

Memory and Knowledge

Tom Stoneham,
Department of Philosophy,
University of York,
Heslington,
York YO10 4JW
U.K.

+44-1904-433258

twcs1@york.ac.uk

Text: 7,180 words

Total: 7,821 words

Memory and Knowledge

ABSTRACT It is argued that having an accurate memory of something which one once knew is not sufficient for now knowing it.

In Plato's *Theaetetus* we find the following exchange (163d5-8):

SOCRATES: ... what I mean to ask is whether someone who has come to know something and remembers it might not know it.

THEAETETUS: How could he, Socrates? What you are describing would be monstrous.

In honour of the young man who thought its denial monstrous, let us call this Theaetetus' Principle:ⁱ

(TP) If someone knows something at t₁, and remembers it at t₂, he knows it at t₂.

I

On a traditional empiricist theory of memory, (TP) would be false, for on such a theory memories are current experiences which provide evidence for beliefs about the past, and clearly it is possible for such evidence to be outweighed by other considerations. But that theory of memory is wrong. Sometimes what is remembered provides evidence for beliefs about the past, as when the detective remembers that the poker was on the left of the fireplace and infers that someone left-handed had been tending the fire, but this is not the normal case. Furthermore, if memories were only ever to serve as evidence for beliefs about the past, the content of the memory could not be that such-and-such happened, for if one remembered

that such-and-such happened, the memory itself would be a belief about the past.

Yet even memories of one's own experiences are memories that such-and-such is or was the case, and many memories are not memories of experiences but memories of facts. For example, some of my memories of a childhood sweetshop are certainly memories of delighted anticipation as I chose the sweets on which to spend my pocket money, but other memories of the same event include the layout of the village, the location of the shop, the name of the shopkeeper, all of which are memories of facts about the world, and not of experiences.

This point is implicit in both Socrates' and my use of the variable 'it' in stating (TP): the very same thing is known at t1, remembered at t2 and known at t2, so whatever it is that is remembered, it must also be something that can be known. For simplicity, let us assume that the content of an episodic or experiential memory, that is the memory of something happening from the inside, can be paraphrased as remembering that I experienced such-and-such. Thus, when I remember the first time I watched the sunset from a mountain top while alone, I remember that the sunset looked a certain way, that I was totally alone and that I experienced it in a certain way. Having a memory with this content may not be sufficient for having the experiential memory,ⁱⁱ but if it is not, the difference does not lie in there being any more to the content of the memory. Intuitively, the difference between remembering that all these things happened, including remembering that I had certain experiences, and remembering the experiences themselves lies in the presentation of the information: the episodic memory involves traces or images which allow me to partly relive the experience. But in evaluating (TP), we need only concern ourselves with the content of memory which can also be the content of knowledge.

We have seen how (TP) is false on the (false) empiricist theory of memory,

and we should also note that on one rival theory it looks trivially true. This rival theory is that memory is nothing more than the preservation of experiences, beliefs and knowledge. On this view we should not think of the operation of memory as some sort of causal process with an input and an output, rather what we call remembering something is nothing more than having known or believed it at an earlier time and neither having changed my mind nor forgotten.ⁱⁱⁱ When, for example, I try to recollect my phone number, I am not presented with some datum, a memory experience, which I can either accept or reject. What happens is that I try to say what my phone number is, and if I get it right, I can be said to have remembered, and if I get it wrong, I have not remembered. There are two ways I might fail to remember, for I might confidently come up with an incorrect number, in which case I have misremembered, or I may simply not be able to offer any number with confidence, in which case I have forgotten though not misremembered. Whichever happens, I may succeed at a later time. Call this the purely preservative view of memory.

If we were to think of memory as some sort of storage process, it might seem puzzling how I might forget or misremember something at one time, but remember correctly later: if I later remember correctly, then the information was sitting there in my 'memory bank' all along, so why could I not access it or why did I distort it? On the purely preservative account of memory this question cannot be formulated, because my getting it wrong or getting it right is not based upon a memory, rather it is what my remembering or misremembering consists in. Memory is not something which provides a basis for knowledge; if I remember something, then I know it, and if we want to ask how I now know it, we should look to how I knew it in the first place. On this account, then, (TP) is trivially true, for if I do not now know it, I cannot have

remembered it.

