



The brittleness of expertise and why it matters

Daniel Kilov¹ 

Received: 30 March 2020 / Accepted: 27 October 2020 / Published online: 9 November 2020
© Springer Nature B.V. 2020

Abstract

Expertise has become a topic of increased interest to philosophers. Fascinating in its own right, expertise also plays a crucial role in several philosophical debates. My aim in this paper is to draw attention to an important, and hitherto unappreciated feature of expertise: its brittleness. Experts are often unable to transfer their proficiency in one domain to other, even intuitively similar domains. Experts are often unable to flexibly respond to changes within their domains. And, even more surprisingly, experts will occasionally be outperformed by novices when confronted with novel circumstances within their domains of expertise. In section 1, I marshal the evidence in favour of brittleness. In section 2, I argue that appeals to brittleness can advance the dialectic in debates on skilled action and provide reasons to reject a powerful recent argument offered by Christensen et al. (*Philos Psychol* 32(5): 693–719, 2019). In section 3, I appeal to brittleness to argue against a common conception of philosophical expertise, according to which philosophers possess a domain-general set of reasoning skills. Although my argument in this section is largely negative, there is a twist. Recalibrating our understanding of philosophical expertise opens new avenues of research for defenders of the so-called ‘expertise defence’ against the findings of experimental philosophy.

Keywords Expertise · Brittleness · Social epistemology · Skilled action · Expertise defence

*“Some people regard the former as one who knows a great deal about a very little and who keeps on knowing more and more about less and less until he knows everything about nothing. Then he is a scientist...
...Then there are the latter specimen who knows a little about very much and he continues to know less and less about more and more until he knows nothing about everything. Then he is a philosopher.”*

✉ Daniel Kilov
Daniel.Kilov@gmail.com

¹ The Brittleness of Expertise and Why It Matters, Australian National University, Canberra, ACT, Australia

-Robert E. Swain on the difference between a scientist and a philosopher (1928)

1 Introduction

Expertise has, in recent years, become a topic of increased interest to philosophers. Expertise is a fascinating subject, worthy of study in its own right. But it also plays a crucial role in several philosophical debates. For instance, epistemologists have become increasingly sensitive to the ways in which we are dependent on experts for many of our beliefs. The usual case is not, as Descartes may have believed, one where we derive our beliefs from first principles. Nor is it the case, as Empiricists suggested, that we justify our beliefs by direct observation. Rather, the standard case is one wherein we defer to experts, and this is becoming ever more common as human knowledge specializes. Meanwhile, experimental philosophers and moral philosophers have wondered about philosophical expertise, and philosophers of action and philosophers of mind have looked towards expertise to inform debates about intention and skill, among other things. My aim in this paper is to draw attention to an important feature of expertise and to draw out the implications for some of these debates. Specifically, I argue that expertise is remarkably brittle in the following ways. Expertise is domain-specific. That is, there is very little transfer from proficiency in one domain even to other, seemingly similar domains. Even more remarkably, experts are often unable to flexibly respond to changes *within* their domains. Furthermore, this inflexibility is not limited to situations in which the domain itself is altered but is also observed when the domain remains intact but novel choices are presented. In Sect. 1, I marshal the empirical evidence in support of brittleness. In Sect. 2, I consider the implications of brittleness for debates on the nature of expert skill and knowledge. In Sect. 3, I use the arguments developed in the preceding sections to argue against a widespread conception of philosophical expertise.

Beyond the issues discussed here, the brittleness of expertise is also relevant to debates in social epistemology; in particular to debates about the role of public experts and trust in journalism. Consider the so-called ‘Gell-Mann Amnesia effect’,¹ described by the science fiction author Michael Crichton (Crichton, 2020) in his 2002 speech, ‘Why Speculate?’:

Briefly stated, the Gell-Mann Amnesia effect is as follows. You open the newspaper to an article on some subject you know well. In Murray’s case, physics. In mine, show business. You read the article and see the journalist has absolutely no understanding of either the facts or the issues. Often, the article is so wrong it actually presents the story backward—reversing cause and effect. I call these the “wet streets cause rain” stories. Paper’s full of them. In any case, you read with exasperation or amusement the multiple errors in a story, and then turn the page to national or international affairs, and read as if the rest of the newspaper was somehow more accurate about Palestine than the baloney you just read. You turn the page and forget what you know.

¹ Crichton named the effect ironically, "because I once discussed it with Murray Gell-Mann, and by dropping a famous name I imply greater importance to myself, and to the effect, than it would otherwise have".

If such an effect exists, it would provide at least two excellent examples of brittleness. The first concerns the brittleness of the expert reading the news and the second concerns the brittleness of the journalists writing outside their expertise. Although it is beyond the scope of this paper to explore these issues, additional work should be done to consider the consequence of the brittleness wherever philosophers appeal to expertise.

2 Section 1: Evidence for the brittleness of expertise

First, some preliminary wrangling of terminology. My aim here is not to offer an analysis of expertise. However, I do want to draw a distinction between individuals who have exceptional abilities in some domain and those who merely enjoy social recognition or reputational expertise. It is experts of the former kind that are the topic of this paper. By a domain of expertise, I mean something like the subset of configurations of a system reachable by following the rules of the activity in which one is an expert.² This set is, in general, much smaller than the total state space of the system in which the activity takes place. For instance, the domain of chess consists of the set of positions that can be reached by game-legal moves, which is smaller than the set of positions that can be achieved by placing combinations of the available chess pieces directly on the board. Furthermore, unless otherwise specified, I'll be talking about those parts of the domain that are relevant to realistic performance. So, my account doesn't require that experts display superior performance in relation to random chess positions, even ones that can be reached through legal moves. It is, however, a consequence of my account that individuals can lose their status as experts without any change in their abilities if the rules of the domain change enough or if the set of positions relevant to performance in that domain shift too far (sometimes referred to as the *metagame*). For example, a chess expert who has built her game around a particular opening might lose her status as an expert once an effective counter to that opening is discovered and becomes widespread. In this case, the rules of chess have not changed but the metagame has shifted. This analysis applies more naturally to games like chess than to domains like philosophy, but there are analogies here. In philosophy we make various dialectical moves and explore different parts of the logical space. We have our own set of rules. We (generally) aim for truth, consistency and fruitfulness in our theorising, among other theoretical virtues. Philosophy is different to chess and other games in that the rules can be negotiated on the fly, but philosophical interlocutors must nonetheless share some standards in order to engage with one-another. In practice, these meta-philosophical commitments are often implicit until they become relevant to philosophical debate, but this is not so different to situations in which participants of a game encounter something during play that triggers an argument about the rules of the game. The difference, of course, is that participants in a game can often appeal to certain authorities or canonical rules to resolve disagreements. We are not so lucky in philosophy.

² There is some wiggle room here. Domains are not so fine-grained that any change in rules results in a change in domain. The rules of chess have changed over time, for instance.

The other key notion is brittleness. Like the brittleness of physical materials, the brittleness of expertise is a gradable notion; it admits of degree, and some forms of expertise may be more brittle than others. Similarly, just as the fractural dispositions of different physical materials vary, different forms of expertise may fail in different ways or under different circumstances. For instance, experts in some domains may respond flexibly to novel configurations within their domains but fail when the domains are changed slightly, or vice versa.

Computer scientists sometimes talk about the brittleness of software.³ This too, provides a useful analogy. Software is considered brittle when, despite appearing reliable, it fails when presented with unusual data or when the digital environment is altered in seemingly minor ways. Experts and software alike tend to fail dramatically, when they do, and often at inopportune times. More speculatively, it may be that experts and software sometimes become brittle for similar reasons. For instance, a piece of software can become brittle if its component dependencies are too rigid. When a component is designed to accept only a certain range of inputs, and that range changes, it can cause errors that ripple throughout the system. Further, changing or updating the problematic component may be impossible because too many other parts of the software are built on top of it, producing a form of technological lock-in. As a program grows more sophisticated and the number of inter-related parts increases, so too does the likelihood that this kind of problem will emerge. If expert skills exhibit similar structural features (e.g. if expert skills are built up from or depend on more foundational ones) it may explain why, as we shall see below, experts are sometimes less able to adapt to subtle changes in their domains than novices.

In making an analogy to software, I do not mean to stack the deck in favour of a representationalist account of brittleness. The experimental findings discussed below should be interpreted as instrumentally as possible. I take the results of these studies seriously, but I am agnostic as to what explains them. Although I will be arguing that skilled action involves representation later in the paper, those arguments float free from any particular theory of brittleness.

