

Googled Assertion

J. Adam Carter

adam.carter@glasgow.ac.uk

(University of Glasgow)

Emma C. Gordon

emma.gordon@ed.ac.uk

(University of Edinburgh)

Abstract: Recent work in the philosophy of mind and cognitive science (e.g., Clark and Chalmers 1998; Clark 2010a; Clark 2010b; Palermos 2014) can help to explain why certain kinds of assertions—made on the basis of information stored in our gadgets rather than in biological memory—are properly criticisable in light of misleading implicatures, while others are not.

Keywords: Assertion; extended cognition; extended mind; knowledge

§1.

John is at a dinner party and wants to appear knowledgeable and cultured, in order to impress the others at the table. The after-dinner conversation turns to Mozart, a topic with which all (except John) are already familiar. One individual at the table insists that of all of Mozart's piano Sonatas, only Piano Sonata No. 14 is in a minor key (i.e., C minor). John knows nothing about Mozart, but sees his chance.

He discreetly Googles [Mozart + sonata + minor key] from his iPhone under the table, locates the Wikipedia page for Mozart, and casually asserts, "Actually, when Mozart was in Paris in the summer of 1778, he wrote what would become Piano Sonata No. 8—which is, interestingly, the only other concerto Mozart wrote in a minor key." Just as John hoped, everyone is impressed.

There is obviously *something* criticisable about John's assertion. It would be a mistake, though to say that what's criticisable is that John didn't *really* know what he asserted. John knew that what he said was true. After all, Wikipedia is very reliable source. If you are not convinced, just switch the details of the case and suppose that John consulted (under the table, with his iPhone) a peer-reviewed music journal.

What's criticisable about John's assertion is at *least* this: that John's 'Googled assertion' was a kind of *bluff*. By asserting what he did, John implied that he knew much *more* than he did, in this case, about Mozart. Bluffs, in poker and elsewhere, can be called. In order to call John's bluff, all anyone needed to do was to ask John a quick follow-up question. John's iPhone would have the answer to the follow-up question, but *John* (independent of the iPhone) would not.

§2.

The case of John can be used in the service of making two very different kinds of philosophical points about assertion. The point we will *not* try to make is one of a style that has been employed in recent work by Jennifer Lackey (2011). Consider that if John knows that *p*, and yet, John's assertion appears to be *epistemically* criticisable in light of the implicature that John's epistemic position is better than it is, then this looks very much like evidence against the *sufficiency* leg of the *knowledge norm of assertion* (e.g. Unger 1975; Williamson 2000; DeRose 2002; Hawthorne 2004; Simion 2015)—according to which, knowledge is a sufficient epistemic credential to warrant assertion.

However, as recent work by Matthew Benton (2014) has shown, the strategy of moving from misleading implicatures¹ about an asserter's epistemic position relative to the asserted proposition, to the conclusion that

¹ See Grice (1991).

the sufficiency thesis of the knowledge norm of assertion is false, is far from a straightforward road². It's not a road we'll try to take here³.

Without weighing in on the issue of whether cases like John's are evidence against the knowledge norm, we will simply take John's epistemic bluff at face value in order to make a very different point altogether. In what follows, what we want to suggest is that to the extent that John's assertion is criticisable, it is so in virtue of details about the *way* in which John asserted on the basis of Googling the fact about Mozart in question. And this has to do with complexities about the relationship between the iPhone and John's *cognitive agency*. If the argument we advance is right, then the epistemic criticisability of 'Googled' assertions like John's varies dramatically, depending on how it is that the technology relied upon is *integrated* into the asserter's cognitive architecture. More generally, we will show how recent thinking in the philosophy of mind and cognitive science has some surprising bearing on our assessments of the propriety of assertions.

§3.

Consider now the following pair of cases, due to Andy Clark and David Chalmers (1998), which form a familiar reference point in the contemporary literature on *extended cognition* in the philosophy of mind.

Inga: Inga has a normally functioning biological memory. When she learns new information, she stores it in her biological memory (as is typical and familiar) and it is her biological memory which she consults when she needs old information.

² Arguments to this effect would also have to establish that the impropriety is epistemic rather than merely a failure to satisfy a Gricean maxim (e.g., Goldberg 2013). Cf., Benton (2014a).

³ See Carter and Gordon (2011, forthcoming) and McGlynn (2014) for critical discussions of the sufficiency leg of the knowledge norm of assertion.