This strikes me as incorrect, for (TP) sometimes gets the location of epistemic blame wrong. On the purely preservative view, if there is something I once knew which I do not now know, that must be a failure of memory. This obscures the possibility that my memory is working fine, but there is some other explanation for my loss of knowledge. Take, for example, Fred whose memory is working fine, but he believes it is not. Perhaps he has a memory that he met the Queen as a small child, but now find this so unlikely, he assumes that it is some kind of childhood fantasy which has taken on the guise of memory over time. Fred did meet the Queen, and he knew it at the time. His memory is working fine with respect to this information, but now he does not know it, in fact he does not even believe it. It is a fine point of English usage, and not one of great importance, whether we can say that he remembers that he met the Queen, but it should be quite clear that his memory is not to blame for his loss of knowledge.

If one held the purely preservative view of memory, what one would have to say about Fred's case is that he does remember, and thus know, that he met the Queen as a child, but, since he mistakenly believes his memory to be playing him false, he also believes that he did not meet the Queen. Though he does have the knowledge, his doubts about the reliability of his memory lead him to fail to recognize that he does. The case is similar to that of someone who thinks he cannot remember his phone number, but if asked to make a guess, gets it right. What we say about such cases is that he did remember, he did know, but he did not know that he knew. Since Fred believes that he did not meet the Queen, the proponent of this view is committed to denying that knowledge, or at least that remembering, entails belief. I propose to allow them that for the purposes of argument.

Consider, instead, Sue, who also has an accurate memory of the fact that she met the Queen as a child. Sue does not doubt her memory, she does not think that she has somehow invented this encounter, but she does doubt that she really met the Queen. What Sue thinks has happened is that her mother was a fantasist and that her mother told her at some impressionable age that she, Sue, had met the Queen. At the time Sue believed it and, having a good memory, she now 'remembers' it. But now knowing that her mother was both a fantasist and a Royalist, she concludes that her original belief, the source of her memory, was not justified. Sue is here showing an appropriate sensitivity to the following necessary condition on memory-knowledge (knowledge based on memory):

(PK) If someone is to have memory-knowledge, then she must have previously known what she now remembers (and that prior knowledge must be the source of the memory).

Sue's case is meant to show that if someone doubts that she previously knew what she now appears to remember, then that undermines her memory-knowledge.

Again, the proponent of the purely preservative view can insist that Sue does remember and thus does know that she met the Queen, it is just that she does not (now) believe that she did. If we wanted to block that move, we would need to advert to a stronger necessary condition on memory-knowledge:

(KPK) If someone is to have memory knowledge, then he must now be justified in thinking that he previously knew what he now remembers.^{iv}

If (KPK) is correct, then Theaetetus' Principle is false, as is the purely preservative view of remembering. Sue clearly fails to meet the condition set out by (KPK), so if it is right, she has lost her knowledge without her memory being to blame: she knew something, she remembers it, but she does not now know it. The rest of this essay shall be concerned with defending (KPK). Though I shall use examples, my claim is not that there is some intuition about memory or knowledge which (KP) cannot explain but which (KPK) can explain. Rather, I propose to argue that in some cases, such as perception and testimony, we do have good reasons to deny that it is a necessary condition on knowledge that one not merely satisfy the other necessary conditions but also know that one satisfies them. However, there is a disanalogy between these cases and knowledge by memory, which suggests that (KPK) is true.

II

The most obvious objection to (KPK) is that the move from the uncontroversial claim (PK) to (KPK) looks like just another example of the scepticism-inducing error of insisting that one not merely meet but also know one has met all the necessary conditions for knowledge. For example, it is one thing to claim that if one is dreaming one cannot gain perceptual knowledge, another to say that if one does not first know that one is not dreaming, one cannot gain perceptual knowledge.

Descartes was right in saying that if I am to know that I am sitting at my desk listening to the Weaver birds and writing about knowledge, then I must be awake, but many have questioned the further claim that I must independently be able to verify that I am awake and not dreaming. In effect, the sceptical hypothesis that I am dreaming is rejected on the ground that I know I am sitting at my desk listening to the Weaver birds.^v If Descartes' move from the uncontroversial necessary condition that

I be awake to the condition that I *know* I am awake is illegitimate, then so surely is the move from (PK) to (KPK).

