

Intellectual Skill and the Rylean Regress

Brian Weatherston

In recent work about know how, Rylean regress arguments have largely dropped out of focus. They play little role in the anti-intellectualist arguments of various kinds in the papers collected in Bengson and Moffett (2011). And they are used as something like target practice by intellectualists like Stanley (2011), who uses the first chapter of his book to dispose of them before getting onto the real business. The majority view seems to be that Ginet (1975) basically showed these arguments didn't work, and it's time to move onto other considerations for or against intellectualism.

I think this isn't right. In particular, I think Rylean regress arguments can show that some kinds of intellectualism are false, and others are undermined. In particular, I'll argue that the Rylean regress arguments show that there is such a thing as *intellectual skill*, and this can't be identified with propositional knowledge.

The position I'm defending is a long way from the strongest kinds of Rylean position that Stanley spends most of his book arguing against. In particular, the view I'm defending is consistent with the following four theses.

1. Instances of intellectual skill are usually, and perhaps always, not happily reported using know-how ascriptions.
2. Know-how ascriptions are rarely, if ever, reports of intellectual skill, and are frequently reports of propositional knowledge.
3. Intellectual skill is guided by, and dependent on, propositional knowledge.
4. Propositional knowledge is not behaviourally inert.

Not just is the view consistent with these four, I'm fairly confident that the last three at least are true. But that's all consistent with the view that intellectual skill is not itself propositional knowledge. And it's all consistent with the view that we can learn philosophically significant conclusions from Ryle's regress arguments.

One disclaimer before I start. Although this paper is heavily influenced by Ryle (1945, 1949), and sympathetic interpreters of Ryle such as Hornsby (2011), I make no attempt at Ryle exegesis here. I think there's a decent case to be made that Ryle was sympathetic to the position defended here, but I'm going to leave that debate for another day.

1 Stanley's Version of the Regress

Stanley offers the following helpful formulation of the premises of the Rylean regress argument.

1. The intellectualist view entails that “for any operation to be intelligently executed”, there must be a prior consideration of a proposition.
2. “The consideration of propositions is itself an operation the exercise of which can be more or less intelligent, less or more stupid.” (Stanley, 2011, p. 12)

Stanley isn't quite as clear what he takes the conclusion of the regress argument to be. At times he takes it the argument to be “directed against the intellectualist view that an action is intelligent if and only if it is guided by propositional knowledge” (Stanley, 2011, p. 2) That's not what I'm taking the conclusion of the argument to be. Indeed, I think that intellectualist view is true, and is consistent with everything we learn from the regress argument.

My primary conclusion is a little more modest. It is that intelligent action must commence with something that is not identifiable with propositional knowledge, something I'm calling intellectual skill. Intellectual skill will usually be guided by propositional knowledge; it just won't be identified with it.

But despite our differences in what we take the argument to purport to show, I think Stanley's formulation of the argument is itself rather helpful. That's because his discussion of the ways to understand the phrase “considering a proposition” usefully brings out the possible interpretations of the argument.

It's true, as Stanley immediately notes, that premise one is wildly implausible if we understand consideration to require conscious awareness. Stanley notes, following Carl Ginet (1975), that we can skilfully open a door without consciously considering what we know about door-opening. On that interpretation, premise 1 is just false.

Stanley then considers an alternative reading of the argument, according to which considering a proposition should just be understood as triggering a representation. On this view, premise 1 seems plausible enough. If intelligent action is driven by propositional knowledge, then it must be preceded by the triggering of a representation of that knowledge. Stanley's objection is to premise 2. He writes

Triggering a representation can be done *poorly* or *well*. But this does not show it can be done *intelligently* or *stupidly* ... Since triggering representations is something we do automatically, replacing “contemplating a proposition” by “triggering a proposition” in Premise

2 of the regress argument results in a manifest implausibility. (Stanley, 2011, p. 16)

This hardly seems right to me. Let's consider two examples that suggest it isn't correct.