Otto: Otto suffers from Alzheimer’s disease, and like many Alzheimer’s patients, he relies on information in the environment to help structure his life. Otto carries a notebook around with him everywhere he goes. When he learns new information, he writes it down. When he needs some old information, he looks it up. For Otto, his notebook plays the role usually played by a biological memory (1998, 8).

First, a similarity, and then a dissimilarity. The most striking similarity between Otto and Inga is a *functional* one; both rely on something (a notebook and a biological brain, respectively) to play the functional role of information storage and retrieval. But there appears to be an important difference. The elephant in the room is that what’s playing this functional role for Otto is something in the world—*outside of his head*—a paper notebook, with which he is physically interacting. The traditional way of thinking about the case (e.g., Adams and Aizawa 2001; 2008; 2010; Rupert 2004) is familiar: that Inga in the former case is relying on her *memory*, but Otto in the latter is not⁴, even though the notebook plays a memory-like role for him. The notebook, after all, is not a constituent of Otto’s mental life.

According to proponents of extended cognition, however, this traditional diagnosis rests on ‘bioprejudice’. They suggest that, rather than to focus on material constitution or physical location, we should let the following ‘parity’ principle guide our judgments of what counts as a part of an individual’s cognitive process⁵:

⁴ The primary argument which Adams and Aizawa have leveled against extended cognition is what they call the coupling constitution fallacy—viz., the alleged fallacy of moving from causal coupling of some object to a cognitive agent, to the conclusion that that object or process is part of the agent’s cognitive processing. Cf., Clark (2010) for a reply.

⁵ As Clark (2007) puts it, ‘the actual local operations that realize certain forms of human

Parity Principle: If, as we confront some task, a part of the world functions as a process which, were it to go on in the head, we would have no hesitation in accepting as part of the cognitive process, then that part of the world is part of the cognitive process.

With reference to the parity principle, it looks as though—given that Otto’s notebook is playing the same functional role *vis-à-vis* information storage and retrieval as is Inga’s biological brain—we should include Otto’s notebook as part of his memory provided we include Inga’s biological memory as part of *her* memory (which, of course, we do).

A familiar objection to this more egalitarian approach to the bounds of cognition is that too much is let in. As Sean Allen-Hermanson (2012) puts the worry: ‘If a notebook counts as part of one’s mind, then why not the yellow pages, the internet, or even parts of the natural world that supply information and support cognition?’ (2012, 792) In response to this kind of ‘slippery slope’ worry—the objection from *cognitive bloat*—Clark has put forward some additional integration conditions, which must be satisfied by external artifacts to be included within a cognitive process:

Clark’s “Trust and Glue” Integration Conditions

- (1) “That the resource be reliably available and typically invoked.”
- (2) “That any information thus retrieved be more-or-less automatically endorsed. It should not usually be subject to critical scrutiny. [...] It

cognizing include inextricable tangles of feedback, feed-forward and feed-around loops: loops that promiscuously criss-cross the boundaries of brain, body and world’ (Clark 2007, §2) According to some more recent defences of extended cognition, especially due to Palermos (2011; 2014), the position can be motivated entirely with reference to dynamical systems theory (Beer 1995; 2000), without reference to the more standard functionalist argument from the parity principle.

should be deemed about as trustworthy as something retrieved clearly from biological memory.”

(3) “That information contained in the resource should be easily accessible as and when required.”

While Otto’s notebook arguably satisfies these conditions, resources such as an almanac in one’s study which is only occasionally consulted, do not, and so—as this line of defence goes—the parity based rationale according to which we count Otto’s notebook as part of his memory needn’t also rule-in the almanac as part of your memory, simply because you’ve consulted it on occasion. Otto’s notebook satisfies cognitive integration conditions which the almanac, in cases of ordinary use, does not⁶.

Unsurprisingly, the hypothesis of extended cognition, is a controversial one⁷. It is beyond our present scope to argue conclusively for the view. Rather, we submit that the view is one which is gaining traction not only in the philosophy of mind, but also more recently in epistemology⁸, and further, that this view has some interesting applications (as we’ll see) in debates about assertion.

§4

Let’s return to the dinner party, with John. Suppose that Otto, from Clark and Chalmers’ case, is present. Otto, coincidentally, has been watching documentaries on Mozart and has jotted down pages of facts in his notebook,

⁶ For a detailed discussion of cognitive integration conditions, and in particular, how cognitive integration matters with respect to extended knowledge-generating cognitive abilities, see Pritchard (2010).