In discussions of knowledge of the past, this scepticism avoiding move is often made about the necessary condition that one's memory be reliable. If one had to ascertain that one's memory was reliable, one would have to compare what one appears to remember with independent knowledge of the past. The independent knowledge would either have to be based on memory, leading directly to a vicious regress, or on some other record such as a diary. The reliability of a record such as a diary can only be ascertained by comparing it with one's memory, so the regress is not avoided. The choice seems clear: either knowledge of the past is impossible, or one is epistemically entitled to rely upon one's memory without independently ascertaining that it is reliable. David Wiggins expresses the view very clearly:

it is the epistemological role of experiential remembering – as it is of direct perception, and as it is of one's memory of what one already believes (each of these three things need the support of the other) – to help provide us with a starting point for any further inquiry about how things are in the world. This need not be a philosophically indubitable or infallible starting point. [...] What it must be is a place I can start straight out from without first making an inference from something else, which would itself have to deploy other materials which would themselves have to come from ... Well, where would these come from?^{vi}

While it is contingent that my memory is reliable, I must begin by trusting it, even if I may later come to doubt its reliability. Wiggins is, in this passage, not being

particularly careful to distinguish the psychological issue from the epistemological one, and someone might object that while it is true that psychologically we need a starting point, epistemologically there can be no foundations. But this objection misses the point that epistemology is not only about the justification of static sets of beliefs, but also about evaluating different ways of extending those beliefs. Wiggins' point is that it is perfectly reasonable to believe what we perceive and remember without first assessing the reliability of those sources. He is not saying that in the more abstract project of trying to give final justifications we may simply set aside the question of the reliability of perception and memory. Rather, the thought is that we are rational to trust perception and memory before we have any evidence as to their reliability, and when we do have such evidence, we may discover that the initial trust was misplaced. The entitlement here has a default status but is also defeasible.

It should be easy to see how this idea could be applied to (PK). It is contingent that my previous beliefs were justified, but if I am to be able to use my memories as a starting point for further inquiries, I cannot be expected first to establish that I had been previously justified. For how else could I do that than by a memory of what my prior justification had been, and this begins a regress that will never end. Of course, in a particular case, such as Sue's memory of meeting the Queen, one may question whether this assumption holds, but if memory is to play a role in the current task of extending and developing one's knowledge, we must in general be entitled to rely upon our prior justifiedness^{vii} without independent evidence. An analogy with testimony is very tempting, for memory can be seen as the testimony of past selves. And it seems clear that if testimony is ever to be a source of knowledge in its own right, we must be entitled to assume by default that, in general, our interlocutors are justified in holding the beliefs they share with us.

Similarly, it would seem, we must be entitled to assume that our past selves were justified in holding the beliefs we now remember. So while (PK) is true, (KPK) is not, and (TP) is shown to be true with only minor qualification:

(TP') If someone knew that p at t1, and has an accurate memory that p at t2 (and has no reason to doubt either the accuracy of her memory or the justifiedness of the prior belief), then she knows that p at t2.

III

There are two problems with using this strategy to deny (KPK). The first is a disanalogy between memory and the paradigm case of testimony, where someone, be it parent, teacher or friend, tells me something. Justification is relative to information available and information available changes, and accumulates, over time. For example, we usually accept our doctor's testimony about what is good for our health, but we recognize that the state of medical knowledge changes. So ten years ago your doctor may have told you that all alcohol was bad, whereas today he might tell you that moderate consumption of red wine is good for the health. If we have no other considerations to go on, we should always choose the current testimony over the past, because we think it is based on a larger evidential base.^{viii} But memories are always of things believed in the past, so remembering is much like only having out of date medical information. Of course, if you are like me in only having access to medical information through the testimony of doctors, out of date testimony is better than none at all, but that is not the case with respect to memory. We need to focus our attention away from cases in which we have in fact acquired new evidence which has led to a change of mind, to cases where our only grounds

for current belief in the proposition is our memory. In the case of testimony, someone who defended the equivalent of (KPK) would be arguing that we should not believe what we are told unless we were justified in thinking the testifier knows. But knowing nothing about the methodology of medical research, I will be in no position to judge this of my doctor's opinions on medical matters. The problem is that even if I were to find out what the evidence was on which he based the judgement, I might not be in a good position to judge whether he was justified in making that judgement.^{ix} The situation is very different with respect to memory, for if I am to have memory knowledge, I must have once known what I remember. That may have been knowledge by testimony, but still it was my knowledge.

It seems that there are two questions I am always able to ask about my memories, namely 'How did I come to know that?' and 'Do I (now) find that plausible?'. Because memories are like testimonies of one's past self, as opposed to testimony from someone who may be in a much better epistemic situation than oneself, it seems that these questions are always pertinent. If I know nothing about the methodology of medical research, it is fatuous for me to wonder how it is known that drinking wine is good for the health, but I can still take my doctor's word for it, and if I have no medical knowledge myself, whether I find it plausible or not is irrelevant. In contrast, what I remember are things which I ought to be able to know in my own right.

