact, once we consider ascribing knowledge of (2), we change the context so that the possibility that Bill's ticket is drawn is relevant. As Lewis would put it, we are paying attention to the possibility that Bill's ticket is drawn, so we are not properly ignoring that possibility.<sup>8</sup> In this context, I cannot be said to know (2), nor can I be said to know (1). No matter what context, if I know (1) then I know its consequent (2). SPC is preserved within each context.

Here shifting one's attention to a new type of question, whether Bill would win the lottery as opposed to whether Bill would be rich, shifted the context and thus led to a seeming violation of deductive closure. Do all seeming violations of closure involve

---

<sup>7</sup> The example of feckless Bill is due to Lewis (1996, p. 565).

<sup>8</sup> See Lewis on the Rule of Attention (Lewis 1996, p. 559).

shifts of attention and attendant shifts of context? Hawthorne (2002; see also 2004, pp. 94-98) has shown that such a view runs into trouble. Even if we keep our attention tightly focused on questions like “Will Bill ever be rich?”, unqualified FPC (i.e., MPC) will lead to trouble if we allow knowledge of propositions like (1).

Consider the following situation: Alice has 5000 feckless friends, each of whom holds one ticket in tomorrow’s lottery. The only way any of Alice’s friends will become rich this year is to win that lottery. The lottery has 5001 tickets, one held by Dr. Evil, who is not Alice’s friend. Sarah asks Alice in turn, of each of her friends, “Will Bill be rich this year? Will Harry be rich this year?” etc. Alice replies, in each case,

( $3_{\text{Bill/Harry/etc.}}$ ) Bill[Harry/etc.] will not be rich this year.

In fact, Dr. Evil’s ticket wins, so none of Alice’s friends is rich this year. Each of her statements (3) turns out to be true. Looking back at year’s end, should we say that Alice knew that Bill would not be rich, that Harry would not be rich, etc.?

If contextualism is to support fallibilism about lottery cases, the contextualist must say that there is a standard for knowledge by which, when we consider whether Bill will be rich this year, we may ignore the possibility that Bill’s ticket wins. Alice, however, seems to stick to one standard as she considers whether Bill will be rich, whether Harry will be rich, etc. If, by a single standard, Alice knows ( $3_{\text{Bill}}$ ) and ( $3_{\text{Harry}}$ ) and the rest, then by MPC within that standard she knows

(4) None of the 5000 friends will be rich by year’s end.



On Alice's evidence, however, (4) has only a 1 in 5001 chance; (4) will not be true unless Dr. Evil's ticket wins. Though, looking back, we know that Dr. Evil's ticket did win, it is outrageous to say that Alice was in a position to know (4).

Hawthorne points out that the contextualist can wriggle out of this problem by positing that Alice does shift standards. One could say: On the Bill-standard for knowledge, one may properly ignore the possibility that Bill's ticket will win, but not that Harry's ticket will win, or Jerry's, etc. On the Harry-standard for knowledge, one may properly ignore the possibility that Jerry's ticket will win, but not Bill's, or Jerry's, etc. When evaluating (3<sub>Bill</sub>), the Bill-standard is appropriate, so it is proper to say that Alice knows (3<sub>Bill</sub>). When evaluating (3<sub>Harry</sub>), the Harry-standard is appropriate, so it is proper to say that Alice knows (3<sub>Harry</sub>). But MPC only governs premises that are all known by the same standard. (4) is the conjunction of one premise that is known by the Bill-standard, one that is known by the Jerry-standard, one that is known by the Harry-standard, etc.; Alice's knowledge by each of these standards may be deductively closed without her knowing (4) by any standard.

This multi-standard solution preserves fallibilism and unqualified MPC/FPC, but it has little else to recommend it. As Hawthorne points out, this "solution to our lottery puzzle require[s] rapid context shifting where, initially, context shift was far from noticeable" (Hawthorne 2002, p. 251). Worse yet, it wreaks havoc on the role of deduction in our epistemic practices. Suppose that Bill and Harry are going on vacation together, and Alice is wondering whether they will be able to afford a certain hotel. She reasons:

(3<sub>Bill</sub>) Bill will not be rich this year.