⁷ For a sample of recent essays for and against this thesis, see Menary (2010).

⁸ See, for example, Carter (2013; forthcoming); Carter & Kallestrup (2016); Carter & Palermos (2015); Carter and Pritchard (forthcoming), Carter et al (2014); Clark (2015); Kelp (2013a); (2013b); Palermos (2014); (2015a); (2015b); Pritchard (2010); Vaesen (2011); (2013).

one fact of which he asserts as a natural follow-up to John's remark: 'Ah yes, Mozart left Paris a few months later, in September of that year, and travelled to Munich.' We may suppose further (to retain similarities across the cases⁹) that Otto's consultation of his notebook in this circumstance mirrored the discreetness of John's consultation of his iPhone in the original case.

Let's consider now some similarities between Otto's assertion and John's assertion, before noting some important differences. Regarding similarities: both Otto and John asserted on the basis of consulting *non-biologically* stored information; neither disclosed this fact; both, by asserting, presented themselves to their interlocutors as knowing the asserted fact; and, crucially, given the specialised nature of the informational content of the assertion, each conversationally implied *that* they know further facts about Mozart's life and music.

John's assertion, from §1, was a clear case of epistemically criticisable bluffing—viz., bluffing because John did *not* know further facts about Mozart which he implied he did. What about Otto's assertion—is it a case of epistemically criticisable bluffing, too? It seems, initially at least, that we must say yes. After all, John was bluffing precisely because he presented himself as knowing more than he actually did¹⁰. And, it looks very much as though Otto is equally guilty of generating the false implicature that he knows more about Mozart than he actually does. (After all, neither John nor Otto can answer a

⁹ See §6 for further discussion on this point.

¹⁰ It is worth pointing out that bluffs can potentially be criticised on the basis of different kinds of reasons, some epistemic (which are the kinds of considerations at issue here) and others moral. In the moral case, for instance, it is noteworthy that bluffs typically involve some kind of *intent* to deceive, a feature that can incur moral disapprobation independent from epistemic evaluations of bluffing. Furthermore, it's important to note that even if one asserts *p* with the *intention* of bluffing in the sense at issue in the cases described--viz., with the intention of communicating via implicature that one knows more than one actually does--one might in fact fail to bluff if one actually knows more than one takes oneself to know. Thanks to an anonymous referee at *Philosophical Psychology* for requesting clarification on this point.

single follow-up via appealing to biological memory). And yet, this diagnosis is not very satisfying.

To bring this point into sharp relief, suppose for instance that following each assertion, the relevant background facts were brought to light to all the interlocutors—viz., including facts about the roles that John’s iPhone and Otto’s notebook play in John’s and Otto’s processes of information storage and retrieval more generally, their history of interaction with their devices, etc. It seems as though the more facts are brought to light, the stronger the corresponding intuition that John’s assertion is criticisable in a way that Otto’s assertion is not, and that this is so *despite* the similarities between the two assertions noted above.

However, this intuition stands in tension with the *explanation* for why John’s epistemic bluff appeared to be criticisable: John’s assertion generated the *misleading implicature that he knew more about Mozart than he actually did*¹¹. And, yet, even bringing all relevant facts to light, it looks as though Otto—no *less than John*—is guilty of generating the implicature that he (Otto) knew more than he did. Just as, if we asked John a follow-up question, he’d need to get out is iPhone, so if we asked Otto a follow-up question, he’d have to

¹¹ According to Grice (e.g., 1975, 49-50) *q* is conversationally implied by a speaker’s utterance of *p* provided (i) the speaker is presumed to be a cooperative speaker; that (ii) in order to maintain this assumption it must be presumed that the speaker believes that *q*, and that (iii) the speaker takes it that both speaker and hearer mutually know that the hearer can work out that to preserve the assumption that the speaker is cooperative, *q* is in fact required (cf., for discussion, Levinson 1983, 113 and Hirschberg 1985, §2). Further, any speaker is being cooperative (as per (i)) only if respecting the Maxim of Quality, according to which one must not say false things, or things for which one lacks good evidence. Misleading implicatures of the sort given by John violate this maxim, for the assumption that John is a cooperative speaker is maintained only if it is also presumed that John really does know other items on the topic on which he asserted a specific piece of ‘quasi-expert’ information, and this is something that both John and his audience plausibly know that the audience can work out. Thanks to a referee at *Philosophical Psychology* for requesting clarification on this point.

get out his notebook.