Another, simpler way of putting the point is as follows. Perhaps it is a reasonable assumption that my belief forming practices have, by and large, been rational and thus that my past beliefs were, by and large, justified. It is a completely different, and much less reasonable, assumption that those justifications would stand up now given all the other things I know. For example, Sue comes to realize that her

mother is a fantasist, and this has the immediate effect that many of her beliefs which appeared fully justified at the time, can no longer be regarded as knowledge. Justifications can be undermined by the acquisition of new information, and so too can knowledge be lost. Given this, it is far from obvious that it is a good starting point in epistemology to assume not only that my past beliefs were justified, but also that they would still be justified from my current perspective. Similarly, I should not assume that I really did know what I then took myself to know. Someone who accepted (PK) but denied (KPK) would be making just that assumption.

The second problem with the envisaged denial of (KPK) is, I think, more fundamental, but a lot less straightforward. The basic point is that having a default entitlement to take our memories as reliable is in part explained by the fact that the reliability of an individual memory is not something within our control. When I have an apparent memory that p, either my memory is accurate or it is not, and there is nothing I can do to change that. In contrast, it seems that whether I have had at some point a justification for believing that p is something squarely within my control.

To see this we need to start with a general distinction between two types of epistemic failure. The first is when the conditions for acquiring the information are not in place, such as when one is under extreme duress, or partially sighted (temporarily or permanently) or one has a poor memory, either in general or for specific sorts of information such as numbers or names. Failings resulting from these conditions do not usually attract blame: if someone did not notice the safety catch was off because they were being forced to do their job too quickly, we do not blame them for the ensuing accident. Similarly, if someone has a very poor memory, we do not blame them for forgetting our name after being introduced. (It is a different question whether someone with a poor memory should be blamed if they

fail to realize that they have a poor memory and take suitable steps to overcome that problem in important cases. Thus, if I have a terrible memory for names, I ought to recognize this and write down the names of the spouses and children of my colleagues.) The second sort of epistemic failure is culpable, for it is when the subject ignores evidence, takes insufficient care or has inadequate epistemic standards. If the victim fails to pick out his assailant in an identity parade because he has a poor memory and does not recognize the culprit, that is not culpable. However, if the victim has no trouble picking out his assailant, but insists that this could not be the right person, despite the similarity in looks, because this person has polished shoes and people with polished shoes do not commit muggings, then we *can* blame the victim for his mistake: he has seriously faulty epistemic standards.

This difference in culpability, and the foundational assumption that our memories are in general reliable have a common element in their explanations, namely that the accuracy of our memory in any particular case is outside our control, it is not something we can in any way affect. Of course, there are strategies one can employ to improve the accuracy of one's memory, but what these do is to improve the percentage of accurate memories one has. In any given case, when confronted with the need to recall, there is nothing one can do to increase the chance of recollecting accurately. There are many heuristics we employ, such as taking our time, concentrating hard and working through other related memories, but it is entirely contingent whether these will help or not. If I have lost my keys, it may help to recall all the things I did since I last had them, but that strategy may just draw a blank, and instead, hours later, while I am thinking of something completely different, it suddenly comes to me where I left them.

The sense in which the accuracy of our memories is beyond our control is

that, in any given case, we cannot do anything to ensure our recollection is accurate. The content of our memory, what, if anything, we recall, is a datum in our reasoning. We can, if we need to, set about trying to establish whether the memory is accurate, but we cannot set about changing its content into a more accurate one. Similarly with the content of perceptions. It is this lack of control which is essential in explaining why the reliability of memory and perception can be epistemically foundational.

If one doubts how the givenness of memory (and perceptual and testimonial) content can help to explain the foundational assumption, consider the case of dreams. As a matter of fact, most of us do not take the information delivered by dreams to be reliable, but that is because (i) we have experience of their unreliability, and (ii) we have some rough theories of how dreams arise which imply that they will be unreliable. But in several human cultures the information delivered in dreams is taken to be reliable. The question we need to consider is whether, in the absence of evidence to the contrary, it is epistemically reasonable to take dreams to be reliable sources of information. And surely those who believe in a dream world have not gone wrong in their first step, namely that of taking the dreams seriously, but in the way they try to reconcile conflicts between dreams and other sources of information. Dreams are representational in more than the merely semantic sense of having a content, for they present events as actually occurring. Thus they are representational in the stative sense and it is possible to take them at face value. Furthermore, their content is not under our control, it is given, and this seems sufficient to give them the default status of being trustworthy. We are prone to overlook this phenomenon because we also have reasons which defeat this default status in every case, but absent those reasons and we would not be irrational to take

dreams at face value, to take them as telling us how things stand in some part of reality.^x