Billy asks Suzy whether she thinks Jill's party will be a success. There are a lot of things that are common knowledge between the two of them. One is that Jill is a proficient party host. Another is that Jill has invited all of their colleagues, including Jack. Another is that parties which Jack attends are rarely successes. But Suzy thinks for a minute, remembers that Jack is away in Ohio, and says that it will be a success.

It was smart of Suzy to think about Jack's whereabouts. It wasn't, perhaps, necessary. If she'd just reasoned from Jill's general proficiency to the success of the party, she would have got to the right conclusion. But it was better to note a possible complication, and check that it wouldn't actually get in the way.

It would have been stupid to perform the same activity for many other kinds of possible complications. If Suzy had thought to herself, "The party will be a disaster if there's an alien invasion in the middle of it, but there's no reason to think the aliens will invade just now, so I'll keep on thinking it will be a success," that would have been stupid. Other possible complications are not stupid to consider, but they are intellectual mistakes. The party won't be a success if there's a police raid in the middle, based on a mistaken view the police have about where a particular drug dealer lives. Police do make mistakes, so even if Jill isn't a drug dealer, this could be a genuine concern, depending on how nearby the mistakes are. But if the nearest mistake was a botched raid in a neighbouring state in the previous year, it's wrong for Suzy to worry about this before answering Billy's question.

Stanley's view has to be that I've been misusing adjectives systematically through the last two paragraphs. I shouldn't have said that it was smart of Suzy to consider Jack's whereabouts, or that it would have been stupid to consider the alien invasion. Rather, it was just her cognitive system working well when she considered Jack, and would have been working poorly had she considered the aliens, and sub-optimally had she considered the police. This doesn't seem at all the natural way to describe the case to me, in part because I'm not sure I see the difference Stanley is hinting at. Intelligence just is the good operation of the cognitive system, and stupidity its poor operation.

But we haven't considered Stanley's argument from automaticity. Let's run through a different example to see why that doesn't work. Consider the following fictional exchange.¹

Oscar Wilde: I wish I'd said that.

James McNeill Whistler: You will Oscar; you will.

At least in the fiction, that's a clever response. Had someone, implausibly, come up with it on the spot in real life, it would have been clever too. It's an occupational hazard of philosophers to think that the ability to come up with quick, clever responses is somehow central to intelligence. But we can reject that wildly implausible view without thinking that it's wrong to think of these quips as a manifestation of a kind of intelligence.

Now let's think of how someone could have come up with this response. Even before we start researching the neural patterns behind quips like this, we can be pretty sure the following is not what happened in a speaker's brain. They first make an exhaustive list of all possible responses, from "Green ideas sleep furiously" to what is actually said, then figure out which would be best, then produce the best one. On this wildly implausible model, the reply would be intelligent because it would reflect the speaker's ability to properly evaluate this list of responses. That's implausible because the list is simply too big. Indeed, it is in principle infinitely large. The list is too big to survey not just consciously, but subconsciously.

Coming up with a response like this requires first coming up with a narrower list of possible responses, and then evaluating which is best from that list. There's a romantic model of intellect where the list in question consists of just the reply actually issued. On this model the perfect reply appears fully and perfectly formed in the mind of the intelligent person. Now such a model may often fit the phenomenology, but I don't think we should give that much credit. It's an empirical question how many possible replies are represented in the mind in a situation like this, before the chosen reply is issued. What's not an empirical issue is whether the list of possible replies that is represented in the mind is finite or infinite. It simply must be finite, which means that there must be better

¹The exchange is from a Monty Python sketch, from episode 13 of season 3 of *Monty Python's Flying Circus*, first broadcast on January 18, 1973. The text of the sketch is reproduced in Dempsey (2012, p. 741). It's actually striking how few really good off-the-cuff quips there are in recorded history. The famous one attributed to Wilde, "I have nothing to declare but my genius", is probably apocryphal, and in any case sounds prepared. Lists of famous come-backs and ripostes are usually crowded with written responses. (Or with fakes, such as this one.) Word play is hard.

and worse lists to consider. And that suggests that there is some skill involved in coming up with the list.