(3<sub>Harry</sub>) Harry will not be rich this year.

(5) Therefore neither Bill nor Harry will be rich next week.

(6) Therefore they will not be able to afford the hotel.

On the multi-standard solution, Alice knows (3<sub>Bill</sub>) by the Bill-standard and (3<sub>Harry</sub>) by the Harry-standard. Since these are different standards, it does not follow by FPC that she knows (5), even though (5) is a deductive consequence of (3<sub>Bill</sub>) and (3<sub>Harry</sub>). If Alice is to come to know (5) and (6) by deduction from (3<sub>Bill</sub>) and (3<sub>Harry</sub>), she must first rederive her premises under a single standard.

It would be nightmarish to constantly recheck the foundations of our knowledge in this way. The deduction from (3<sub>Bill</sub>) and (3<sub>Harry</sub>) to (5), in particular, is unlike the deduction from (1) to (2), which may reasonably be taken to require rechecking our reasons for believing the premise. When we reason from (1), that Bill will never be rich, to (2), that Bill's ticket will not win, we realize that our acceptance of (1) required ignoring the possibility that (2) might be true; ignoring that possibility might have been reasonable when considering (1), but it is not reasonable when considering (2). Focusing on (2) raises the new doubt, "What if Bill's ticket does win?" No such new doubt is raised in the deduction from (3<sub>Bill</sub>) and (3<sub>Harry</sub>) to (5). To get to (3<sub>Bill</sub>) and (3<sub>Harry</sub>), Alice must have ignored the possibility that their respective tickets win; if this was proper, it is proper to ignore these possibilities when considering (5). For Alice has not refocused her attention on the tickets; she is still considering whether her friends will be rich. If Alice really was in a position to know the premises (3<sub>Bill</sub>) and (3<sub>Harry</sub>), the deduction of (5) from these premises can raise no new worries and can require no rechecking of her premises.

Eventually, it seems Alice will be driven to the conclusion (4), that none of her 5000 friends will be rich at year's end; at no step does adding an extra premise (say,  $3_{\text{Jerry}}$ ) raise a new doubt, and so there is no step where the context shifts so as to destroy knowledge of the conclusion; so it seems.

4. Let us examine the argument for (4) step by step. Suppose that Alice never does slip into contemplating her friends' lottery tickets, nor do we ascribers of knowledge, so there are no context shifts. Alice asks herself "Will Bill be rich this year?" and answers "No." On the standard that applies to Alice's thoughts, the 1 in 5001 chance that Bill's ticket will win the lottery is remote enough to ignore, and so Alice does know ( $3_{\text{Bill}}$ ), that Bill will not be rich this year.<sup>9</sup>

Alice now asks, "Will either Bill or Harry will be rich this year?" and answers "No." Given that she already knows ( $3_{\text{Bill}}$ ), the only way for the true answer to come out "Yes" would be for Harry to become rich, which would require that Harry's ticket win, which is a 1 in 5001 chance that may properly be ignored.<sup>10</sup> Adding this negligible chance of error to something that Alice already knows cannot destroy her knowledge. So Alice knows the conjunction ( $(3_{\text{Bill}}) \& (3_{\text{Harry}})$ ). Alice then asks, "Will Bill or Harry or Jerry be rich this year?" and answers "No." Adding Jerry to the mix adds only a 1 in 5001 chance of error, which negligible chance cannot destroy Alice's knowledge, so Alice knows ( $(3_{\text{Bill}}) \& (3_{\text{Harry}}) \& (3_{\text{Jerry}})$ )....

---

<sup>9</sup> In fact, it is debatable whether this chance is negligible, given the low number of tickets in the lottery. Even a fallibilist might not allow that we know that Bill will never become rich, given the high expected value of a 1 in 5001 chance of becoming rich. If this causes concern, substitute a more tractable example.