That expertise is domain-specific is, in a certain sense, unsurprising. Medical doctors who are expert diagnosticians are no better or worse at troubleshooting dysfunctional washing machines than the rest of us. What is surprising is just how context-dependent expertise is. For instance, there is now a large body of research showing that expertise in surgery is highly local. The ability to perform a particular surgical task derives from specific practice of that task and does not generalise even to similar tasks (Wanzel et al. 2002). Not only do surgical skills not transfer from one procedure to another, they also provide little advantage on tasks explicitly designed to approximate the domain of their expertise. Van Sickle et al. (2007) found no correlation between years of practice or number of laparoscopic procedures completed and performance on a virtual reality laparoscopic training simulator. Instead, what determined performance on the simulator was just the surgeon's accumulated experience on the specific simulator. Similarly, Park et al. (2007) found that performance in training simulators were poor predictors of performance in a clinical setting. These

³ 'Brittle software' struck me as a contradiction in terms until a colleague noted that my immediate environment contained two examples of objects that are both soft *and* brittle: blackboard chalk and biscuits!.

results should not make us sceptical of surgical expertise. Rather, they highlight how sensitive expertise is to the contours of a domain. Perhaps, when looking for a hand surgeon, we should start by asking whether they specialise in the right or left hand!

Consider one more example: being an air-traffic controller is hard. The stakes are high. A wrong decision can lead to disaster. The work is cognitively demanding. It requires the constant updating and integration of complex information with prior knowledge. Air-traffic controllers thus develop dexterity and flexibility in their thinking processes. Yntema and Mueser (1960) wanted to know whether air-traffic controllers had a superior general ability to keep track of many things at once. Subjects in their study undertook a series of memory-based tasks with shapes and colours. The air traffic controllers performed no better than the general population. Their sophisticated cognitive abilities did not translate beyond their professional area. As Feltovich et al. (2006) point out, expertise typically develops in very narrow and highly specific ways such that there is “little transfer from high-level proficiency in one domain to proficiency in other domains—even when the domains seem, intuitively, very similar”.

The studies discussed above all feature skills involving a substantial cognitive component. Philosophers interested in defending a sharp distinction between intellectual activities and motor skills might wonder whether athletic abilities transfer better than intellectual ones. Intuitively, gross motor skills seem particularly good candidates for transfer. As an anonymous reviewer helpfully put it, “A good baseball pitcher seems to be a good candidate for grenade thrower in the army; a skilled batter should be good at hitting many other objects other than baseballs”. Here too, however, context matters. Throwing a grenade is unlike throwing a ball; there is a pin, counting and holding before release and it can explode in one’s hand! A pitcher in a novel, high pressure situation may well see their abilities compromised.

Less speculatively, there is empirical evidence that batting ability does not transfer from baseball, to softball. Professional baseball hitters were unable to hit throws by Olympic softball champion Jenny Finch (A Women’s Softball Pitcher vs. the Top Baseball Hitters...Who Wins? 2020). Although Major League Baseball players routinely hit baseballs travelling at over 153 kph, differences in the field and pitching style meant that they were unable to connect with a larger, slower target (Finch pitches at a ‘mere’ 110 kph).⁴ Thus, even gross-motor skills are often surprisingly brittle.

Expertise is not only brittle in the sense that it cannot be deployed outside of a narrow domain of expertise. Experts are often unable to flexibly respond to changes *within* their domains. For example, Sternberg and Frensch (1992) compared expert and novice bridge players and examined the effects of various arbitrary rule changes on their performance. In general, experts were found to suffer more than novices from any rule change. In another study, expert accountants proved less adept than novices at using a new tax law that rendered obsolete a previous rule concerning tax deductions (Marchant et al. 1991). Similarly, while studying the phenomenon of choking in sports, Beilock et al. (2002) reported that the performance of expert soccer players suffered more than that of novices during exercises that required them to memorise a series of words while dribbling a ball.

⁴ Thanks to an anonymous reviewer for both this and the ‘grenade’ example.

We must, of course, exercise some caution when interpreting these results. After all, my account of expertise requires only that experts display superior performance with regards to those parts of the domain that are relevant to *realistic* performance. Expert soccer players are not required to remember random words in *actual* soccer games and arbitrary rule changes in bridge games are just that—arbitrary! Nonetheless, the fact that novices were able to adapt to the new conditions more readily demonstrates a surprising dimension of expert brittleness: experts are less flexible in response to novel circumstances than we might have thought. Further, the studies above speak to the fact that expert skill is remarkably narrow. Indeed, even very subtle changes can lead to significant decreases in performance. Sims and Mayer (2002) found evidence for extreme specificity of skill in expert Tetris players. Tetris is a video game that requires players to rotate shapes comprised of four equally sized squares to create rows of pieces within a limited timeframe. The researchers compared the general spatial ability of Tetris players of varying levels of skill. The spatial tests variously involved rotation of standard Tetris shapes, shapes similar to Tetris shapes, and other shapes, such as letters and numbers. The results showed that highly skilled players outperformed less skilled players *only in the rotation of Tetris shapes*. In a second phase of the study, novices were trained on Tetris for 12 h. This practice improved the participants' ability to rotate the Tetris-like shapes but had no effect on their more general spatial ability, thereby again highlighting the specificity of the Tetris skill.

Inflexibility of this kind is not limited to situations in which the domain itself is altered. It is also observed when the domain remains intact but novel choices are presented. For instance, Saariluoma (1991) found that experts at blind-fold chess were unable to track the positions of pieces if random—rather than meaningful—moves were performed. This is consistent with the now famous research by Chase and Simon (1973), which demonstrated that the superior recall of chess experts largely disappears when random board configurations, rather than configurations of actual boards, are used in the recall task. Additionally, expertise can sometimes be an obstacle to problem solving. Known as *Einstellung* effects, these are trained responses that prevent the discovery of novel, superior solutions. When an expert's domain-specific representations are activated, certain solutions immediately come to mind. For instance, in a pair of studies on chess problem solving (Bilalić et al. 2008; Saariluoma 1990), players of various levels of expertise were presented with a sequence of chess problems and asked to find the best solution. When the first four problems were solvable by a 'smothered mate' motif,⁵ experts failed to notice that the fifth problem could be solved by other, objectively better means. When the problems were presented on their own, however (i.e. without the *Einstellung* stimuli), almost all the experts were able to find the better solutions.

It's important not to overstate the case here. Clearly, some skills do generalise to many different contexts. Reading is an example. Although I find myself reading mostly philosophy papers, I could, in principle, read works in other disciplines, or in fiction, with negligible drop in performance. However, I suspect that, in general, as situational demands increase, the skills required for expert performance become

⁵ A smothered mate is a checkmate achieved when the mated king is unable to move because he is surrounded (or smothered) by his own pieces.

increasingly specialised/narrow. For example, I probably would struggle to read papers in physics, or philology. Skills may share something like the trade-offs recognised in ecological models between generality and predictive power (Matthewson and Weisberg 2009). Consequently, these findings should make us sceptical of an image of philosophical expertise as consisting in a domain-general set of thinking tools. But this self-conception is widespread. Consider this quote from Ganeri (2018):

It is a remarkable feature about philosophy that, no matter how different their areas of specialization are, philosophers can and do talk to each other. In your regular departmental colloquium, it would be normal for a visiting speaker to be questioned by specialists in Aristotle as much as by the resident philosopher of mathematics. This is because what philosophers share is what I described before as a basic tool kit for the management of disagreement: spotting inconsistencies in an argument or an overlooked alternative explanation that can reconcile apparently contrary assumptions, winking out hidden assumptions and missing steps. This ability to talk to each other across specialization is something philosophers greatly prize.

I concede that Ganeri's view is common-sense (at least among professional philosophers), but theorists need to move beyond anecdote and informal observation and engage with the rich and sophisticated psychological research on expertise. The research reviewed above suggests that experts generally do not have domain-general problem-solving skills. Further, studies of experts illustrate how the processes of thinking are tightly knit to the content of thought. Reasoning effectively requires domain-specific knowledge of the kind philosophers are likely to lack, outside of their areas of speciality. I address these issues in depth in Sect. 3. In the next section, I explore the significance of brittleness to debates about the nature of skilled action and expert performance.

3 Section 2: Skilled action

3.1 Background

Knowledge and skill are intimately related. Skilled painters, dancers, and chess players, for instance, possess an enormous amount of knowledge about their domains in addition to their skills at performance. They may be experts on the historical developments of their domains, on the biographical details of key figures or on the training methods and techniques of their fields. Knowledgeable linguists, chemists and biologists also possess considerable skills of analysis, experimental design and argumentation in their respective domains. Despite this intimate relationship, both analytic philosophers and psychologists have understood knowledge and skill as distinct. Philosophers, at least since Ryle (2009), have taken knowledge-that to be distinct from knowledge-how. Cognitive scientists, meanwhile, cleave a similar distinction between declarative memory and procedural memory. Procedural skills are thought to be non-cognitive, automatic and unconscious. The declarative knowledge side of this dichotomy is generally characterized as cognitive, intentional and conscious. In this section I will argue

that that this dichotomy clusters together traits that come apart in important ways. Although I will be arguing, in line with some recent work by Christensen et al. (2019) that skilled action is cognitive,⁶ I will also argue, contra that work, that skilled action is largely automatic. Finally, I will suggest that expert skill, although largely automatic, is nonetheless conscious.