One could reject this last conclusion. One could try saying that the coming up with a list of possible replies is no manifestation of skill, but the skill is only involved in the evaluation and selection of replies. But this seems to generate a bizarre explanation about why the less skilled interlocutor comes up with worse replies. The model, presumably, is that the lack of skill does not explain having the wrong list of replies to choose between. Rather, what explains their less skilled reply must simply be that they misevaluated the possible replies. But that doesn't fit the observed data. It's much easier to see of someone else's reply that it was clever than it is to come up with a clever reply.

It could also be objected that the model I've suggested is much too simple. It isn't just that the mind issues a list of options, then evaluates them, and then selects the best. A more plausible model involves more recursive steps. The mind first generates a list of options, selects the best, then generates a list of refinements of that best option, selects the best of those, and so on. Perhaps when we consider superficial forms of intelligence, such as quips, it makes sense to consider a 'one-step' model, where a list is generated and evaluated, followed by a speech. But when one is choosing one's words carefully, as in say Wilde's writing, the simple model I've described feels much too simple.

But although the simple model is too simple for considered writing, the general structure must be right. Even a writer working at a leisurely pace, such as Joyce taking decades to write *Finnegans Wake*, doesn't have time to consider, even subconsciously, all possible constructions. There are still too many. And nor is it true that the difference between Joyce's skill and ours is that he realises the value of the sentences we all represent. The rest of us didn't simply misjudge the value of "Nobirdy aviar soar anywing to eagle it" (Joyce, 1939/2012, p. 505); we simply didn't token it. The ability to token mental representations like that is part of what Joyce's genius consists in.

So there are ample reasons to reject the incompatibility of automaticity and intelligence that Stanley's argument relies on.² It's true that we can't consider whether to token something before tokening it. It doesn't follow that what we token is independent of our intelligence; adopting the right starting points in an inquiry is one manifestation of intelligence.

I've focussed so far on cases where it is a priori implausible that human thinkers start by surveying the range of possible things they could do. It is also

²More generally, there are reasons to resist the idea that the involuntary is outside the scope of normative assessment. See Adams (1985), Smith (2005) and, for some applications more directly to epistemology, Ryan (2003), Steup (2008), and Weatherston (2008).

interesting to look at cases where this is in principle possible, but doesn't seem to happen in practice. There have been, traditionally, major differences in the style of play between human and computer chess players. (Since so many young players learn from machines these days, Kasparov (2010) suggests these differences are diminishing.³) This isn't necessarily because humans can't consider all options on the chess board. Usually there will be fewer than a hundred available moves, and a human could consider each. But that isn't, it seems, how humans think. They don't allocate equal resources to working through each of the possible options. As a result, computers often come up with surprising kinds of moves. Now computers are actually very good at chess, so these pre-deliberative allocations of cognitive resources may not have been optimal. Perhaps it would have been better for traditional chess players to spend more time thinking through unlikely progressions of the game. But it is evidence that even when we could use an unintelligent method for beginning inquiry, namely recursively generating the possible options, we prefer to use intelligent methods.

2 Skill and Knowledge

In the previous section I've argued that there is such a thing as intellectual skill; it is crucially involved in the selection of options to seriously consider. It is this skill that stops the Rylean regress. I take this to show that there is something beyond knowledge that is crucially important to intelligent life. In this section I'll respond to the most pressing objection to this conclusion, namely that the skill in question just is a kind of knowledge.

What knowledge could this skill be? One natural thought is that it is knowledge that such and such consideration is *relevant* to the current deliberation. I mean to use the term 'relevant' here in a fairly generic way; someone who thinks that what I'm calling a skill is really a kind of knowledge is surely not committed to the view that the knowledge in question is precisely articulated in English using the ordinary term 'relevant'. But presumably it is something like this; after all, skilled thinkers do on the whole consider relevant options and considerations, and unskilled ones either overlook relevant ones, or consider irrelevant ones.