§5.

The hypothesis of extended cognition offers a neat way to resolve the tension noted in the previous section. The diagnosis, in short, goes as follows: while both John and Otto, by asserting as they did, generate the implicature that each knows quite a bit more about Mozart, John's implicature is false and Otto's is true. That is why John's assertion is criticisable in a way Otto's is not. And this is so even though *neither* can tell you anything further about Mozart without consulting an external artefact. Otto has additional *extended knowledge*, but John does not.

What makes the difference between generating a false implicature and a true implicature, is whether each *does* know other things about Mozart. And, given that each has the information externally stored, the crux of the matter is that only if either possesses additional *extended* knowledge, information stored outside the biological brain, about Mozart does either generate a *true* implicature that either knows further information on the topic. And it is with reference to the hypothesis of extended cognition that we have a principled reason to think that Otto does plausibly have such knowledge, and John does not. And *this* is because, on the extended cognition programme, of an important *difference in cognitive integration* in the two cases.

To appreciate this point, it will be helpful to compare both John and Otto with Lilah, who is also at the dinner party:

Lilah: Suppose Lilah habitually and automatically stores all of her appointments and deadlines onto her Google calendar, which is very readily accessible on an external device—a smart watch—that she takes everywhere she goes. While John and Otto are talking about Mozart, another of the dinner party guests asks Lilah when she'll be giving her

talk at next week’s workshop. She doesn’t have this information—or for that matter any other information about her appointments—stored in biological memory, and so discreetly glances at her smart watch and says ‘Wednesday at 3:30 p.m., on the first floor, room 1.17’.

Lilah isn’t that different from many of us. After all, people nowadays *rarely* rely primarily on their biological memories¹² for information storage and retrieval of this nature¹³. Furthermore, Lilah produces very specific information about her lecture location, and almost immediately, in a way that is suggestive of a general competence with this kind of information.

Of course, Lilah has none of this information in her head. (Asked a follow-up, Lilah would no less than John or Otto need to consult her gadget, without which she would be hopeless.) Lilah is intuitively, however, *not* criticisable, in light of generating a misleading implicature that she knows more than she does; she is, intuitively, *not* bluffing. And this is so even though her assertion is suggestive of her knowing the details of her schedule, details that are simply not in her head.

Extended cognition can explain all of this: Otto’s assertion and Lilah’s assertion are not criticisable—even though John’s is—because Otto and Lilah *do* know further things about the topic in question. This is so even though the information is not in their heads but in, with reference to extended cognition, their extended memory. With reference to Clark’s ‘glue and trust’ integration

¹² For work that illustrates the types of memory limitations that offloading can help us overcome, see e.g. Simons and Levin (1997) and Cowan (2010).

¹³ See Lynch (2014) for an overview of the cognitive offloading trend, and see e.g. Risko (2014) and Chu & Kita (2011) for evidence that cognitive offloading improves performance not just with respect to memory but also perception and spatial reasoning. Meanwhile, see Risko and Gilbert (2016) for an exploration of the social and cognitive consequences of cognitive offloading, and for a review of how this behavior is seemingly triggered by internal cognitive demands associated with a task. Thanks to an anonymous referee at *Philosophical Psychology* for pointing out the relevance of this empirical literature.

conditions, note that Otto's and Lilah's gadgets are cognitively integrated in a way that John's is not, despite some superficial similarities. This point needs some unpacking.

Note, firstly, that in both the case of Otto and Lilah, their interaction with their external devices is automatic and seamless, uncritical, and (as Clark puts it,) 'transparent' as is consulting biological memory. In the case of John, by contrast, what we have is a kind of *deliberate* action—one which takes 'two steps': a plan and an action.

As Clark (2010b) has noted, 'Otto is so used to using the book that he accesses it automatically when bio-memory fails. It is transparent equipment for him just as biological memory is for Inga.' And this is the case for Lilah as well, whose consultation of her smart watch (and inscription in it) is second nature for her.