<sup>10</sup> Recall that, in ascribing this knowledge, we are looking back from the time when Dr. Evil has won the lottery; so none of Alice's 5000 friends become rich this year. It would not be proper to neglect a possibility that actually came to pass.

As Lewis would say, Blind Freddy can see where this is going. Eventually Alice gets to (4), that none of her 5000 friends will win, and we conclude that Alice knows (4), which had at best a 1 in 5001 chance on Alice's evidence (even though it turned out to be true, since Dr. Evil's ticket won). Each step was an application of FPC, and each step seemed innocuous—in fact no step but the first could be less innocuous than any other—but the cumulative effect is disastrous. One conclusion that might be drawn is that contextualist fallibilism is wrong, and that Alice cannot be said to know (3<sub>Bill</sub>) on its own. Another conclusion is that we have described what would happen if FPC were sorites premise and should investigate whether it is one.

I will not even hint at a resolution of sorites paradoxes here; but what they have in common is that an intuitive hypothesis, which causes no apparent trouble when applied a few times, leads to absurd conclusions when applied repeatedly.<sup>11</sup> If we begin with a pile of sand that is definitely a heap, taking away one grain of sand will not yield a non-heap. It is not safe, however, to assume that the result of this operation will always be a heap from which we may again take away a grain, yielding another heap from which we may again take away a grain, etc. Similarly, applying FPC once or twice leads to know trouble. If we are happy to say that Alice knows that a single one of her friends will never be rich, we should be happy to allow her to conjoin a few of these premises and conclude that a few of her friends will never be rich. Trouble ensues when we assume that these known conjunctions may be taken as premises for further applications of FPC,

---

<sup>11</sup> The sorites premise leads to no *apparent* trouble when applied once, but that is not to rule out that it may lead to real trouble when applied once. On epistemicist theories of vagueness (see Williamson 1994), there is a definite dividing line between heaps and non-heaps, but we cannot know where that dividing line is. So there is one application of the sorites premise that takes us from the true to the false; but that application leads to no apparent trouble, since we cannot know which application takes us from the true to the false.

which yield further known conjunctions that may be taken as premises for further applications of FPC, which yield further known conjunctions....

As remarked above, we have the option of concluding that Alice never knew the premises to begin with. This would not be an acceptable option for other sorites paradoxes; we cannot conclude that there are no real heaps, or that a hue is not the same color as itself. But just as the individual inputs to the supposed sorites are open to question here, so is the general premise, FPC. For we had an argument for why FPC should be intuitively appealing: that we can add to what we know by deducing from multiple premises. That argument in fact supports treating FPC as a sorites premise rather than taking it to be true without qualification.

For we have and should have no inclination to perform the sort of deduction that unqualified FPC licenses.<sup>12</sup> The problem with Alice's deduction that none of her friends would become rich is that the negligible chance that each individual premise is false compounds as she conjoins them, so that well before she reaches the conclusion the chance of falsity is non-negligible. No one would ever attempt to perform such a deduction in real life. Even when we are not dealing with actual lotteries, we do not and should not string together arbitrary long conjunctions of ordinary empirical sentences and deduce something that relies on their conjunction. Consider how such an attempted deduction would proceed. It could bring together information from a variety of sources: personal observation, the testimony of others, deductions from people's habits, inferences to the best explanation, etc. Each piece of information, let us suppose, is necessary if the conclusion is to follow. (Below I will discuss redundancy, in which the argument is constructed so that each piece of information is not necessary.) Such an argument will be

---

<sup>12</sup> Thanks to [suppressed for blind review] for emphasizing this point.

look like a conspiracy theory in which elaborate charts are drawn linking the participants through twenty intermediaries, or the solution to an Agatha Christie mystery in which every piece of evidence is explained through some baroque imposture. Someone who made such an argument could rightly be criticized on the grounds that a single flaw anywhere in the edifice of her reasoning would bring the whole structure crashing down. In ordinary cases of empirical knowledge, she may have checked each proposition well enough to establish that she knows it, without having checked it well enough to establish it with the certainty that would be required to make it suitable for use in such an elaborate argument.<sup>13</sup>