One bad argument against the identification of skill with knowledge starts with the idea that skills are active, while knowledge is passive. The thought is that the person who knows a lot is like the Tortoise in Louis Carroll's dialogue (Carroll, 1895), only able to add more premises but never to reach a conclusion. It is only with skill that we can get to the conclusion. Stanley rightly objects to this argument on the grounds that it just isn't true that knowledge is passive in

³Thanks to Bernard Kobes and John Collins for helpful discussions about the chess examples.

the relevant sense. We should not, as Stanley puts it, “over-intellectualize knowing that”. (Stanley, 2012, p. 773). (A similar point is made in Stalnaker (2012).) Knowing that p is not just a matter of having p written in a knowledge box somewhere in the brain; it can be in part be constituted by active dispositions.

A better argument looks at the very different modal profiles of intellectual skill and knowledge of relevance. Someone can know that something is irrelevant and yet lack the skill to ignore it; or they can know that something is relevant and yet lack the skill to consider it in a timely manner. Examples from the other direction, where there is skill without knowledge, are a little more contentious, but we’ll look at some possible cases of those too. But first we’ll run through two examples to show how easy it is to have knowledge without skill.

Alice has spent a lot of money on video-conferencing equipment. But it isn’t working at all well, and she now has to decide whether to try and patch it into something better, or buy a whole new system. She knows the sunk cost fallacy is a fallacy; that buying a new system would make the previous purchases a waste is no reason to not buy a new system, especially if doing so is good value compared to the cost of buying a ‘patch’. But she can’t bring herself to ignore this fact when deliberating. Even though she eventually makes the right decision and buys new equipment, she takes much longer about this than she would have if, say, the existing equipment was old enough that she could easily conceptualise it as obsolete.

Bob is trying to solve a puzzle about the properties of functions from rationals to rationals. He knows that it is often helpful, when solving such puzzles, to transform the puzzle into one about functions from ordered pairs of integers to ordered pairs of integers. He knows that in the sense that if you asked him whether it could be useful to consider that transformation of the puzzle, he would immediately say yes, and this answer would come with the phenomenology of recollection, not of new insight. But no one does ask him that question, and the transformation in question simply never occurs to Bob. Since the untransformed puzzle is very hard, while the transformed puzzle is manageable, Bob never solves the problem.

It seems to me that what’s happened in both cases is that the agent has some knowledge, but is incapable of using it. What they lack is a skill. It’s true that knowledge isn’t completely passive. If Alice never appealed to the fact that the sunk cost fallacy is a fallacy in her reasoning, we wouldn’t say that she knows it. If none of Bob’s answers were guided by the existence of natural and useful transformations between rational numbers and ordered pairs of integers, we wouldn’t say he knows such transformations are natural and useful. I’m here

agreeing with Stanley and Stalnaker, and against Ryle, that knowledge that is itself a kind of disposition. And intellectual skill is a kind of disposition too. But they are very different dispositions. In particular, they have very different triggering conditions. Bob lacks some skill because he does not call to mind this fact about rational numbers right now. He has the salient knowledge about rational numbers because he is disposed to use the facts in question often enough.

So intellectual skill and knowledge of relevance have different manifestation conditions, and so they are not identical. But we can say something stronger than that. The cases of Alice and Bob are not in any way unusual. Examples where we forget the salience of some consideration, or can't get an irrelevant point out of our heads, are frequent. In principle, one could respond to the arguments I've made so far by saying that while knowledge of relevance is not identical to skill, nevertheless the two are as closely linked as, say, a material object and the matter that constitutes it. And if I had to resort to bizarre cases of the kind we torture introductory students with to make my point, I'd say that would be the right response. But given how normal Alice and Bob's cases are, this seems like the wrong move. Skill and knowledge don't just come apart in theory, they come apart in practice, frequently.

3 Three Objections

Objection: Appeal to skill does not stop the regress. If we need to posit something, say a skill, that comes between the possession of knowledge and the use of knowledge in reasoning or action, then we also to posit something that comes between the possession of a skill, and the use of that skill in reasoning or action. (Compare (Stanley, 2011, p. 26)).