Much of the credit for renewed interest in these debates goes to Stanley and Williamson (2001), who argued that our best semantic theories for knowledge-how ascriptions entail that knowing how to φ is knowing that q , where q is a proposition containing a way to φ . Their argument has been the target of criticism, however, with philosophers seemingly divided by methodological fault-lines [for example, see Schwartz and Drayson (2019)]. Noë (2005) described their linguistic approach as GOOP (or ‘Good old Oxford philosophy’), suggesting that we ought to care about the distinction between knowing-how and knowing-that only if it is a feature of our psychological reality, rather than of our speech. After all, we are not after an account of how we use sentences that ascribe knowledge-how. Rather, we care about the truth of these ascriptions. Noë (2005) points to Stanley & Williamson’s ascription of knowledge-how to non-human animals in order to illustrate his objection. As he writes, “the point is not that dogs can’t grasp propositions. The point is that whether or not they can grasp propositions is an open question, one that is debated in cognitive science. The problem for Stanley and Williamson is that their analysis commits them to the strong consequence that dogs can grasp propositions, at least if it is to have any hope of being true” (p. 12). In addition, Noë appeals to evidence from embodied cognition to highlight the ways that skilled action exploits features of an agents’ body and world, eliminating the explanatory role of propositional representations. In a similar vein, Devitt (2011) argues that evidence from cognitive ethology suggests that procedural knowledge is entirely distinct from declarative knowledge and thus non-representational. In defence of this claim he references work on insects done by Gallistel and work on the caching strategies of scrub jays by Clayton. Devitt argues that, although this work reveals animals to have “surprisingly rich cognitive lives”, there is no sign in the literature “of retreat from the received view that procedural knowledge is quite distinct from declarative knowledge” (Devitt 2011, p. 6).

3.2 Cognitive ethology and procedural knowledge

I agree with Devitt (2011) that any attempt to understand procedural knowledge without attention to the relevant science is “deeply misguided”. However, both Gallistel’s (2008) work on bees and their “waggle dances”, and Clayton’s (1998) work on jays actually speak against the anti-intellectualism he advances. Far from suggesting that non-human animals lack representational capacities, Gallistel notes that “the behavioural evidence implies that the nervous system possesses a read–write memory mechanism that performs the same essential function performed by the memory registers in a computer” (p. 228). In addition, Gallistel suggests that “The results of behavioural experiments on nonhuman animals have increasingly implied that much learned behaviour is informed by enduring temporal and spatial representations (e.g.,

⁶ In the sense that sophisticated representations play a causal role in their manifestation.

Menzel et al. 2005), as even some prominent advocates of associative theories have recently acknowledged (Clayton et al. 2006)” (p. 227). For example, in one typical experiment, foraging bees returned from a food source on a rowboat in the middle of a pond or small lake. The returning foragers dance, but fail to recruit other bees. If the past experiences of the on-looker bees, as represented on their maps, indicate that there is no food to be found at the coordinates indicated by the dance, the bees decide not to go to the indicated coordinates (Gallistel 2008). This indicates that what the dance communicates is not instructions for flight paths but map coordinates, and the observing bees consult their own cognitive maps before deciding whether or not to act on the communicated information. Gallistel suggests that we have massively underestimated the representational and computational powers of “brains as small as the head of a pin” (p.235) and that we are laboring under a falsehood when we believe that these brains can “get along without the symbolic memory mechanisms that make representation possible” (p. 227).

Hutto and Myin (2017) have recently challenged Gallistel’s account, arguing that talk of representations, understood in terms of semantically evaluable content, is explanatorily superfluous. They claim that Gallistel’s bees should be understood as exploiting systematic isomorphism between features of their cognitive systems and their environment. As Hutto and Myin put it:

“Gallistel invokes talk of representations in the explanations he offers of such behavior, but the only items doing actual load-bearing work in his explanations are systematic structure-preserving correspondences—correspondences that hold between certain features of the organisms and certain features of their environments. Although the bees’ exploitation of those correspondences feature heavily in Gallistel’s explanations of their navigational behavior, contentful representations make no appearance at the level of cognitive drivers of such behavior. In discussing this very case, Rescorla (2012) makes clear that, “Explanatory power resides solely in the ‘functioning isomorphism’ between mind and world. There is no obvious reason why ‘functioning isomorphism’ must have truth conditional content. ... The burden of proof lies with those who claim that functioning isomorphism suffices for truth conditions” (Rescorla 2012, 96; see also Tonneau 2011/2012).” (Hutto and Myin 2017, p. 110)

Their argument is subtle and a full evaluation of it is beyond the scope of this paper. Nonetheless, let us suppose, for the sake of argument, that structural isomorphisms are insufficient for representation. Does it follow that we can do away with talk of truth conditions (and thus representation) when explaining the behaviour of Gallistel’s bees? I think not, for the following reason: as the bees learn and their internal isomorphisms are modified in response to environmental stimuli, they are not changed randomly or wholesale. Rather, at least when things go well, they are updated in a way that makes them more accurate. But once we admit talk of accuracy/inaccuracy, the notion of truth conditions begins to do real explanatory work.

As brain size increases from bug to bird, the case for representation is even more compelling. A series of ingenious experiments conducted by Clayton & Dickson (1998) on the food-caching of jays imputes to them the ability not just to remember and recall specific episodes (of the kind considered declarative memory) but also

to modify behavior on the basis of inferences and semantic memory. The jays in the caching experiments were able to remember what kind of food they had hidden in each cache, which other jays might be watching and whether those jays might be likely to steal from their caches. They kept track of time and were aware that the contents of certain caches would expire before others. Although the autonotic consciousness (i.e. the awareness of self) that accompanies episodic recall has no obvious non-linguistic indicators, and is thus probably undetectable in many species, the cache recovery pattern of scrub jays “fulfils the three, ‘what’, ‘where’ and ‘when’ criteria for episodic recall and thus provides, to our knowledge, the first conclusive behavioural evidence of episodic-like memory in animals other than humans” (Clayton and Dickinson 1998). As Gallistel (2008) writes, “the information drawn from memory that is combined to inform current behaviour comes from a mixture of episodic memories (“Three days ago, I hid meal worms there, there and there, and 5 days ago, I hid peanuts there, there and there”) and declarative memories (“Meal worms rot in 2 days; peanuts don’t rot”)” (p. 236).

Although Devitt (2011) insists that “Psychology presents a picture of procedural knowledge as constituted somehow or other by embodied, probably unrepresented, rules” (p. 6), the evidence he appeals to actually suggests that our understanding of animal behaviour is best served by supposing a symbolic memory which is capable of encoding and representing information about the world and “carrying that information forward in time in a computationally accessible form” (Gallistel, 2008, p. 234).

Devitt and Noë (and Ryle, for that matter) might be right that there is some form of irreducibly non-representational knowledge.⁷ But the work discussed above gives a central role to representations. Thus, their arguments give us no reason to deny that skilled behaviour is mediated by rich representational structures. As we shall see, there is compelling psychological evidence that this is the case.

3.3 Long term working memory and expert skill

Experts in a variety of fields can perform feats which, to the uninitiated, appear to border on the supernatural. In 2011, the German chess player, Marc Lang, set a world record by playing 46 simultaneous games of blindfold chess.⁸ The Chinese memory athlete, Zou Lujian, can memorize the order of a shuffled deck of playing cards in 13.96 s. Shashank Jain, a human calculator from India, can mentally calculate the square root of an eight-digit number in 1 min and 25 s. Performing such cognitively complex tasks requires access to large volumes of information. Chess experts who play blindfolded, for instance, must keep track of the location of all the pieces on the board. They must also be able to update that information as the game progresses. Similarly, human calculators must not only represent the problems they are working on but also store and update intermediate answers. Experts thus make use of greater volumes of information than do novices in guiding their actions. Somewhat paradoxically, experts

⁷ Indeed, I think they probably *are* right. The evidence I have given above for the brittle and context-sensitive nature of expertise is well explained by accounts that emphasise the embedded and extended nature of know-how.

⁸ He ended up with 25 wins, 19 draws and only 2 losses.

also display much greater speed in making their decisions. A successful theory of expertise must be able to explain how such rapid and complex learning can occur given well known limitations on working memory. Working memory involves the short-term storage of information and is assumed to have a limited capacity of around seven pieces of information plus or minus two (Miller 1956). This capacity does not vary between experts and novices and places strict limits on what kind of operations can be performed.

Traditional theories of chunking predict that since chunks are retrieved from long-term memory (LTM) and maintained in short-term working memory, interruptions to working memory should result in reduced performance. Research has shown, however, that experts are substantially less likely than novices to have their performances derailed by interruption or delays imposed prior to recall on memory tasks (Charness 1976). Consequently, Chunking Theory is unable to account for the relatively rapid learning displayed by experts in dynamic environments. In addition, traditional Chunking Theory does not seem able to account for the ability of experts to move beyond the standard capacity of seven or so chunks (Ericsson and Kintsch 1995).