The point about cognitive integration can be made, likewise, with respect to *dynamical systems theory* (e.g., Beer 1995; 2000; Chemero 2009, Froese et al. 2013, Sutton et al. 2008, Theiner et al. 2010, Tollefsen & Dale 2011, Palermos 2014b), according to which, in order to claim that two (or more) systems give rise to an overall extended, or coupled, system, this requires non-linear relations that arise out of *mutual interactions* between the contributing parts¹⁴. As Palermos (2011; 2014a; 2014b) has noted, on the basis of dynamical systems theory, we are in a position to claim that, in order to have an extended cognitive system (i.e., Otto and his notebook; Lilah and her calendar), all we need is that the internal and external parts interact mutually with each other, and the existence of such processes of mutual interactions is a requirement that goes unsatisfied in the case of John and *his* phone.

¹⁴ For an expanded discussion on this point, see Carter, Gordon and Palermos (2015). The application of dynamical systems theory to the epistemology of extended cognition has been developed primarily by Palermos.

Putting this all together: initially, it seem utterly mysterious why *some* cases of asserting on the basis of information not stored in one's head (but rather, stored in a gadget) are, intuitively, cases of epistemic bluffing and others aren't. John, Otto and Lilah all make assertions that generate the implicature that they possess further background knowledge which none of them actually has in their heads, but rather on external devices. Yet, John's assertion is intuitively criticisable in a way that Otto's and Lilah's isn't. As we've seen, extended cognition can neatly explain this. In short, if extended cognition is true, John is bluffing, but Otto and Lilah are not. His implicature that he possesses further knowledge on the topic is false, and theirs is true.

§6 Objections and replies

Objection: What work is the inconspicuousness of consulting the external artifacts doing, exactly, in these cases? More specifically: why is it built in to the cases discussed that John, Otto and Lilah are being discreet in their consultation of the device when asserting on the basis of it? Why not just make these cases where the individuals consult their external storage devices in front of their interlocutors?

Reply: Nondiscreetness would cancel the relevant implicatures. Consider that, in the circumstances, as described, John is clearly bluffing, but Otto and Lilah aren't, and this was a phenomenon that needed explaining (given that none uses their brains for information storage on the topic). However, if (for example), John was entirely transparent about getting the information from his phone, he would no longer be bluffing. This is because his making explicit the source of his assertion would cancel the implicature that he is knowledgeable on the topic of Mozart. The interesting case, then, is the discreet case, one where John *is* bluffing but (with reference to extended

cognition) Otto and Lilah, equally discreet, are not, as they generate an implicature (that they are knowledgeable on the respective topics) that, unlike John's, is true.

Objection: Even if the Clark/Chalmers line is right, and Otto's and Lilah's notebook and phone, respectively, function as a kind of extended memory for each, why should we grant that what they have is extended *knowledge*? This seems to matter, given that their assertions are claimed to generate the implicature that they are knowledgeable about the respective topics. But this implicature could still be false, *even if* they have extended memory, if the status of the information stored in their external devices falls short of knowledge.

Reply: This is right; if Otto and Lilah do not have extended knowledge, then like John, they generate false implicatures. We can imagine tweaks on these cases where they (like John) fail to have extended knowledge, even if the information stored in their external devices is, by the lights of extended cognition, part of their extended memory. For example, if a jokester tampered with Otto's notebook, switching around the dates corresponding with Mozart's Sonata's, then *even if* the jokester accidentally switched the dates back to the right order, the information in Otto's notebook is not extended knowledge; the 'knowledge' is Gettiered¹⁵. Thus, extended memory without extended knowledge is certainly possible. However, *if* we embrace the extended cognition reading of the cases discussed, there is no obvious *epistemic asymmetry* that would justify a difference in epistemic assessment in the extended cases relative to biological counterpart versions of these cases. (As soon as we attempt to point to various possibilities of error, for Otto and Lilah,

¹⁵ See Gettier (1963). For a discussion of Gettier cases for extended cognition, see Carter (2013). Cf. Jarvis (2015).

the proponent of extended cognition can note point out that if the possibilities of error are *greater* than they would be were they to rely on normally functioning biological memory, then these aren't bona fide cases of extended cognition, as the cases would fail to satisfy the parity principle¹⁶).

Objection: Let's suppose that the Clark and Chalmers reading of the Otto case is right, and further that we accept that Otto counts as knowing what is in his notebook, and Lilah the information in her smart watch calendar. What if John interacted with his iPhone, for the purposes of Googling, in a way that was as automatic as the way Otto and Lilah consult their external devices. Imagine, for example, a twist on the original case of John, where we hold fixed all details except tweak (dramatically) the extent to which John's iPhone is cognitively integrated, *vis-à-vis*, Googling. John never attempts to store information in biological memory and always automatically Googles a range of trusted sites, in a way that is as fluid and transparent as the way Otto automatically consults his notebook. In such a circumstance, might John *not* be bluffing, when he makes a comment of the sort which he is described as making in §1? Might he count as having extended knowledge?