Reply: What I'm going to say here is similar to what Jeremy Fantl (2011) said in a response to an earlier version of Stanley's argument, so I'll be brief. Skills are dispositions. We don't need to posit anything that comes between the disposition and its triggering. If a string is disposed to produce a middle C when struck, and it is struck, we don't need to posit an extra intermediary between the striking and the note. Dispositions stop regresses.

But, you might insist, couldn't the same be true of knowledge? After all, on a broadly functionalist construal of the mental, knowledge is a kind of disposition. My reply is in theory knowledge could stop such a regress, but in practice it is unlikely. An agent could be facing a problem where the possible considerations and options can be enumerated without using any particular skill, and the options are few enough that they can be each considered in turn. That is the situation an agent playing a relatively simple game might face. But it isn't the general human condition. In practice, we face problems every moment where it requires skill to

bring the right considerations to bear, at least given the processing capacities we have available.

Objection: Even if all of this is true, there may still be a sense in which intellectualism is true. After all, it could still be that knowledge guides action in a suitable way. (Compare (Stanley, 2011, p. 2).)

Reply: This could be true. Whether it is a win for intellectualism depends a bit on the boring question of how we settle the term ‘intellectualism’, and a bit on more interesting questions about priority. Let’s start by distinguishing five theories we might call intellectualist.

Identity Intellectualism - The possession of an intellectual skill just is the possession of a piece of knowledge.

Constitution Intellectualism - The possession of an intellectual skill is, always, constituted by a piece of knowledge.

Weak Constitution Intellectualism - The possession of an intellectual skill is, often, constituted by a piece of knowledge.

Causal Intellectualism - The possession of an intellectual skill is, always, caused by the possession of a piece of knowledge.

Weak Causal Intellectualism - The possession of an intellectual skill is, often, caused by the possession of a piece of knowledge.

This paper has been arguing against Identity Intellectualism. I think the falsity of this is as much as we could reasonably hope to prove using regress arguments. (I think I’m here agreeing with Wiggins (2009) and Hornsby (2011).) And the considerations behind the regress argument do, I think, show it to be false. If someone wants to insist that by intellectualism, they mean something weaker than this, I’m not going to quarrel over terminology. I’ll just note that Identity Intellectualism is an interesting, and false, thesis.

The arguments here are clearly not arguments against either form of Weak Intellectualism. Indeed, they are naturally understood as the kind of cases that confirm Weak Intellectualism. Mathematics students, like Bob, train by learning a lot of mathematical facts. And it’s hard to see how they could develop the relevant skills without knowing some important facts. This is, I suspect, the general case. Skillfully bringing the right considerations to bear on a problem requires, and is probably the causal consequence of, knowing a lot of relevant facts.

But what of the other two intellectualist theories? Do we have reason to think that there are some skills that are not constituted by, or not caused by, the possession of factual knowledge? One way to quickly show that would be to show that there can be skills without the related knowledge. Perhaps that's not just sufficient for rejecting Constitutive/Causal Intellectualism, but necessary. If knowledge without skills is possible, as in Alice and Bob's cases, and skills without knowledge were impossible, that asymmetry would call out for explanation. And something in the vicinity of Constitutive or Causal Intellectualism would be a very good candidate explanation.

There are (at least) two promising routes to showing that there can be skills without knowledge. One is due to Imogen Dickie (2012). She argues that since there are so many different routes to skill than there are to knowledge, we should expect that there will be cases of skill that are causally prior to knowledge. Jason Stanley (2012) replies that Dickie's argument assumes an overly narrow conception of propositional knowledge. This is a fascinating debate, but I don't have anything useful to add to it, so I'll just note the existence of this route, and move on.

The other route is due to Yuri Cath (2011). He suggests that know how isn't defeated by the same kind of considerations as know that is defeated. I'm going to sketch a Cath-style argument that we can have intellectual skills without knowledge. I think the argument has some force, though there are more ways to resist it than there are to resist the argument against Identity Intellectualism.