In light of this, Ericsson and Kintsch (1995) developed Long-Term Working-Memory Theory (LTWM Theory) according to which a cognitive process is “viewed as a sequence of stable states representing end products of processing. In skilled activities, acquired memory skills allow these end products to be stored in long-term memory and kept directly accessible by means of retrieval cues in short-term memory” (p. 211). Experts, therefore, overcome the limits of short-term working memory by exploiting prior knowledge to encode new information in terms of familiar, well-organized concepts represented in long term memory. In doing so, experts take advantage of “both the content and structure of an elaborate semantic memory network to create meaningful memory codes that created multiple potential cues and avenues for retrieval” (Ericsson & Staszewski, 1989, p. 239).

Although LTWM theory is not the only contemporary theory of expert performance,⁹ it should be of special interest to those interested in giving a philosophical theory of expertise because of its wide scope. As philosophers, we are interested not only in giving an account of expertise in one or other domain, but of expertise generally. LTWM theory has been applied to domains as diverse as mental calculation and medical diagnosis. On this view, expertise in a given domain is the product of acquired mental processes. These processes involve the development of hierarchically organized patterns and schemas, stored in long term memory, which can be rapidly accessed via developed retrieval structures. According to Ericsson and Lehmann (1996), these representations are developed over many years of ‘deliberate practice’ which consists of targeted, effortful practice at the limits of one’s abilities. Empirical support for LTWM theory comes from research indicating that deliberate practice is indispensable in the attainment of expert performance, as well as studies employing ‘think aloud’ protocols Ericsson and Simon (1984), which probe the cognitive processes involved in performance. Experts in these studies report employing rich cognitive processes (Ericsson 2006), providing additional evidence against the view that expert skill is unconscious and nonconceptual.

⁹ See, for example, (Gobet 1998).

Christensen et al. (2019) argue that Ericsson’s research also casts doubt on the idea that expert skills are largely automatic. This is because Ericsson’s notion of deliberate practice requires that experts maintain the ability to intervene on their performance, allowing them to “construct and refine increasingly complex cognitive mechanisms that allow higher levels of control, self-monitoring, and performance evaluation” (p. 698). They argue, however, that more needs to be said about why expert performance demands cognitive control in the first place. Specifically, they develop an argument to show that even a rich set of automatic responses is likely to be inadequate to meet the requirements of expertise. They call this the *Domain Size Argument*. In the next section, I examine the argument in detail and show why the brittleness of expertise poses a serious problem for Christensen et al.’s account. I then offer an alternative account of expertise that emphasises the way that experts transform tasks so they can rely on pre-learned, automated responses.

3.4 The domain size argument against automaticity

According to the *Domain Size Argument*, if expert skill is to be accounted for in terms of automatic responses, then experts will need to have encountered the relevant parts of their domain many times. This is because automatic responses are acquired slowly and over the course of repeated exposure. That experts could have had such exposure is extremely unlikely, however, since the state space of even simple systems is enormously complex. Christensen et al. (2019) offer the following example:

...even a system with relatively few elements can have a large state space because those few elements can enter into many combinations...

Moreover, for many types of systems, there will be a tendency for the size of the state space to increase nonlinearly with linear increases in the number of elements. The state space for a system with a relatively modest number of elements can, consequently, be extremely large. For instance, chess has 64 squares and 32 pieces, but the number of possible chess positions has been estimated to be between 10^{40} and 10^{50}

(Steinerberger 2015, p. 699).

Implicit in their argument is a commitment to the idea that experts can flexibly deploy their expertise across the state space of an activity (like chess). But as we saw in Sect. 1, this is false. Christensen et al. (2019) thus find themselves between Scylla and Charybdis: if they insist that experts can flexibly respond to a large enough part of the state-space, they run afoul of the evidence in favour of brittleness. If they relax their requirement to say that experts need only retain skilled performance across the parts of the state space relevant to realistic performance, then it’s hard to see why the total size of the state-space is at all important. If it’s only the subset of that total state space relevant to performance that matters, then the threat of combinatorial explosion is much less real.

Rubik’s cube might be an illustrative example here: there are 43 quintillion possible combinations of the cube. But the sheer number of combinations isn’t indicative of how difficult the cube is, or how large the relevant space is. In fact, as a puzzle, Rubik’s

cube is easy, mainly because it is possible to find algorithms that just act on very few of the pieces without disturbing the rest of the cube. Even elite speed solvers only use between 78 and 109 algorithms to solve the cube. If the states relevant to performance can be only a very small subset of the total state space, then there is no problem with accounting for them in terms of repeated exposure to previously encountered combinations of factors.

Christensen et al. (2019) might reply that Rubik's Cube is, perhaps atypical. Most domains are more complex. The vast size of the chess domain, for instance, means that no player could ever be exposed to more than a tiny fraction of the potential contingencies. Nonetheless, these contingencies play a vital role in chess. Chess is competitive, and players stand to gain an advantage by presenting their opponents with novel situations. Chess experts must therefore be able to flexibly cope with novel situations if they are to count as experts at all. There is some evidence to support this idea. Meta-analysis of studies on chess experts shows a small but clear memory advantage, even for random (and thus novel) board position (Gobet and Simon 1996). However, this small advantage is likely best explained by the expert's ability to discover even small regularities in otherwise random positions by matching the presented boards against a vast repertoire of patterns in long term memory—a repertoire which could contain as many as 300,000 chunks (Gobet and Simon 2000). So, this advantage is circumscribed in exactly the way brittleness would predict. As Lewandowsky and Thomas (2009) put it, "Accordingly, when the degree of randomness (defined by the extent to which basic game constraints are violated) is manipulated, players with greater expertise have been found to be better able to exploit any remaining regularities than players with lesser expertise Gobet and Waters (2003). Thus, the specificity of expertise extends to highly subtle regularities indeed" (p. 51).

Rubik's cube might be unusually simple, as a domain, but I maintain that it is nonetheless typical as an illustration of how experts rely on pre-learned patterns to transform unfamiliar and complex scenarios into familiar and manageable ones. The memory techniques of expert mnemonists provide additional illustration and are of special interest since mnemonic techniques played an important historical role in the development of Long-Term Working-Memory theory to which Christensen et al. (2019) appeal. Indeed, memory techniques were the paradigm case in Ericsson and Kintsch (1995).

Speed cards is the prestige event in memory competitions. The event requires competitors to memorise the order of a randomly shuffled deck of cards as quickly as possible. Records for this event are broken every year and the current world record is 12.74 s, held by the Mongolian mnemonist Shijir-erdene Bat-enkh. Card memorisation is, by the standards of the domain size argument, an incredibly large domain. The number of possible combinations of a deck of cards is 52 factorial. This is a number so large that if every star in our galaxy had a trillion planets, each with a trillion people living on them, and each of these people has a trillion packs of cards and somehow they manage to make unique shuffles 1000 times per second, and they'd been doing that since the Big Bang, they'd only just now be starting to repeat shuffles ("Jumble," 2012). So, there is an obvious sense in which every deck of cards a mnemonist encounters is completely unfamiliar. However, there is also an important sense in which mnemonists *don't* memorise unfamiliar strings of cards at all. Rather,

they transform combinations of cards into distinctive, pre-learned images which they place in pre-learned memory palaces. A large part of mnemonic training consists in building and learning sophisticated mnemonic systems. When mnemonists lack task specific encoding strategies, their performance on memory tasks is no better than the general population (Maguire, Valentine, Wilding, & Kapur, 2003).

Chess, speed cubing and memory sports are, on the traditional way of dividing things up, all largely cognitive skills. How does my account fair with more embodied skills, where expertise appears to have less to do with discrete operations within a formalised set of rules and more to do with improvisational performance, as in sports? This question assumes that a principled distinction can be drawn between the performances of chess players, speed cubers and mnemonists on the one hand, and improvisational performances on the other. But this is precisely what I deny. To make the point vivid, consider Olympic judo. Olympic judo is a fast paced, dynamic combat sport in which competitors can win by throwing their opponents, pinning them or submitting them with an armbar or stranglehold. The official curriculum of Kodokan judo lists 68 throwing techniques and 32 grappling (newaza) techniques, though in reality, there are many more. Nonetheless, a study of 39 World and/or Olympic judo champions found that the average champion utilizes a mere six throwing techniques and two grappling techniques in competition (Weers 1997). Further, as the author notes:

...elite players would forego the opportunity to use one newaza skill in favor of something else. i.e. Pass-up a hold down to work for an arm lock. This should not be surprising! Players seek their favorite throws in spite of the opportunity for another throwing skill on a regular basis. Why shouldn't a player prefer one type of newaza over another?