The matter is very different with respect to the question of whether we are justified in believing a given memory, for the epistemic effectiveness of memory depends upon more than just reliability: the 'garbage in – garbage out' rule also applies to memory. If we are not so justified, then we cannot have memory knowledge.^{xi} Suppose I have a memory that the population of some city, say Nairobi, is now more than 1 million. If I have absolutely no idea how I might have acquired this information, if I have no interest in Kenya or in human geography, then it would be epistemically irresponsible of me to take myself as being justified in believing it. If I acted on it, and that action went awry, I would be culpable for acting without proper justification. The problem is not that I doubt that I ever had good justification for believing it, for it is not such an esoteric fact that I could not have picked it up from a newspaper or other reliable source. Rather, the problem is that I lack any positive reason for thinking that I was once justified, for ruling out that the source of this information is my own invention or a late night conversation with an unreliable stranger. I have no idea whether I was once justified in believing this, so how could I be now justified in believing it on the basis of a memory alone? If I am not now justified in believing it, I do not now know it: (TP) fails. Furthermore, it seems that in this case I lack memory knowledge because I fail to satisfy the condition set out in (KPK).

Whether an accurate memory provides me with knowledge is something which is, to a certain extent, within my control. For any particular thing that I remember, there is something I can do to improve the chances of being justified in believing it: what I can do is check that I was previously in a position to be justified in

believing it, that is, check that I could have known it. This is not the requirement that I recover the original justification, but the weaker requirement that I have a reason to think that there was an original justification which would still hold.

We can see this condition operating in the challenge to justify one's memory based assertions. Suppose as part of a game of Trivial Pursuit I answer an obscure question such as 'Who was the first Norman Archbishop of Canterbury?', my answer (Lanfranc) conflicts with the answer on the card and I insist that the card is wrong. Other players might wonder how I know it was Lanfranc, to which the reply 'I remember it' would be woefully inadequate and the reply 'I have a reliable memory' would be irrelevant. The answer which would explain how I know it was Lanfranc, is that I once studied that period of history at school. If I cannot come up with any explanation of how I might once have known it, then we would doubt whether I really do know it now.

The argument for (KPK) has run like this. For some of the necessary conditions on memory-knowledge, such as accuracy of memory, it is not necessary that one know those necessary conditions are fulfilled, for one is defeasibly entitled simply to trust that they are. The explanation of this is that we lack a relevant sort of control over the accuracy of our memories, which is evidenced by our non-culpability for failing to remember. However, we are culpable for taking ourselves to have memory-knowledge when we lack evidence of prior justification, and this does seem to be something over which we have partial control. Hence, we cannot treat the necessary condition (PK) like the condition of accuracy, as something which must not be doubted but which needs no independent verification. Focussing on examples such as remembering one's phone number, or one's siblings' names, obscures this point, for in those cases it is obvious how one would have acquired the

knowledge in the first place, so the consequent of (KPK) is easily satisfied. In cases where this condition is harder to satisfy, it does turn out to be necessary for memory-knowledge.

IV

There are two lines of objection to this argument which need to be considered. One is that we are sometimes blamed for memory failures, which undermines the disanalogy on which the argument turns. The other is that, by making it a condition on memory-knowledge that I know some necessary condition for that knowledge has been satisfied, I have reintroduced the vicious regress which leads to scepticism.

The first line of objection is motivated by examples like the following.^{xii} Suppose that Fred is very bad at remembering dates, such as birthdays and anniversaries. What it means to be bad at remembering dates like this (as opposed to remembering the date of the Battle of Agincourt), is for contemplation of that date (day and month) or one close to it, not to serve as a cue for remembering the anniversary. Thus, if your own birthday is the 10th August and someone suggests doing something on the 9th August, you are likely to be cued to remember that it is your birthday the following day. Similarly, if you note, for some reason, that today is the 9th August, that will equally remind you that your birthday is imminent. What is wrong with Fred is that these cues rarely work: he can look at a calendar, notice that the current date is 27th October and not be prompted to remember that it is his mother's birthday, or whatever. If he sets out to recollect when his mother's birthday falls, he may well remember, but that does not help much when he has forgotten to ring her on the day.