This means that all knowledge is not created equal. Some things a person knows are suitable without reservation for use as premises for further few-premise deduction. These will be things that she knows non-inferentially, or that she has inferred using only a few premises that are not virtually certain.<sup>14</sup> Care must be taken when using other things she knows as premises for deduction. These are the conclusions of deductions from several different premises of the first kind. Deduction from a few of these premises may generate something that is not determinately knowledge, if the deduction is constructed in such a way that the possibilities of error in each of the premise compound in the

---

<sup>13</sup> Compare the Preface Paradox of Makinson (1965), also discussed by Hawthorne (2002), in which an author asserts in a preface that there is at least one mistake in the main body of her work. If the empirical assertions in the body have been verified independently, say through examination of separate probate records and the testimony of many eyewitnesses to separate incidents, then it is rational for the author to believe that she has made at least one mistake even though it is rational for her to believe, for each individual assertion, that she knows that that is true. If we do allow that she knows each individual assertion in the body, we should still deny that she knows the conjunction of every assertion in the body, since it is rational for her to believe that this conjunction is false. Again, the checking that is sufficient to establish knowledge of each individual assertion is not enough to establish that the assertion is suitable for use in arbitrary conjunctions, unless the conjunction is constructed so the possibility of error in each assertion does not compound when the premises are conjoined. See also n. 15, below.

<sup>14</sup> In keeping with the focus on *deductive* closure, we might treat non-deductive inferences as enthymematic deductive inferences, in which the suppressed premise is known (if the conclusion is true in a non-Gettierized fashion) with less than virtual certainty.

conclusion. The wise knower will not make such deductions, or will recheck her conclusions when she does; as, according to the contextualist, the deduction that Bill will never win the lottery from our knowledge that Bill will never be rich destroys our knowledge of both by obliging us to reconsider our reasons for thinking that Bill will never be rich.

It would be disastrous if the result were that we could never engage in long chains of inference. Fortunately, the only long deductions that are endangered by treating FPC as a sorites premise are those that compound the possibilities of error for each premise, and it is possible to construct a deduction in which the possibilities of error do not compound. What is necessary is to build in some redundancy in the engineer's sense, in which parts of a structure are supported in more than one way, so that the failure of one part will not bring down the whole. We can obtain knowledge by structuring our deductions so that the falsity of a single premise does not undermine the whole. Conclusions obtained through such deductions may be as suitable for use in new deductions as the original premises were.

Building in this redundancy avoids the sorites paradox, because the intermediate steps can be verified to be determinate instances of knowledge. Compare the original sorites premise. It may be that  $n$  applications of the sorites premise, for some heaps, will take you from a definite heap to something that it would be absurd to judge a heap. If, however, after taking away  $n$  grains of sand from one particular heap, you are in a position to verify that the remainder is definitely a heap, then you can take away another grain of sand without fearing that the result will be something that it would be absurd to judge a heap. Similarly, when an argument has redundancy built in, the intermediate

steps will be more secure than knowledge gained through FPC sometimes is, so applying FPC to these intermediate steps will not lead to a sorites paradox.

Indeed, real arguments from empirical premises have this redundancy. Consider the famous physics paper with 99 authors discussed by Hardwig (1985). The research behind the paper incorporates a huge number of observational premises, but the paper's conclusions surely do not depend on every single observation being accurate. Part of the reason that it is necessary to make so many observations is to keep a single measurement error from spoiling the whole conclusion. Every individual observation may be made in conditions that allowed for knowledge of the result, for instance that the apparatus is working and the experimenter's vision is sound, etc. Still, we would forbear from drawing a conclusion that required each of these observations to be accurate, and we would advise anyone who drew such a conclusion to do more double-checking.<sup>15</sup> In mathematics we do build long deductive chains, in which the failure of any step would bring down the whole; but in mathematics the premises are non-empirical and virtually certain, so they do not require us to invoke FPC.