Ross and Rachel are economics students taking an exam. They are given a hard question asking about the likely effects of an exogenous shock, say an earthquake affecting an area the supplies crucial raw materials, on some related markets. The question is hard, with the relevant causal pathways being interconnected and often opposing. The only plausible way forward is to use a model and search for equilibrium points in the model. That's what Ross and Rachel have both been taught to do. And in fact both of them quickly select the right kind of model, with just the right amount of complexity in it to answer the question without being overburdened, and set out on the difficult algebra involved in solving the question.

So far it looks like both Ross and Rachel have shown intellectual skill. Now it turns out Ross and Rachel have very different views about the role of models in economic thinking. (My own thinking about models has been heavily influenced by (Strevens, 2008, ch. 8) and Davey (2011), and I rely on their insights in what follows.) These models involve, as all models do, some serious idealisation. Most notably, they assume that all the relevant actors are perfectly rational utility maximisers. Rachel hasn't given much thought to this assumption, though

she knows it to be literally false. But if pressed, she would say some reasonably sensible things about why she was using the model. For one thing, the familiar failures of human rationality aren't obviously relevant to the puzzle being presented. For another, they've been taught that using these models is a good way to solve problems, and that testimonial evidence carries some weight. And for another, it's an exam, and it is likely that questions have been selected to test how well students can use the models they have been taught. If those are her background, implicit, views, I think it is plausible to say that Rachel knows that the model is relevant to the exam question, even if she couldn't produce a theory of idealisations in economics of the standards of the best philosophers.

Ross's views about models are rather different. He thinks the familiar models in economics work, when they do, because the background assumptions are strictly and literally true. He thinks economic agents are utility maximisers, and the apparent evidence to the contrary is due to sloppy experimental design. He thinks markets are always in general equilibrium. And so he thinks that the only sources of error in predictions we can make about markets are from errors about things like the costs of extracting raw materials after the earthquake. This perspective is, of course, grossly mistaken. Moreover, Ross thinks that if the assumptions were not correct, there would be no point in using the models. This too is a mistake, though perhaps not as dramatic as his other mistakes.

Now even if Ross and Rachel aren't thinking about these philosophical views about the nature of models, I think they are relevant to whether each of them know that the models are relevant to the puzzle. In particular, I think Rachel does know that the models are relevant, while Ross's belief that they are relevant is more like a lucky guess than a piece of knowledge. Still, I think we should say that Ross showed skill in using this model rather than a more or less complex model, or a different kind of model, or no model at all. So he is a case of intellectual skill without knowledge of relevance.

I don't think this case is conclusive. I can think of at least four ways someone might reasonably object to the case.

1. It might be argued that despite his false views about why the models are relevant, he really does know that they are relevant. In other words, we would have another counterexample, to be added to those discussed by Warfield (2005) and Luzzi (2010), to the theory that false beliefs cannot generate knowledge.
2. It might be argued that Ross is not really skilled, since it is a matter of luck that the falsity of his beliefs does not lead him to false conclusions here.

3. It might be argued that although Ross doesn't know that this model is relevant, his skill is constituted by, or caused by, some other knowledge he has.
4. It might be argued that the broad picture of the role of idealisations in scientific reasoning that I'm adopting from Strevens and Davey is mistaken, and this fatally undermines my use of the case to argue against intellectualism.

I don't think these arguments are going to ultimately work. But it's clear we are a long way from Rylean regress arguments here. And that's where I think the debate about regress arguments should end. We have a good argument against Identity Intellectualism. And we have some suggestive considerations that seem to tell against Constitutive and Causal Intellectualism, but whether these arguments ultimately work will depend on considerations independent of the regress.

Objection: There are semantic arguments that attributions of know how are attributions of propositional knowledge. This shows that Ryle was wrong to draw a broad distinction between know how and know that.

Reply: I'm not making any claims about either know how or about 'know how'. I am making some claims about skill, and those imply some claims about 'skill'. But I'm sympathetic to the idea that reports of know how are often reports of some kind of practical propositional knowledge. I certainly haven't offered any arguments, nor I think any considerations in the direction of an argument, against this view.