By the reasoning I gave above, if Fred fails to remember a birthday, he is not

culpable, he just has a poor memory. However, it is indisputable that if Fred forgets his wedding anniversary, his wife will certainly be aggrieved and hold him to blame. His terrible memory for dates will explain the oversight, but it will not excuse him. It appears that Fred can be blamed for this particular failure of memory, even if he is not to blame for generally having a bad memory for dates.

This is an interesting example because it is the importance of the information not remembered which is the crucial factor in apportioning blame. There are lots of one-off events in one's life, from the first day at school to the first pint of beer, for which it is totally unimportant whether one remembers their anniversary or not, but a wedding is rather different. There are three explanations for this, only one of which threatens my argument. First, we might think that wedding anniversaries are so important that one should not rely on memory alone, especially if one is not certain of having a good memory. For important dates, one is obliged to take extra steps to back-up possibly fallible memory, such as making a note in a diary. Unlike tying a knot in a handkerchief, a diary entry is more than just another cue, for it carries the information to be remembered and thus supersedes the memory. Some philosophers might want to regard diaries as just an extension of memory, but there is an important difference between memory itself and a cognitive tool like a diary, for a tool is deliberately chosen for the purpose and it should be chosen in part for its reliability. Handheld computers can save one from having to remember all sorts of things, and can offer unsolicited reminders at pertinent times, but if someone is to use one as a cognitive tool in place of remembering, they are obliged to not merely choose a reliable model, but to do so knowingly. So we can explain Fred's wife's anger at his forgetting their wedding anniversary to be directed not at his failure of memory itself, but at his failure to supplement his memory in a way which would

ensure important anniversaries do not go unmarked.

The second explanation we can offer is that one's wedding is such an important event that the date should be firmly impressed upon even the weakest memory. The idea here is that, given the way our memory works, trivial or unimportant matters are less likely to be remembered, but very significant events are hard to forget. It is an empirical matter whether our memories work this way, but folk-psychology and some empirical research suggests they do.^{xiii} The implication of Fred's forgetting his wedding anniversary would be that he did not take the original event, the wedding ceremony, to be particularly significant. So the blame is again not focussed on his failure to remember, but instead on his original attitude towards the fact to be remembered. If he had had the right attitude, he would have remembered.

The third explanation is very similar to the second, but it does support the idea that Fred is to blame for the memory failure itself. It might be thought that if some past event is very important to Fred now, whether or not it was important to him at the time, it should be easier for him to remember it. Hence, if he cared enough about his marriage now, he would be able to remember the wedding anniversary. If this were right, there would be something he could do to ensure, or improve, the accuracy of his particular memory, namely to care a great deal about the event that he is trying to remember. I take it that it is an empirical matter whether valuing something highly now makes it easier to remember, but, once distinguished from the first and second explanations, it looks pretty implausible. It is part of the frustration of having a poor memory that even when it is really important to remember, one cannot. This third explanation of why Fred's wife blames him for forgetting their anniversary, even though he has a notoriously poor memory for dates, is the only

one which supports the objection to my argument, but it is also the least plausible.

V

The second objection is that by requiring not merely (PK) but the stronger (KPK), I have reintroduced the vicious regress of justifying memories by reference to further memories. Now I am happy to accept that it cannot be a necessary condition on knowledge by memory that one know one's memories are reliable, because that would make all such knowledge impossible. But it does not follow that (KPK) equally leads to a vicious regress.

Our memories are multi-faceted and there is no reason to think that a single epistemological story should apply to all facets. In particular, I think we should distinguish three distinct types of memory: experiential memory, working memory and factual memory. Experiential memory was discussed earlier (p.3), and distinguished by its content: having an experiential memory of a past event involves, at the very least, remembering having experiences of the event. Factual memory has no such constraints upon its content, and prospective memory of future events and appointments falls into this category. In effect, one can have a factual memory of anything one can know, whereas one can only have an experiential memory of something one has experienced. Working memory is rather different, but is rarely distinguished from factual memory by philosophers. One can generate phenomenologically obvious examples of working memory by asking someone to do moderately hard mental arithmetic, such as adding 384 and 195. This process involves several steps and the original numbers plus the results of each step need to be 'kept in mind' if one is to get the correct answer.^{xiv} And any mental task which takes enough time for one to be distracted, and that is a very short amount of time

indeed, will need to use similar storage of inputs, calculations and goals. Sometimes psychologists characterize working memory by distinguishing short-term from long-term memory, but one can see that if working memory was just very short-term factual memory, there would not be much philosophical mileage in the distinction. The thought here would be that interim conclusions, such as '9+8=17 so I should put a 7 in the 10s column and carry one to the 100s' in the example above, are kept in working memory as long as they are needed, but could in theory stay in memory indefinitely.