---

<sup>15</sup> Here it is instructive to compare Mark Kaplan's argument that the paradox of the preface (see note 13 above) is no paradox; that an author should "be prepared to assert in the context of inquiry that everything she asserts in her book is true," even though she knows that it is very unlikely that everything she has asserted is true (Kaplan 1996, p. 118). Kaplan argues that, since any ambitious empirical theory or historical narrative is overwhelmingly likely to be false, and since we cannot give up theorizing or narration, we must be prepared to assert and believe things that are overwhelmingly likely to be false. (Kaplan connects this to a proposed revision of the notion of belief; see Kaplan 1996, pp. 112-148).

In response, we may have reason to assert every component of some ambitious theory or narrative, but this does not mean that we have reason to assert any consequence of such theories that relies on every single component of that theory. If failure of any plank whatsoever in the narrative would lead to failure of the conclusion we wish to draw, then we are best advised not to draw that conclusion on the basis of the narrative alone. Fortunately, the conclusions we actually draw from ambitious theories and narratives will usually have redundancy built in. If I read a book about the French Revolution, my conclusions about the causes of the Revolution may well survive the discovery that the author got some details of the history wrong. So we lose nothing by refusing to assert that our ambitious theories and narratives are completely error-free.



The intuitive basis for MPC and FPC encourages us to treat FPC as a sorites premise rather than to accept it without qualification. MPC seemed tremendously intuitive because of the idea that we can gain new knowledge by deduction from many pieces of old knowledge. Indeed, we can gain new knowledge by this deduction, so long as our arguments do not require repeated applications of FPC. Arguments that do require repeated applications of FPC, however, are either arguments that we do not and should not use to extend our knowledge, unless we are in a position to verify that negligible chances that each premise is mistaken do not compound so that there is a non-negligible chance that the conclusion is mistaken.

So the argument that Alice knows that none of her friends will be rich poses no particular problem for contextualist fallibilism. Contextualism aspires to account for our willingness to say that Alice knows that Bill will never be rich and not that he will not win the lottery in terms of the different contexts in which the two propositions are considered. (I have not addressed whether that argument succeeds on its own merits.) The argument that Alice can know that none of her friends will be rich cannot be blocked in the same way, by positing that the context shifts in the course of the argument. But that argument can be blocked without positing a context-shift, because it uses FPC repeatedly in a way that runs afoul of the sorites paradox.

5. The argument over whether FPC is a sorites premise in fact has little to do with contextualism. Fallibilism leads to apparent failures of SPC, which the contextualist explains away by adverting to context shifts. Hawthorne shows that fallibilism also leads to apparent failures of MPC, and that context-shifting maneuvers cannot be used to

explain these away. My response is that these are real failures of MPC, resulting from repeated use of FPC in a sorites-like fashion; so the contextualist is under no burden to explain them away. My argument lends equal support to non-contextualist fallibilist theories, which are also freed from the burden to explain apparent failures of MPC; and Hawthorne's argument creates trouble for non-contextualist fallibilist theories as much as for contextualism.

Since Hawthorne himself acknowledges that MPC can cause trouble for any non-skeptical view, we should consider his suggestions for accommodating MPC.<sup>16</sup> Hawthorne proposes that we treat knowledge in terms of the knower's practical environment: "Insofar as it is unacceptable... to use a belief that *p* as a premise in practical reasoning on a certain occasion, the belief is not a piece of knowledge at that time" (Hawthorne 2004, p. 176).<sup>17</sup> The practical environments in which we might call on individual premises are different from those in which we might call on long conjunctions. In the first sort of environment, we know the premises; by MPC we also know their conjunction, but only because we could never put them to use. In an environment in which we might use the conjunction, we do not know either the conjunction or the premises.<sup>18</sup> When Alice should be worrying about whether any one of her friends will be rich, say because she is concerned with whether that friend will be able to afford a certain

---

<sup>16</sup> See, for instance, (Hawthorne 2004), p. 181: "[t]he path [to salvaging MPC] is... a difficult one." Hawthorne's project is not so much to advocate any particular theory of knowledge as to illuminate the strengths and weaknesses of the various theories, so it may not be fair to ascribe the views discussed in my text to him, although he does suggest that these are the proposals he would adopt if you put a gun to his head (p. 188).