Reply: This depends importantly on two things. Firstly, how we fill out further details of the case, and secondly, which philosophical rationale for extended cognition one is embracing: commonsense functionalism or dynamical

¹⁶ See also Pritchard (2010) and Palermos (2014; 2015b) for arguments to the effect that satisfying cognitive integrations for extended cognition is at the same time to satisfy the epistemological conditions requisite for one's externally aided belief forming process to qualify as a knowledge-generating *cognitive ability*. On Pritchard and Palermos's line, note that knowledge doesn't require that one's correctness be *primarily* creditable to one's cognitive ability (cf., Greco 2003; 2010; Sosa 2009) but just that the correctness of one's belief be significantly creditable to one's exercise of cognitive ability. In this way, Pritchard and Palermos are distancing themselves from robust virtue epistemologists such as Greco and Sosa. For related discussion, see Kallestrup and Pritchard (2012; 2013)

systems theory (see §5). Even if John seems to satisfy Clark and Chalmers' parity principle (which is not clear¹⁷), he fails extended cognition conditions according to dynamical systems theory (e.g., Palermos 2014) so long as there are not—as there are in the cases of Otto and Lilah—feedback loops between John and the device he is using to Google. In short, if the direction of causation is primarily 'one-way', from Google to John, then he fails to count as mutually interacting with the source of the information like Otto and Lilah are—i.e., where it is second nature to be *affecting* their source, by continually updating it. *However*, we can imagine an idiosyncratic kind of Googling where such feedback loops are more plausibly present than they are in the original case. Suppose, for instance, that the site John queries on his phone is a collaborative music Wiki, such as classical.net, which allows users to update and edit content¹⁸. Suppose further that John constantly updates the site, has been doing so for a long time, is familiar with the entries, goes to the site automatically, etc. In such a scenario, John is arguably much closer to satisfying the feedback loop condition, under dynamical systems theory, than he does in the version of the case described in §1. More generally, on the extended cognition approach, the question of whether an external artefact should be included within an individual's cognitive process is always going to be underdetermined in light of facts about the material constitution of the gadget or whether a particular kind of search engine is used. What matters in the main is cognitive integration, which can't be settled without looking closely at the details about the gadget's functional role and the way the agent typically interacts with it.

§7 Concluding remarks

¹⁷ After all, even if information storage and retrieval *vis-à-vis* one's biological memory is the intracranial functional analogue of Otto's process of consulting the notebook, it's not clear what the intracranial analogue process is to Googling.

¹⁸ <http://www.classical.net>

As cognitive offloading increases, we can expect an influx in assertions that are based on the information stored in our gadgets rather than in our heads. Some of these assertions are properly criticisable; they can—as with the case of John in §1—generate the false implicature that the asserter is more knowledgeable on the topic of the assertion than the asserter actually is. Interestingly, though, other gadget-assisted assertions which do not draw from any information stored in our heads nonetheless do *not* seem criticisable, and this is so even when such assertions generate the very same implicature that the asserter is knowledgeable and competent on the topic in question. While this is a genuine puzzle for those who embrace traditional thinking about the bounds of cognition, we’ve shown how the resources of extended cognition provide a neat solution. More generally, we hope to have shown a straightforward way in which recent thinking in the philosophy of mind and cognitive science can cut some ice in discussions about assertoric norms¹⁹.

¹⁹ The authors would like to thank S. Orestis Palermos, Mona Simion and Duncan Pritchard for helpful discussion. Gordon’s research in connection with this paper has benefitted from the ‘Intellectual Humility MOOC’ grant funded by the John Templeton Foundation, hosted at the University of Edinburgh’s Eidyn research centre.