Just like chess players, speed cubers and mnemonists, expert judoka work to create opportunities in which they can execute a small set of highly automatised techniques. This does not mean we should deny that competitive judo involves improvisation. Rather, we should recognise that improvisation involves considerable recourse to patterns stored in long term memory.

The domain-size argument is supposed to undermine the default assumption in psychology that expert skill is characterised by a high degree of automaticity. It fails to do so, however.

For an expert to properly count as flexible, she must be able to perform at a level far higher than the general population in novel circumstances. As I noted in the discussion of speed-cards above, however, there is some ambiguity with regards to what counts as 'novel'. If we are sufficiently fine-grained in how we differentiate circumstances of performance, then of course every performance takes place in novel circumstances. To borrow from Heraclitus, no expert performs in the same river twice. From this vantage point, the extraordinary ability of experts to discover and exploit even small regularities in their domains might be seen as facilitating flexibility in different situations, rather than displacing it. Is it possible, then, that my disagreement with Christensen et al. (2019) is merely verbal? I think not, for the following reason: how fine-grained we go is not merely a matter of preference. We need to be able to provide reasons for focusing on flexibility at particular levels. In the next section I provide such reasons,

arguing that the primary drivers of expertise manifest at the level of technique and not at a situational level.

3.5 Levels of control

I've argued that expert skill is highly automated at the level of technique, where automatic pattern recognition plays a central role. But expert performance often requires operating at several different scales. Christensen et al. (2019) describe a three-levelled hierarchy. In addition to the level of technique (which they call *implementation control*), they argue that experts maintain flexibility at what they call *strategic* and *situational* levels of control.¹⁰ Strategic control involves the governance of an extended course of action so that it achieves broader goals. In a mixed-martial arts match, for example, this might involve submitting an opponent with a chokehold or joint lock to win the match. Situational control involves determining what actions need to be performed in the immediate situation in order to achieve or move towards strategic goals. To continue with the above example, this might involve throwing an opponent to the mat so they can establish positional control and work to execute their submission techniques. I agree with Christensen et al. (2019) that experts maintain control at these levels. But performance at the strategic and situational level is not what makes the difference between experts and novices under normal conditions.

In 1998, Gary Kasparov organised the first 'cyborg chess' tournament, in which humans paired with chess computers. As Epstein (2019) writes:

Years of pattern study were obviated. The machine partner could handle tactics so the human could focus on strategy...In chess, it changed the pecking order instantly...Kasparov settled for a 3–3 draw with a player he had trounced four games to zero just a month earlier in a traditional match. "My advantage in calculating tactics had been nullified by the machine." The primary benefit of years of experience with specialized training was outsourced, and in a contest where humans focused on strategy, he suddenly had peers. (p. 49)

Higher-level strategic thinking was not the primary driver of Kasparov's expertise. What set him apart from his peers were his skills at the technical level (referred to in chess as tactics). When his years of pattern training were neutralised, so too was a significant part of what made him an expert.

The kinds of heuristics that help navigate at the strategic and situational level are often comparatively easy to acquire. One of the frustrating things about learning a new skill is that we often know what we need to do—we just can't do it yet. Again, this is a consequence of brittleness; employing putatively flexible or general-purpose heuristics requires domain-specific knowledge and skill. There are at least two obstacles. The first is that we may lack the subject-specific knowledge required to recognise what kind of problem we are facing, and thus fail to recognise that a certain heuristic we possess is appropriate to the task in front of us. Here is a simple example: Imagine I want to take a trip from Canberra to Sydney, Australia, along a route of 250 km.

¹⁰ This hierarchy is to be understood as an idealisation. The number of levels may vary from competition to competition and it might not always be possible to differentiate them in practice.

Imagine, further, that I want to document my trip. I have enough room on my smart phone for 100 photos. If I want them evenly spaced, I should take one photo every how many kilometres? Many people, on first hearing this problem, divide the distance by the number of photos and to suggest that I take a photo every 2.5 kilometres. But if I take my first photo in Canberra (at km 0) then I'll take my last photo 2.5 km short of Sydney. Here is an easy way to see the error: suppose I only want to take two photos. Now how many kilometres should separate my photos? If I divide 250 by two, I get 125 km. I'd take my first photo in Canberra and my second only half-way along the trip. Thus, the formula to solve this problem is not number of kilometres divided by number of pictures. It's the number of kilometres divided by the number of pictures minus 1.¹¹

Someone who might get the problem above correct is someone who builds fences. Buying the right number of fenceposts is analogous to solving this problem. Solving a problem is often not a matter of critical thinking but of being able to recognize what kind of problem one is dealing with. Recognizing deep structural similarities requires deep knowledge of a subject. To be useful, recognitional abilities must also be practiced until they are fast enough to be deployed in real time.

Another reason merely possessing a heuristic is inadequate for expertise is that even if we recognise that a heuristic is appropriate, we may lack the subject-knowledge required to apply it. Suppose I want to evaluate the quality of a study on a particular intervention designed to improve math outcomes in the classroom. I have a powerful, general-purpose heuristic: I know that an important feature of a study is that it has an appropriate control group. The children in the control group should be the same in all relevant respects to those undergoing the intervention. But here is the problem: there are an almost innumerable number of ways in which the kids might differ that, at least to me, seem like they could plausibly be relevant to assessing their mathematical abilities (IQ, personality, highest level of education of parents, socio-economic status, ethnicity, time spent playing board games in early life with numerical concepts etc.). Not all these things will be controlled for in even the best study. Of course, that doesn't matter since some of these things are going to be much more relevant than others. But the only way I could know which factors the important ones to control for are is by having specialist knowledge of the literature around maths education. So, my putative general-purpose critical thinking heuristic doesn't work in the absence of a great deal of domain-specific skill and thus can't be readily applied.

The time has come to address an apparent tension in the account I've been developing. In arguing that expert skill is mediated by rich representations, I drew on research according to which expertise is acquired after many years of deliberate practice (Ericsson, Hoffman, Kozbelt, & Williams, 2018). But this body of research suggests that what differentiates experts is precisely that experts resist automaticity, allowing them to continue to shape and refine their skills. How can expert performance involve both more and less automaticity at the level of technique than novice performance? The answer is that although experts must be able to intervene during practice, they are more likely to rely on their finely tuned, automatized routines during performance. Hájek (2014) provides an instructive illustration of the distinction:

¹¹ This example was adapted from Willingham (2012).

John Searle tells a story of when he was a ski racer, and he had an Austrian coach to whom he turned for advice after doing a run. The coach's advice was simple: "Schneller!" ("Faster!"). The coach's point was that Searle should not overthink what he was doing. Instead of being preoccupied with his weight distribution or hand position, he should just think fast, and let his body do the rest. (p. 312)

This distinction, between practice and performance is something of an idealisation. Experts often work to simulate competition conditions *in* their practice. Given the limitations of transfer characteristic of expert brittleness, this makes sense: experts know that they will perform the way they train, a sentiment captured by a saying, common among judoka, that 'the best training for judo is judo'. Nonetheless, this sort of simulated performance comprises only a small part of expert practice. Ericsson et al. (1993) found that, among elite musicians, targeted solo practice, rather than time spent playing and performing songs (either alone or with other musicians) was the most significant driver of elite performance. This is a recurring theme in the research on deliberate practice. Musicians do not play scales on stage. Practice, in practice, differs from performance.

Marking the distinction between practice and performance also allows us to explain the seemingly-paradoxical observation that individuals who engage in more deliberate practice also experience greater levels of flow, a state characterised by complete absorption, effortlessness and spontaneity (Von Culin et al. 2014). Although the phenomenology of flow-like states has sometimes been appealed to as evidence that representational states couldn't possibly meet the demands of skilled action (Dreyfus and Dreyfus 2005), LTWM, with its hierarchically organised retrieval structures, was developed precisely to explain the extraordinary speed of expert reasoning. So, I don't take it as a challenge to the view I've developed here.

In this section, I've appealed to the notion of brittleness to defend a view of expertise according to which expert skill is representational and mindful, but still largely automated. In the next section I apply these ideas to debates concerning the nature of philosophical expertise.

4 Section 3: expert philosophers

4.1 Recent debates concerning philosophical expertise

Debates about philosophical expertise have recently taken centre stage in discussion of philosophical methodology. Experimental philosophers have discovered systematic differences between the responses of philosophers and non-philosophers about canonical cases and uncovered troubling sensitivities and biases in the judgements of professional philosophers (for a summary, see Machery 2017, chap. 2).