<sup>17</sup> Although Hawthorne presents this as a form of invariantism, he points out that contextualism may also make use of practical environment, see (Hawthorne 2004), p. 188n53.

<sup>18</sup> See Hawthorne 2004, p. 183. Hawthorne invokes similar considerations, though these do not advert to practical environment, in his original argument for MPC: "the settings where we most vividly wish to deny that one can know the conjunction... are settings where we are also somewhat unwilling to say that the conjuncts are known" (Hawthorne 2004, p. 49).

hotel, then she does know that that friend will never be rich. By MPC she can come to know that none of her friends will every be rich, and that Dr. Evil's ticket will win. But if that conjunction is relevant to her practical reasoning, for instance if she is trying to figure out a fair price for Dr. Evil's lottery ticket, then she does not know that her friends will never be rich, nor even that any one of her friends will ever be rich. When she starts to reason about whether her friends will be able to afford the hotel again, she again knows that her friends will never be rich. As Hawthorne says, the practical environment view "will make knowledge come and go with ease" (Hawthorne 2004, p. 176).

It is not clear, however, that the practical environment in which we use long conjunctions always will be different from that in which we use the individual premises. Take the following situation: Alice parks cars on-street overnight for customers, noting their location. An average of one car per day is stolen in the neighborhood.<sup>19</sup> Alice also has a friend who runs urban orienteering contests, where players must follow instructions such as "the next clue is chalked on the sidewalk next to a green Volkswagen." Alice's friend often asks her if she knows where she has parked certain cars or groups of cars, to help him design the course that he will run that night. If there is a mistake in the course, it is not catastrophic, though it will spoil the game for the night.

Today Alice has parked one thousand cars, none of which have in fact been stolen. Her friend is about to give her a list of cars whose locations he wants to know. The list may contain one car or one thousand. Is Alice in a practical environment such that she knows where all thousand cars are, or is she in an environment such that she does not know where any individual car is? Neither answer is satisfactory. If Alice does not know where any of the individual cars are, then she cannot answer her friend if he asks

---

<sup>19</sup> The example of the possibly stolen car is due to Vogel (1990).

her if she knows the location of a single car; yet the stakes are low enough that Alice ought to be able to answer this question without rechecking whether the car has been stolen. Yet it is as odd to say that she knows the conjunction of the statements of every car's location, for on her evidence that statement is likely to be false (though in fact it is true); if her friend uses all these cars in designing his course, then she has reason to expect that his game will be spoiled. Nor will it do to say that the practical environment depends on the question she is asked. If the fact that Alice is about to be asked a question on a topic is not enough to determine her practical environment, then it is hard to see how her practical environment can be fixed with respect to  $p$  in any case in which she is not explicitly considering using  $p$  in practical reasoning. If practical environments are underdetermined in this way, then the practical environment proposal will not help settle the truth of most knowledge ascriptions.

By contrast, treating FPC as a sorites premise provides a satisfactory account of Alice's knowledge. For any individual car, Alice knows where that car is. For any two cars, Alice knows where those cars are. Alice does not know where all thousand cars. For some intermediate values of  $n$ , Alice's belief concerning where  $n$  of the cars are is a borderline case of knowledge. All this is as usual with sorites premises. Furthermore, Alice can give her friend a map of where all thousand cars are and say "I know that most of the cars are as marked on this map"; the deduction of this statement from premises she knows has the required redundancy, since the theft of one car would not falsify the conclusion.