This deflationist view of working memory does not mesh well with empirical evidence that the mechanisms underlying short and long-term memory are distinct. But nor does it fit with our everyday understanding of recollection. Consider, for example, this comment of Locke's:

The same Idea, when it again recurs without the Operation of the like Object on the external Sensory, is Remembrance; if it be sought after by the Mind, and with Pain and Endeavour found, and brought again in view, it is Recollection.^{xv}

On Locke's distinction, remembrance occurs when prior knowledge or experience is currently available, whereas recollection is an activity the product of which is remembrance. Clearly recollection is not necessary for remembrance, so there remains the possibility that there are some things which we remember but which we cannot recollect, that is, what determines their current availability is not the activity of recollection. Now it would seem that, on this understanding, though we sometimes do remember the contents of our working memories, there is no sense in trying to

recollect them. For example, if I add a column of figures, there is a determinate way I perform the intermediate steps of the process, but after it is complete and I am no longer thinking of the process, I cannot recollect how I did it. I might be able to reconstruct the steps from knowledge of my favourite arithmetic techniques, or by inference from an experiential memory of looking at particular figures, but I will not recollect the actual process, I will not be able to retrieve the contents of my working memory.

These comments are not intended to be decisive, but to make more plausible the proposal to give different epistemological treatments to experiential, working and factual memory. Now, it does seem that for working memory to perform its essential role in all our reasoning, (KPK) would have to be false. Whatever the epistemological conditions for working memory to give knowledge, they cannot include a requirement to know something which can only be known through a process which involves working memory.^{xvi} So if (KPK) applies to anything, it only applies to experiential memory and factual memory. Does it create a vicious regress here?

Interestingly, making (KPK) a necessary condition for knowledge based on experiential memory creates no regress of any type, for the content of an experiential memory includes information not only about what happened by also about how I know that it happened. Thus the requirement in (KPK) is trivially satisfied. For episodic memories, the question of how I knew what I now claim to know on the basis of memory is adequately answered by saying that I remember it happening, I remember experiencing it.

Non-episodic, factual memory is a different matter altogether. If (KPK) is true of that type of memory, then to have memory-knowledge I must have some grounds

for thinking that I once knew, or was at least justified in believing, what I now remember. This could be very general. For example, the explanation of how I came to know the names of my sisters, or my own telephone number, or that Paris is the capital of France, would advert to very general facts about the society I was brought up in. Contrast that with the explanation of how I came to know, and am thus able to remember, my National Insurance number or the name of the nearest pub to Earl's Court station. Here the explanation would have to advert to quite specific facts about my past, for only they would explain how I came to know these things. Whatever the form of the explanation, it would seem that regress beckons, because, barring exceptional circumstances, I would have to remember the general facts about how children are educated or the specific facts where I used to meet a certain friend.

But it is far from obvious that this is a vicious regress rather than a harmless coherence condition. Standard objections to coherentist approaches to memory-knowledge make assumptions about the scope of the claims which are not true of the current proposal. (KPK) entails that factual memories need to be supported by memory-based evidence of how the original belief might have been justified. As such they need to cohere with each other, but they will also need to cohere with experiential memories, and we have seen that (KPK) does not create a regress there. Nor is a problem created by the process of justifying factual memories by reference to other memories itself relying on working memory, for we have seen that (KPK) is not true of working memory. The moral is that analogies and regress arguments tempt us to apply the purely preservative view of memory to all forms of remembering, but if we are more careful, we can see that they do not in fact rule out (KPK) from applying to experiential and factual memory. And in the case of factual memory, (KPK) imposes a condition for memory-knowledge which is often not met,

despite the well-functioning of our memories.

As a final point, I would like to return to the apparent analogy between memory and testimony. If (KPK) is true, but no such similar condition holds for testimony (p.10 above), then it follows that knowledge by memory is ever so slightly harder to come by than knowledge by testimony. This might seem surprising given the strongly individualist tradition that has permeated epistemology, but it should not be too surprising. After all, testimony, but not memory, is a source of new information. While the ability to hang on to our knowledge is clearly essential for any sort of progress, the ability to gain knowledge must come first. In our retention of knowledge, I have argued that success requires some sensitivity to our limitations and fallibility. In the process of gaining knowledge we are not required to be so sensitive to the limitations of our sources, instead we (defeasibly) trust them, for the simple reason that we cannot set about improving them in any particular case.