Some philosophers have argued that these findings undermine traditional philosophical methodology. Others have argued that we need not find such variation troubling (Hales 2009; Ludwig 2007; Williamson 2011). This variation is to be explained, the story goes, by philosophical expertise. Arguments for philosophical expertise generally appeal by analogy to other disciplines. Williamson (2011), for instance, invites us

to consider ‘the hypothesis that professional physicists tend to display substantially higher levels of skill in cognitive tasks distinctive of physics than/laypeople do. The hypothesis could be tested by systematic experiment. But even before that has happened, one can reasonably accept it’ (p. 220). Likewise, it’s reasonable to assume that philosophers also have expertise in their field, in the absence of experimental evidence. Devitt (2011) appeals to the expert intuitions of palaeontologists, psychologists and scientists. Ludwig (2007) makes an analogy to mathematical expertise. This argument—that differences between the case judgements of philosophers and those of the folk are to be accounted for in terms of the former’s expertise—is known as the ‘expertise defence’.

If expertise can be taken for granted in these other domains, then there is no need to worry about philosophical expertise. I take this argument seriously but note that the analogy often speaks *against* the assumption of expertise. There is a significant body of evidence demonstrating the surprising ways in which putative experts fail to outperform novices in their supposed domains of speciality: so-called “political experts” for example, including professors of political science, political journalists and professional politicians do no better than the average reader of the New York Times when it comes to making political predictions (Tetlock 2017). Stock-pickers, too, are subject to an illusion of expertise. The year to year correlation between the outcomes of mutual funds is barely higher than zero (Kahneman 2011). And among mental health professionals, clinical judgement remains the principle tool for predicting patient outcomes, despite meta-analyses of decades of studies speaking to the inferiority of clinical predictions to actuarial methods (Ægisdóttir et al. 2006; Grove et al. 2000).

In addition, there is empirical evidence that speaks directly against the expertise defence. Studies suggest that the training of professional philosophers does not appear to inoculate them against ordering effects (Schwitzgebel and Cushman 2012), actor/observer bias (Tobia et al. 2013) or personality effects (Schulz et al. 2011).

Although the debates prompted by experimental philosophy are new, concerns about philosophical expertise are not. Moral philosophers, in particular, have long worried about whether there are moral experts.

In this section, I will argue that the brittleness of expertise poses an important challenge to widely held views about what it takes to be an expert in philosophy. The structure is loosely historical; I begin by building a case against moral experts before extending my account to other philosophical domains. Naturally, my argument will also apply to philosophical expertise more broadly construed. If no-one meets the conditions required to be an expert in any sub-discipline of philosophy, then a fortiori they won’t meet the conditions required to be an expert in philosophy as a whole.

I will also argue, however, that taking the brittleness of expertise seriously opens a possible research avenue for advocates of the expertise defence. By recalibrating our understanding of the scope of philosophical expertise, we can refine the tools we use to look for it.

4.2 Moral expertise

Can someone be a moral expert in the same way that one can be an expert chemist or expert painter? Moral philosophy is analogous to these other disciplines in several ways: it is a professional discipline, entry into which generally requires many years of formal training. It has a distinctive set of methods and a considerable literature with which students must familiarise themselves. Like professional chemists, moral philosophers attend conferences and workshops and publish research findings. Like painters, moral philosophers develop their own style. A rare few even enjoy a degree of public recognition. In addition, as Singer (1972) points out, moral philosophers have a range of advantages over the laity when it comes to making moral judgements. As he writes, “Someone familiar with moral concepts and with moral arguments, who has ample time to gather information and think about it, may reasonably be expected to reach a soundly based conclusion more often than someone who is unfamiliar with moral concepts and moral arguments and has little time. So moral expertise would seem to be possible” (p. 117). Nonetheless, the prevailing mood has tended towards scepticism about moral expertise (Archard 2011; Coady 2012; Cowley 2005). Some of this scepticism has been fuelled by recent empirical findings. For instance, moral philosophers are no less subject to various biases in their judgements than non-philosophers (Machery 2017). Moral philosophers are also no more likely to behave morally than other people (Schönegger and Wagner 2019; Schwitzgebel and Rust 2016). Scepticism about moral expertise has also come from more traditional quarters, emphasising the widespread disagreement found between moral philosophers. Bambrough (1971), for instance, writes that moral philosophers “disagree so much and so radically that we hesitate to say that they are experts” (p. 164).

These arguments are each animated by different presuppositions about the nature of expertise. Arguments based on evidence of biases assume that expertise in the moral domain consists in superior capacities for judgement or discernment regarding cases. But philosophers plausibly do more than make case judgements or act as moral exemplars. They also make arguments, draw distinctions and construct theories. Similarly, arguments from disagreement assume a veristic view, according to which expertise consists in possessing a greater store of moral truths. Widespread disagreement entails that many putative experts will have systematically wrong beliefs but philosophers who want to analyse expertise in terms of superior abilities won’t be troubled by this.

My aim in this section is to offer an argument against moral expertise that works regardless of one’s theory of expertise.¹² The argument, roughly, is this: the domain of moral philosophy is broad and heterogenous. Because expertise is brittle, it does not transfer to different contexts. Thus, although someone can be an expert with regard to particular moral issues, it is unlikely that anyone could possess enough knowledge, or abilities (or whatever you think expertise consists in) to be a moral expert. I elaborate on each of these claims below.

The range of issues discussed by moral philosophers is vast. Topics range from the treatment of non-human animals to sexual ethics, just distribution of wealth, the

¹² A similar argument, but one which applies only to veristic accounts of expertise, can be found in Coady (2012).

ethics of cloning and so on. The growth of applied philosophy and the accompanying proliferation of methods and philosophical machinery speaks to how different these areas are. In addition to specialised methods, moral philosophers spend a great deal of time accumulating facts relevant to their areas; bioethicists must be able to consume and digest medical papers in addition to their usual diet of philosophy. Political philosophers must be similarly omnivorous, drawing not only from philosophy, but from politics and economics as well.

We might be tempted to think that there will still be an underlying, core set of concepts and theories that unite these areas. There may be some, but these will be too few to wholly constitute moral expertise. Firstly, many apparently similar concepts will vary in important ways. Take a philosophically important notion like ‘function’. There are almost certainly several different concepts of function that play different roles in different parts of moral philosophy. For instance, although Sterelny and Fraser (2016) appeal to an evolutionary notion of function in their defence of moral realism, such analyses are considered ‘unattractive’ in bioethics, since the biomedical sciences seem to employ a different conception of function (Reiss 2016). This makes sense; our concepts are tools and we fashion tools to be suitable to particular jobs. Secondly, even when philosophers want to apply their general theories to particular issues, they will still have to acquire a good deal of domain-specific knowledge (recall the discussion of general-purpose heuristics in Sect. 2).

Philosophers, can, of course, acquire the relevant methods and accumulate the relevant facts. And although it’s hard work, they can and do apply their general theories to specific moral issues. But this gives them claim only to expertise on those moral issues and not to moral philosophy at large.

4.3 Extending the argument

It should be obvious that much of what I’ve said about moral expertise will generalise to other areas of philosophy. Given the brittleness of expertise, it should not surprise us that experts in Gettierology are not always experts in formal epistemology. Further, if expertise doesn’t generalise from one issue to another in any sub-discipline of philosophy, then a fortiori it won’t generalise from one sub-discipline to another. Thus, it’s even less likely that anyone meets the standards required to be an expert on philosophy than that they meet the standards to be a moral expert, for example.

On the one hand, this is probably uncontroversial. Few philosophers would claim that anyone has ever been an expert on the whole of philosophy (except perhaps Aristotle, at his time). Philosophy is simply too large and motley a discipline for anyone to know everything. Despite this, there is a widespread conception among philosophers that expert philosophers have domain-general skills that can be fruitfully applied to any problem from aesthetics to Zeno’s paradoxes. These skills include things like spotting inconsistencies, reconciling seemingly incompatible views, seeing what follows from what and so on. This is also an image we try sell to our students: even if you don’t plan on being a philosopher, studying philosophy will help you learn to think critically and reason effectively, regardless of your future occupation. I suspect that part of what allows for this act of mental gymnastics is an implicit commitment

to the independence of skill from knowledge. As we saw in Sect. 2, however, such an independence does not obtain.

The problem is not merely that philosophical knowledge is highly circumscribed. I doubt anyone expects that an expert on Kantian ethics should also be an expert on formal epistemology. The problem is that having expert level skills *depends* on the possession of relevant knowledge. As noted in Sect. 2.3, according to LTWM, expert skill is the product of acquired mental processes. These processes, in turn are the result of highly organised knowledge schemas. So, we can't draw a sharp distinction between expert knowledge and expert skill. This, on its own, does not rule out that there might be some kind of knowledge that can provide general-purpose reasoning skills of the kind philosophers often take themselves to have. But I think it's very unlikely for the reasons I provided in Sect. 2.5. High level knowledge is rarely what drives expert skill. I provided two reasons why this might be the case. The first is that we may lack the subject-specific knowledge required to recognise what kind of problem we are facing, and thus fail to recognise that a certain heuristic/piece of higher level knowledge we possess is appropriate to the task in front of us. The second is that even if we recognise that a heuristic is appropriate, we may lack the subject-knowledge required to apply it. So, I don't believe that philosophers do possess expert level domain-general skills.