In general, it is not clear that reliance on practical environment be of any use if we wish to salvage MPC. Hawthorne is undoubtedly correct to say that, when some practical

issue turns on a long conjunction of fallible premises, the conjunction should not be used as a premise in practical reasoning. That seems simply to mean that the possibility that the conjunction is mistaken is too likely to be neglected. There is no guarantee, however, that such a practical environment will rule out using the individual conjuncts as premises for practical reasoning. For if practical issues also turn on the individual conjuncts, we may well be able to use those conjuncts as premises for practical reasoning. The possibility of error for each conjunct may be negligible with respect to the practical issues at hand, while the possibility of error for the conjunction is non-negligible. That is exactly what we would expect if FPC were a sorites premise.<sup>20</sup>

The practical environment approach, then, is no help in reconciling fallibilism with unqualified MPC. Considering the practical environment does allow us to avoid skepticism; you may know where your car is without ruling out the possibility that it has been stolen, and you may know where you will be next week without ruling out the possibility that you will suddenly die, etc. It will not, however, confirm that within a single environment you can know arbitrary conjunctions of the things you know.

Treating FPC as a sorites premise helps here, as in other non-skeptical contexts, because

---

<sup>20</sup> Hawthorne raises another problem for MPC that he takes to cause trouble for the practical environment approach. The problem is this: Suppose that you come to know, in the ordinary fashion, that each of 150 people will be at the APA. Then someone announces (falsely) that one of the 150 has just died, without specifying which. Though you have every reason to trust the informant, you ignore her. Hawthorne argues that this counterevidence destroys knowledge of the conjunction “All 150 people will be at the APA” but doubts that this knowledge of any of the individual conjuncts. (See Hawthorne 2004, pp. 48-50 and 183-4).

I do not share Hawthorne’s intuition that knowledge of the individual conjuncts remains intact. Nevertheless, my approach can accommodate this intuition. After the false announcement, the probability that S will not be at the APA for each individual S may increase from a negligible number to a number that is larger but still negligible. Then, if the announcement does not make salient the possibility that S has died for any individual S in a way that prevents the subject from properly neglecting it, the subject may still know that S will be at the APA, while the conjunction is not known. In this case, it will be the case that for some  $n$ ,  $n$ -conjunct conjunctions that were known before the announcement will be borderline cases or unknown after the announcement.

it does not simply forbid us from deducing new knowledge from old. Most deductions from old knowledge will yield new knowledge; we must simply avoid the sorites effect.

The fate of MPC, then, is independent of the fate of contextualist fallibilism.

Whether or not we conclude that apparent failures of single-premise closure are due to a context-shift, we should treat few-premise closure as a sorites premise. Within a single standard for knowledge, we can know many individual propositions without knowing their conjunction, or other propositions that are deduced from them in a non-redundant way. This treatment of deductive closure grants us ordinary knowledge of the future without committing us to knowledge of unlikely events and respects the way we use deduction to gain new knowledge from old.

### **Works Cited**

Cohen, Stewart (1988). "How to Be a Fallibilist." *Philosophical Perspectives* 2, 91-123.

Dretske, Fred I. (2000). "Epistemic Operators." In *Perception, Knowledge, and Belief* (Cambridge: Cambridge University Press).

Hardwig, John (1985). "Epistemic Dependence." *Journal of Philosophy* 82, 335-349.

Hawthorne, John (2002). "Lewis, the Lottery, and the Preface." *Analysis* 62, 242-251.

Hawthorne, John (2004). *Knowledge and Lotteries*. Oxford University Press, Oxford.

Kaplan, Mark (1996). *Decision Theory as Philosophy*. [publisher]

Lewis, David (1996). "Elusive Knowledge." *Australasian Journal of Philosophy* 74, 549-567.

Makinson, D.H. (1965). "The Paradox of the Preface." *Analysis* 25, 205-7.

Vogel, Jonathan (1990). "Are There Counterexamples to the Closure Principle?" In *Doubling: Contemporary Perspectives on Skepticism*, ed. M. Roth and G. Ross. Kluwer, Dordrecht.

Williamson, Timothy (1994). *Vagueness*. Routledge, London.