REFERENCES

- Baddeley, A. (1999): *Essentials of Human Memory*, (Hove: Psychology Press).
- Blaney, P. (1986): 'Affect and Memory: a Review', *Psychological Bulletin*, 99, 229-46.
- Boghossian, P. 1989: 'Content and Self-Knowledge', *Philosophical Topics*, 17, 5-26
- Burge, T. 1988: 'Individualism and Self-Knowledge', *Journal of Philosophy*, 85, 649-63, reprinted in Cassam (1994).
- Cassam, Q. (ed.) 1994: *Self-Knowledge*, (Oxford: OUP)
- Dancy, J. (ed.) 1988: *Perceptual Knowledge*, (Oxford: OUP)
- Dokic, J. 2001: 'Is Memory Purely Preservative?', in Hoerl & McCormack 2001
- Evans, G. 1982: *The Varieties of Reference*, (Oxford: OUP)
- Hitch, G. (1978). 'The Role of Short-term Working Memory in Mental Arithmetic'.
Cognitive Psychology. 10. 302-23.
- Hoerl, C. & McCormack, T. 2001: *Time and Memory*, (Oxford: OUP)
- Hopkins, J. & Savile, A. (eds.) 1993: *Psycho-analysis, Mind and Art* (Oxford: Blackwell)
- James, W. (1890): *Principles of Psychology*.
- Locke, D. 1977: *Memory*, (London: Macmillan)
- Locke, J. 1690. *An Essay Concerning Human Understanding*.
- Plato 1973: *Theaetetus*, (Oxford: OUP), translated by McDowell, J.
- Quinton, A. (1962): 'Spaces and Times', *Philosophy*, ,130-147.
- Wiggins, D. 1993: 'Remembering Directly', in Hopkins & Savile (1993)

NOTES:

ⁱA similar claim is taken as equally obvious by Paul Boghossian (1989, p.23).

ⁱⁱFor an argument that the content of experiential memory involves reference to my experiences, see James (1890, vol.i, p.650). For an argument that the content of episodic memories must go beyond this, see Dokic (2001). I am not convinced by Dokic, but the point does not matter for the present paper.

ⁱⁱⁱe.g. Locke (1971, p.95-6).

^{iv}In response to an earlier version of this paper, Philip Percival wondered whether there was a consistent notion of pure memory-knowledge in (PK) and (KPK). The problem is that (KPK) seems to say that memory-knowledge rests on an independent justifying element, and is thus not pure, but once we allow that, (PK) does not look so obvious any more. The notion of pure memory-knowledge I am working with is that the relevant necessary condition memory-knowledge that p is that one have no current evidence that p. (KPK) does not violate this, because it requires one to have evidence that one was justified in believing that p, which one can have without having evidence that p. For example, if I remember that X's phone number is nnnnnnnn, (KPK) is satisfied by my justified belief that I used to phone X often, which provides no evidence about what X's phone number is, only that I once knew it.

^vThe view is now quite common in the epistemology of perception. See Burge (1988, p.70) for a succinct statement and Dancy (1988) for a collection of papers on the topic.

^{vi}Wiggins (1993, p.348). See also Evans (1982, ch.7.5)

^{vii}Justifiedness is the property of being justified, which is distinct from being able to provide a justification.

^{viii}Of course, many testifiers may be conveying beliefs based on very out of date evidence, but the principle still holds because, even if a testifier at time t2 does not take into account all the evidence available at t2, a testifier at t1 cannot take into account evidence only available after t1.

^{ix}One example: the easy availability of medical data on the internet leads many not trained to make clinical judgements to question their doctor's opinions. I have heard this called 'cyberchondria', and it is clearly a form of epistemic irrationality.

^xAnthony Quinton (1962) once used dreams to argue for the possibility of an individual experiencing two unrelated spaces. What is striking about his discussion is that he sees the task of making it reasonable to believe the events of the dream really happened as a task of removing the normal conflicts with waking experience.

^{xi}Given the distinction between justifiedness and providing a justification, this claim does not commit me to any particular analysis or theory of knowing.

^{xii}This example, and the objection, were suggested by Ward Jones.

^{xiii}See Blaney (1986), which according to Baddeley (1999, p.186), is still current.

^{xiv}A seminal experiment which looks at the role of memory in such tasks is Hitch (1978).

^{xv}Locke (1690), II.xix.1. Interestingly, this is quoted by the OED in the entry on

recollection.

^{xvi} For a more detailed account of this argument, see my 'Memory and Inference' (MS).