How, then, do I explain Ganeri's (2008) observation that 'no matter how different their areas of specialization are, philosophers can and do talk to each other'? The first thing to say is that, insofar as this observation is true, it needn't be explained by domain-general skills. None of what I've said above is incompatible with the fact that philosophers often work productively in multiple areas of philosophy. Nor is it incompatible with the observation that philosophers sometimes import ideas and arguments from one domain to another (see Hájek (2016) for compelling examples). However, when they can do so, it's because they've spent time developing new domain-specific knowledge and skills.

The second thing to say is that although casual reflection may seem to support the claim that philosophers are able to ask probing questions, generate counter examples and alternative explanations with regard to issues they know little about, personal reflection is not the right way to establish this claim. This is a case of marking our own homework. Our memories are unreliable and prone to bias. For instance, we are much more likely to remember the great questions and arguments than the ones that fall flat. Further, Ganeri's narrative is a self-serving one and so may be subject to confirmation bias. This point can be made vivid with an analogy. Suppose you are a medical doctor and you want to know what proportion of people who contract a certain disease make a full recovery. An obviously bad way of working this out would be to reflect on how many people you know who have contracted this disease and count how many you think made a full recovery. A better way would be to collect a random sample from the afflicted population, take systematic measurements and perform the appropriate statistical analysis. Similarly, someone looking to provide evidence for the view that philosophers possess expert-level domain-general skills would do well to record all the questions asked in a sample of seminars and then to check what proportion of them were asked by people without domain-specific expertise.

To clarify, I am not arguing that philosophers lack *any* kind of expertise. In fact, my argument relies on it being the case that philosophers *do* possess genuine expertise

with regards to particular issues. If they didn't possess expertise of any kind, then evidence for the brittleness of expertise would be irrelevant when trying to understand the shape and limitations of philosophical activity.

Proper appreciation of the brittleness of expertise also raises interesting questions about experimental evidence showing the unreliability of philosophical judgement. It may be that one reason we've failed to find evidence of expertise regarding case judgements is that we've treated philosophers as a more or less undifferentiated group. Schwitzgebel and Cushman (2012, 2015) show some sensitivity to these issues by testing philosophers who specialise in ethics, specifically, but even this may be too coarse grained. Given the considerable investment of time required to attain expertise, it may be that philosophers attain expertise regarding specific philosophical issues only at the cost to their abilities in other areas. It's interesting, for instance, that philosophers inside a particular subfield of philosophy sometimes hold quite different views from the profession at large (Bourget and Chalmers 2009). Future research should examine whether specialists on particular issues, as opposed to areas, are subject to the same influence as non-specialists. On the other hand, brittleness imposes a significant constraint on any version of the expertise reply since, even if individual expert intuitions are reliable with regard to a narrow set of questions, it does not mean we can assume that they will be reliable more broadly. Philosophers interested in defending the armchair methods of their discipline may, ironically, find the best resources to make their case in the tools and findings of experimental philosophy.

Whether or not philosophers can lay claim to any kind of expertise is, of course, a matter to be settled by empirical research. However, because of the brittleness of expertise, we should not expect that a philosopher who can expertly draw distinctions, construct theories or provide guidance on one issue to perform as well in regard to another issue with which they have no special experience. If this is the bar, we wish to set for ourselves, it may turn out that there are no experts in philosophy.

5 Conclusion

That expertise is brittle in surprising ways has thus far been underappreciated by philosophers. In this paper I have argued that expertise generally doesn't transfer to novel tasks, even ones that are intuitively similar. In addition, experts are often unable to flexibly cope with changes to their domains or to novel parts of their domains. After outlining the evidence that expertise is brittle, I argued that brittleness gives us reasons to reject the domain-size argument against automaticity of expert skill and advances the dialectic in debates on the nature of expert skill more generally. In Sect. 3, I used the evidence for brittleness to develop a novel argument against philosophical expertise. Because philosophy is broad and heterogenous, it is unlikely that expertise on one topic will transfer to other topics in that area. Thus, even if philosophers can claim expertise regarding certain issues, it is extremely unlikely that there are moral experts, expert epistemologists and so on. This may, however, offer new hope for advocates of the expertise defence.

Acknowledgments Many thanks to Alan Hájek who provided invaluable feedback on several drafts of this paper. Thanks also to John Sutton, under whose guidance I originally developed many of the ideas in Sect. 2 of this paper. I am also indebted to several colleagues at the ANU for their insightful discussion: Don Nordblom, James Willoughby and Ross Pain. Finally, thanks to two excellent anonymous referees at Synthese. The final version of this paper was much improved by their thorough and thoughtful criticism.

Funding Not applicable.

Compliance with ethical standards

Conflict of interest The authors declare that they have no conflict of interest.

References

- Ægisdóttir, S., White, M. J., Spengler, P. M., Maugherman, A. S., Anderson, L. A., Cook, R. S., et al. (2006). The meta-analysis of clinical judgment project: Fifty-six years of accumulated research on clinical versus statistical prediction. *The Counseling Psychologist*, 34(3), 341–382.
- Archard, D. (2011). Why moral philosophers are not and should not be moral experts. *Bioethics*, 25(3), 119–127.
- Bambrough, R. (1971). Plato's political analogies. In *Plato* (pp. 187–205). Springer.
- Beilock, S. L., Carr, T. H., MacMahon, C., & Starkes, J. L. (2002). When paying attention becomes counterproductive: impact of divided versus skill-focused attention on novice and experienced performance of sensorimotor skills. *Journal of Experimental Psychology: Applied*, 8(1), 6.
- Bilalić, M., McLeod, P., & Gobet, F. (2008). Inflexibility of experts—Reality or myth? Quantifying the Einstellung effect in chess masters. *Cognitive Psychology*, 56(2), 73–102.
- Bourget, D., & Chalmers, D. (2009). The PhilPapers surveys: results, analysis and discussion. Retrieved from *PhilPapers website* <http://philpapers.org/surveys>.
- Charness, N. (1976). Memory for chess positions: Resistance to interference. *Journal of Experimental Psychology: Human Learning and Memory*, 2(6), 641.
- Chase, W. G., & Simon, H. A. (1973). Perception in chess. *Cognitive Psychology*, 4(1), 55–81.
- Christensen, W., Sutton, J., & Bicknell, K. (2019). Memory systems and the control of skilled action. *Philosophical Psychology*, 32(5), 693–719.
- Clayton, N. S., & Dickinson, A. (1998). Episodic-like memory during cache recovery by scrub jays. *Nature*, 395(6699), 272–274.
- Clayton, N. S., Emery, N., & Dickinson, A. (2006). The rationality of animal memory: Complex caching strategies of western scrub jays. In Hurley (Ed.), *Rational animals?* (pp. 197–216).
- Coady, D. (2012). *What to believe now: Applying epistemology to contemporary issues*. Hoboken: Wiley.
- Cowley, C. (2005). A new rejection of moral expertise. *Medicine, Health Care and Philosophy*, 8(3), 273–279.
- Crichton, M., 2020. *Why Speculate?*. [online] Geer.tinho.net. Retrieved 19 August, 2020 from <http://geer.tinho.net/crichton.why.speculate.txt>.
- Devitt, M. (2011). Methodology and the nature of knowing how. *The Journal of Philosophy*, 108(4), 205–218.
- Dreyfus, H. L., & Dreyfus, S. E. (2005). Peripheral vision: Expertise in real world contexts. *Organization Studies*, 26(5), 779–792.
- Epstein, D. (2019). *Range: Why generalists triumph in a specialized world*. London: Riverhead Books.
- Ericsson, K. A. (2006). Protocol analysis and expert thought: Concurrent verbalizations of thinking during experts' performance on representative tasks. *The Cambridge handbook of expertise and expert performance*, 223–241.
- Ericsson, K. A., Hoffman, R. R., Kozbelt, A., & Williams, A. M. (2018). *The Cambridge handbook of expertise and expert performance*. Cambridge: Cambridge University Press.
- Ericsson, K. A., & Kintsch, W. (1995). Long-term working memory. *Psychological Review*, 102(2), 211.
- Ericsson, K. A., Krampe, R. T., & Tesch-Römer, C. (1993). The role of deliberate practice in the acquisition of expert performance. *Psychological Review*, 100(3), 363.