REFERENCES

- Adams, Fred, and Ken Aizawa. 2001. 'The Bounds of Cognition'. *Philosophical Psychology* 14 (1): 43–64.
- . 2008. *The Bounds of Cognition*. John Wiley & Sons.
- . 2010. 'The Value of Cognitivism in Thinking about Extended Cognition'. *Phenomenology and the Cognitive Sciences* 9 (4): 579–603.
- Allen-Hermanson, Sean. 2012. 'Superdupersizing the Mind: Extended Cognition and the Persistence of Cognitive Bloat'. *Philosophical Studies* 164 (3): 791–806. doi:10.1007/s11098-012-9914-7.
- Beer, Randall D. 1995. 'A Dynamical Systems Perspective on Agent-Environment Interaction'. *Artificial Intelligence* 72 (1): 173–215.
- . 2000. 'Dynamical Approaches to Cognitive Science'. *Trends in Cognitive Sciences* 4 (3): 91–99.
- Benton, Matthew A. 2014a. 'Gricean Quality'. *Noûs*. doi:10.1111/nous.12065.
- . 2014b. 'Expert Opinion and Second-Hand Knowledge'. *Philosophy and Phenomenological Research*, May, 1–17. doi:10.1111/phpr.12109.
- Carter, J. Adam. 2013. 'Extended Cognition and Epistemic Luck'. *Synthese* 190 (18): 4201–4214.
- . Forthcoming. 'Epistemic Luck and the Extended Mind.' *The Routledge Handbook of Theories of Luck* (ed.) I.M. Church, London: Routledge.
- Carter, J. Adam and Emma C. Gordon. 2011. 'Norms of Assertion: The Quantity and Quality of Epistemic Support.' *Philosophia* 39 (4): 615–635.
- . Forthcoming. 'Knowledge, Assertion and Intellectual Humility'. *Logos & Episteme*.
- Carter, J. Adam, Emma C. Gordon and S. Orestis Palermos. 'Extended Emotion.' *Philosophical Psychology* 29(2): 197–218.
- Carter, J. Adam and Jesper Kallestrup. 2016. 'Extended Cognition and Propositional Memory.' *Philosophy and Phenomenological Research* 92 (3): 691–714.
- Carter, J. Adam, Jesper Kallestrup, Orestis Palermos and Duncan Pritchard. 2014. 'Varieties of Externalism.' *Philosophical Issues (supplement to Noûs)*, 24 (1): 63–109.
- Carter, J. Adam and S. Orestis Palermos. 2014. 'Active Externalism and Epistemic Internalism', *Erkenntnis* 80: 753–772.
- Carter, J. Adam and Duncan Pritchard. Forthcoming. 'Extended Entitlement.' In *New Essays on Entitlement*, (eds.) P. Graham & N. Pedersen, (Oxford UP).
- Chu, M. and Kita, S. 2011. 'The nature of gestures' beneficial role in spatial

- problem solving.' *Journal of Experimental Psychology* 140: 102-116.
- Clark, Andy. 2010a. 'Coupling, Constitution and the Cognitive Kind'. In *The Extended Mind*, edited by Richard Menary, 81–99. Cambridge, M.A: MIT Press.
- . 2010b. 'Memento's Revenge : The Extended Mind Extended'. In *The Extended Mind*, edited by Richard Menary, 43–66. MIT Press.
- . 2015. 'What "Extended Me" Knows'. *Synthese*, April, 1–19. doi:10.1007/s11229-015-0719-z.
- Clark, Andy, and David Chalmers. 1998. 'The Extended Mind'. *Analysis* 58 (1): 7–19.
- Cowan, N. 2010. 'The Magical Mystery Four: How is Working Memory Capacity Limited, and Why?' *Current Directions in Psychological Science* 19: 51-57.
- DeRose, Keith. 2002. 'Assertion, Knowledge, and Context'. *Philosophical Review* 111 (2): 167–203.
- Gettier, Edmund. 1963. 'Is Justified True Belief Knowledge?' *Analysis* 23 (6): 121–23.
- Goldberg, Sanford. 2013. 'Disagreement, Defeat, and Assertion'. In *The Epistemology of Disagreement: New Essays*, edited by David Christensen and Jennifer Lackey, 167–89. Oxford: Oxford University Press.
- Greco, John. 2003. 'Knowledge as Credit for True Belief'. In *Intellectual Virtue: Perspectives From Ethics and Epistemology*, edited by Michael DePaul and Zagzebski, Linda, 111–34. Oxford: Oxford University Press.
- . 2010. *Achieving Knowledge*. Cambridge: Cambridge University Press.
- Grice, H. Paul. 1975. 'Logic and Conversation', In Peter Cole and Jerry Morgan, eds., *Syntax and Semantics*, volume 3: Speech Acts, 43–58. New York: Academic Press.
- . 1991. *Studies in the Way of Words*. Harvard University Press.
- Hawthorne, J. 2004. *Knowledge and Lotteries*. Oxford: Oxford University Press.
- Hirschberg, Julia. 1985. *A Theory of Scalar Implicature*. Ph.D. thesis, University of Pennsylvania.
- Jarvis, Benjamin. 2015. 'Epistemology and Radically Extended Cognition'. *Episteme* 12 (04): 459–78.
- Kallestrup, Jesper, and Duncan Pritchard. 2012. 'Robust Virtue Epistemology and Epistemic Anti-Individualism'. *Pacific Philosophical Quarterly* 93 (1): 84–103.
- . 2013. 'Robust Virtue Epistemology and Epistemic Dependence'. In *Knowledge, Virtue and Action: Putting Epistemic Virtues to Work*, edited by Henning, Tim and David P. Schweikard, 209–26. London: Routledge.
- Kelp, Christoph. 2013a. 'Extended Cognition and Robust Virtue Epistemology'. *Erkenntnis* 78 (2): 245–52.
- . 2013b. 'Extended Cognition and Robust Virtue Epistemology: Response to Vaesen'. *Erkenntnis*, 1–4.
- Lackey, Jennifer. 2011. 'Assertion and Isolated Second-Hand Knowledge'. In