Indeed, there are a lot of anti-intellectualist positions that I'm not arguing against here. Anti-intellectualism is often tied up with the view that there is an important distinction between theoretical and practical fields. The arguments I've developed here suggest that if there is such a distinction, then proving mathematical theorems is on the 'practical' side. I think that's a strange enough conclusion that it is time to change our terminology. That's why I've talked about the distinction between intellectual skills and knowledge, not the distinction (if such there is) between know how and know that, or between praxis and theory.

References

- Adams, Robert Merrihew. 1985. "Involuntary Sins." *Philosophical Review* 94:3–31.
- Bengson, John and Moffett, Marc (eds.). 2011. *Knowing How*. Oxford: Oxford University Press.
- Carroll, Lewis. 1895. "What the Tortoise Said to Achilles." *Mind* 4:278–280.

- Cath, Yuri. 2011. "Knowing How without Knowing That." In Bengson and Moffett (2011), 113–135.
- Davey, Kevin. 2011. "Idealizations and Contextualism in Physics." *Philosophy of Science* 78:16–38, doi:10.1086/658093.
- Dempsey, Luke (ed.). 2012. *Monty Python's Flying Circus: Complete and Annotated...All the Bits*. New York: Black Dog & Leventhal Publishers.
- Dickie, Imogen. 2012. "Skill Before Knowledge." *Philosophy and Phenomenological Research* 85:737–745, doi:10.1111/j.1933-1592.2012.00638.x.
- Fantl, Jeremy. 2011. "Ryle's Regress Defended." *Philosophical Studies* 156:121–130, doi:10.1007/s11098-011-9800-8.
- Ginet, Carl. 1975. *Knowledge, Perception and Memory*. Dordrecht: Riedel.
- Hornsby, Jennifer. 2011. "Ryle's *Knowing How*, and Knowing How to Act." In Bengson and Moffett (2011), 80–98.
- Joyce, James. 1939/2012. *Finnegans Wake*. Oxford: Oxford University Press, oxford world's classics edition.
- Kasparov, Garry. 2010. "The Chess Master and the Computer." *The New York Review of Books* 57. Online at <http://www.nybooks.com/articles/archives/2010/feb/11/the-chess-master-and-the-computer/>.
- Luzzi, Federico. 2010. "Counter-Closure." *Australasian Journal of Philosophy* 88:673–683, doi:10.1080/00048400903341770.
- Ryan, Sharon. 2003. "Doxastic Compatibilism and the Ethics of Belief." *Philosophical Studies* 114:47–79.
- Ryle, Gilbert. 1945. "Knowing how and Knowing that." *Proceedings of the Aristotelian Society* 46:1–16.
- . 1949. *The Concept of Mind*. New York: Barnes and Noble.
- Smith, Angela M. 2005. "Responsibility for Attitudes: Activity and Passivity in Mental Life." *Ethics* 115:236–271.
- Stalnaker, Robert. 2012. "Intellectualism and the Objects of Knowledge." *Philosophy and Phenomenological Research* 85:754–761, doi:10.1111/j.1933-1592.2012.00640.x.

- Stanley, Jason. 2011. *Know How*. Oxford: Oxford University Press.
- . 2012. “Replies to Dickie, Schroeder and Stalnaker.” *Philosophy and Phenomenological Research* 85:762–778, doi:10.1111/j.1933-1592.2012.00641.x.
- Steup, Matthias. 2008. “Doxastic Freedom.” *Synthese* 161:375–392.
- Strevens, Michael. 2008. *Depth: An Account of Scientific Explanations*. Cambridge, MA: Harvard University Press.
- Warfield, Ted A. 2005. “Knowledge from Falsehood.” *Philosophical Perspectives* 19:405–416, doi:10.1111/j.1520-8583.2005.00067.x.
- Weatherson, Brian. 2008. “Deontology and Descartes’s Demon.” *Journal of Philosophy* 105.
- Wiggins, David. 2009. “Knowing How to and Knowing That.” In Hans-Johann Glock and John Hyman (eds.), *Wittgenstein and Analytic Philosophy: Essays for P. M. S. Hacker*, 263–277. Oxford: Oxford University Press.