- Ericsson, K. A., & Lehmann, A. C. (1996). Expert and exceptional performance: Evidence of maximal adaptation to task constraints. *Annual Review of Psychology*, 47(1), 273–305.
- Ericsson, K. A., & Simon, H. A. (1984). *Protocol analysis: Verbal reports as data*. Cambridge: The MIT Press.
- Ericsson, K. A., & Staszewski, J. J. (1989). Skilled memory and expertise: Mechanisms of exceptional performance. *Complex Information Processing: The Impact of Herbert A Simon*, 2, 235–267.
- Feltovich, P. J., Prietula, M. J., & Ericsson, K. A. (2006). Studies of expertise from psychological perspectives. *The Cambridge handbook of expertise and expert performance* (pp. 41–67). New York: Cambridge University Press.
- Gallistel, C. R. (2008). Learning and representation. *Learning and Memory: A Comprehensive Reference*, 1, 227–242.
- Ganeri, J. (2018). *Taking Philosophy Forward*. Retrieved from <https://lareviewofbooks.org/article/taking-philosophy-forward/>.
- Gobet, F. (1998). Expert memory: A comparison of four theories. *Cognition*, 66(2), 115–152.
- Gobet, F., & Simon, H. A. (1996). Templates in chess memory: A mechanism for recalling several boards. *Cognitive Psychology*, 31(1), 1–40.
- Gobet, F., & Simon, H. A. (2000). Five seconds or sixty? Presentation time in expert memory. *Cognitive Science*, 24(4), 651–682.
- Gobet, F., & Waters, A. J. (2003). The role of constraints in expert memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 29(6), 1082.
- Grove, W. M., Zald, D. H., Lebow, B. S., Snitz, B. E., & Nelson, C. (2000). Clinical versus mechanical prediction: A meta-analysis. *Psychological Assessment*, 12(1), 19.
- Hájek, A. (2014). Philosophical heuristics and philosophical creativity. *The Philosophy of Creativity: New Essays*, p. 288.
- Hájek, A. (2016). Philosophical heuristics and philosophical methodology. *The oxford handbook of philosophical methodology*, pp. 348–373.
- Hales, S. D. (2009). *Relativism and the Foundations of Philosophy*. Cambridge: MIT Press.
- Hutto, D. D., & Myin, E. (2017). *Evolving enactivism: Basic minds meet content*. Cambridge: MIT press.
- I. Lorimer (Director). (2012). Jumble [Television series episode]. QI, QI.
- Jugs Sports. 2020. A Women'S Softball Pitcher Vs. The Top Baseball Hitters...Who Wins?. [online]. Retrieved 1 August, 2020 from <https://jugssports.com/blog/a-womens-softball-pitcher-vs-the-top-baseball-hitterswho-wins/>.
- Kahneman, D. (2011). *Thinking, fast and slow*. New York: Macmillan.
- Lewandowsky, S., & Thomas, J. L. (2009). Expertise: Acquisition, limitations, and control. *Reviews of Human Factors and Ergonomics*, 5(1), 140–165.
- Ludwig, K. (2007). The epistemology of thought experiments: First person versus third person approaches. *Midwest Studies in Philosophy*, 31(1), 128–159.
- Machery, E. (2017). *Philosophy within its proper bounds*. Oxford: Oxford University Press.
- Maguire, E. A., Valentine, E. R., Wilding, J. M., & Kapur, N. (2003). Routes to remembering: the brains behind superior memory. *Nature Neuroscience*, 6(1), 90–95.
- Marchant, G., Robinson, J., Anderson, U., & Schadewald, M. (1991). Analogical transfer and expertise in legal reasoning. *Organizational Behavior and Human Decision Processes*, 48(2), 272–290.
- Matthewson, J., & Weisberg, M. (2009). The structure of tradeoffs in model building. *Synthese*, 170(1), 169–190.
- Menzel, R., Greggers, U., Smith, A., Berger, S., Brandt, R., Brunke, S., et al. (2005). Honey bees navigate according to a map-like spatial memory. *Proceedings of the National Academy of Sciences USA*, 102(8), 3040–3045.
- Miller, G. A. (1956). The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological Review*, 63(2), 81.
- Noë, A. (2005). Against intellectualism. *Analysis*, 65(4), 278–290.
- Park, J., MacRae, H., Musselman, L. J., Rossos, P., Hamstra, S. J., Wolman, S., et al. (2007). Randomized controlled trial of virtual reality simulator training: transfer to live patients. *The American Journal of Surgery*, 194(2), 205–211.
- Reiss, J. A., Rachel. (2016, 2016). Philosophy of Medicine. *The Stanford Encyclopedia of Philosophy*. Summer 2016. Retrieved from <https://plato.stanford.edu/archives/sum2016/entries/medicine/>.
- Rescorla, M. (2012). Millikan on honeybee navigation and communication. In D. Ryder, J. Kingsbury, & K. Williford (Eds.), *Millikan and her critics*. Wiley.

- Ryle, G. (2009). *The concept of mind*. Abingdon: Routledge.
- Saariluoma, P. (1990). Apperception and restructuring in chess players' problem solving. *Lines of Thought: Reflections on the Psychology of Thinking*, 2, 41–57.
- Saariluoma, P. (1991). Aspects of skilled imagery in blindfold chess. *Acta Psychologica*, 77(1), 65–89.
- Schönegger, P., & Wagner, J. (2019). The moral behavior of ethics professors: A replication-extension in German-speaking countries. *Philosophical Psychology*, 32(4), 532–559.
- Schulz, E., Cokely, E. T., & Feltz, A. (2011). Persistent bias in expert judgments about free will and moral responsibility: A test of the expertise defense. *Consciousness and Cognition*, 20(4), 1722–1731.
- Schwartz, A., & Drayson, Z. (2019). Intellectualism and the argument from cognitive science. *Philosophical Psychology*.
- Schwitzgebel, E., & Cushman, F. (2012). Expertise in moral reasoning? Order effects on moral judgment in professional philosophers and non-philosophers. *Mind and Language*, 27(2), 135–153.
- Schwitzgebel, E., & Cushman, F. (2015). Philosophers' biased judgments persist despite training, expertise and reflection. *Cognition*, 141, 127–137.
- Schwitzgebel, E., & Rust, J. (2016). The behavior of ethicists. *A companion to experimental philosophy*, p. 225.
- Sims, V. K., & Mayer, R. E. (2002). Domain specificity of spatial expertise: The case of video game players. *Applied Cognitive Psychology: The Official Journal of the Society for Applied Research in Memory and Cognition*, 16(1), 97–115.
- Singer, P. (1972). Moral Experts. *Analysis*, 32(4), 115–117. <https://doi.org/10.2307/3327906>.
- Stanley, J., & Willamson, T. (2001). Knowing how. *The Journal of Philosophy*, 98(8), 411–444.
- Sterelny, K., & Fraser, B. (2016). Evolution and moral realism. *The British Journal for the Philosophy of Science*, 68(4), 981–1006. <https://doi.org/10.1093/bjps/axv060>.
- Sternberg, R. J., & Frensch, P. A. (1992). On being an expert: A cost-benefit analysis. In *The psychology of expertise* (pp. 191–203). Springer.
- Swain, R., (1928). Quoted in 'knows much about little: that is one definition given of scientist by chemist'. *The Ogden Standard-Examiner*, [online] p.1. Retrieved 19 August, 2020 from <http://Newspapers.com>.
- Tetlock, P. E. (2017). *Expert political judgment: How good is it? How can we know?—New edition*. Princeton: Princeton University Press.
- Tobia, K., Buckwalter, W., & Stich, S. (2013). Moral intuitions: Are philosophers experts? *Philosophical Psychology*, 26(5), 629–638.
- Tonneau, F. (2011/2012). Metaphor and truth: A review of representation reconsidered by W. M. Ramsey. *Behavior and Philosophy*, 39(40), 331–343.
- Van Sickle, K. R., Ritter, E., McClusky, D., Lederman, A., Baghai, M., Gallagher, A., et al. (2007). Attempted establishment of proficiency levels for laparoscopic performance on a national scale using simulation: The results from the 2004 SAGES Minimally Invasive Surgical Trainer—Virtual Reality (MIST-VR) learning center study. *Surgical Endoscopy*, 21(1), 5–10.
- Von Culin, K. R., Tsukayama, E., & Duckworth, A. L. (2014). Unpacking grit: Motivational correlates of perseverance and passion for long-term goals. *The Journal of Positive Psychology*, 9(4), 306–312.
- Wanzel, K. R., Hamstra, S. J., Anastakis, D. J., Matsumoto, E. D., & Cusimano, M. D. (2002). Effect of visual-spatial ability on learning of spatially-complex surgical skills. *The Lancet*, 359(9302), 230–231.
- Weers, G., 1997. *Skill Range Of The Elite Judo Competitor* | *Judo Info*. [online] [judoinfo.com](https://judoinfo.com/weers1/). Retrieved 19 August, 2020 from <https://judoinfo.com/weers1/>.
- Williamson, T. (2011). Philosophical expertise and the burden of proof. *Metaphilosophy*, 42(3), 215–229.
- Willingham, D. T. (2012). *When can you trust the experts?: How to tell good science from bad in education*. Hoboken: Wiley.
- Yntema, D. B., & Mueser, G. E. (1960). Remembering the present states of a number of variables. *Journal of Experimental Psychology*, 60(1), 18.