- Assertion: New Philosophical Essays*, edited by Jessica Brown and Herman Cappelen, 251–76. Oxford: Oxford University Press.
- Levinson, Stephen C. 1983. *Pragmatics*. Cambridge: Cambridge University Press.
- Lynch, Michael Patrick. 2014. 'Neuromedia, Extended Knowledge and Understanding'. *Philosophical Issues* 24 (1): 299–313.
- Menary, Richard, ed. 2010. *The Extended Mind*. Cambridge, M.A: MIT Press.
- McGlynn, Aidan. 2014. *Knowledge First?* London: Palgrave MacMillan.
- Palermos, Orestis. 2011. 'Belief-Forming Processes, Extended'. *Review of Philosophy and Psychology* 2 (4): 741–65.
- . 2014. 'Loops, Constitution, and Cognitive Extension'. *Cognitive Systems Research* 27: 25–41.
- Palermos, S. Orestis. 2014. 'Knowledge and Cognitive Integration'. *Synthese* 191 (8): 1931–51.
- . 2015a. 'Active Externalism, Virtue Reliabilism and Scientific Knowledge'. *Synthese* 192 (9): 2955–86. doi:10.1007/s11229-015-0695-3.
- . 2015b. 'Spreading the Credit: Virtue Reliabilism and Weak Epistemic Anti-Individualism'. *Erkenntnis*, May, 1–30. doi:10.1007/s10670-015-9741-2.
- Pritchard, Duncan. 2010. 'Cognitive Ability and the Extended Cognition Thesis'. *Synthese* 175 (1): 133–51.
- . 2012. 'Anti-Luck Virtue Epistemology'. *Journal of Philosophy* 109: 247–79.
- Risko, E.F. et al. 2014. 'Rotating with rotated text: a natural behavior approach to investigating cognitive offloading.' *Cognitive Science* 38: 537-564.
- Risko, E.F. and Gilbert, S. J. 2016. 'Cognitive Offloading.' *Trends in Cognitive Science* 20: 676-688.
- Rupert, Robert D. 2004. 'Challenges to the Hypothesis of Extended Cognition'. *Journal of Philosophy* 101 (8): 389–428.
- Simion, Mona. 2015. 'Assertion: Knowledge is Enough.' *Synthese*, doi: 10.1007/s11229-015-0914-y
- Simons, D. J. and Levin, D. T. 1997. 'Change Blindness'. *Trends in Cognitive Sciences*. 1: 261-267.
- Sosa, Ernest. 2009. *A Virtue Epistemology: Apt Belief and Reflective Knowledge (Vol. 1)*. Oxford: Oxford University Press.
- Unger, Peter K. 1975. *Ignorance: A Case for Scepticism*. 105. Oxford University Press.
- Vaesen, Krist. 2011. 'Knowledge without Credit, Exhibit 4: Extended Cognition'. *Synthese* 181 (3): 515–29.
- . 2013. 'Critical Discussion: Virtue Epistemology and Extended Cognition: A Reply to Kelp and Greco'. *Erkenntnis* 78 (4): 963–70.
- Williamson, Timothy. 2000. *Knowledge and Its Limits*. Oxford University Press.