Action and the problem of evil

Heine A. Holmen

Published online: 16 Dec 2015.

Submit your article to this journal

View related articles

View Crossmark data

Action and the problem of evil

Heine A. Holmen*

Institute of Philosophy, University of Tromsø, Tromsø, Norway

Most contemporary action theorists deny the possible existence of intentionally evil actions or diabolic agency. The reason for this is a normative interpretation of agency that appears to be motivated by action theoretic concerns, where agents are conceived as necessarily acting *sub specie bonie* or under 'the guise of the good'. I argue that there is nothing in human agency to motivate this view and that diabolic evil is not at odds with inherent features of our nature.

Keywords: evil; action theory; intentional action; moral psychology; practical knowledge; Anscombe

No system of philosophy which proposes to explain the mysteries of existence can leave untouched the undeniable and perplexing fact of warps and imperfections in the fabric of our life, seemingly inherent in the very tissue of which it is woven.

C.M. Chilcott[1]

Introduction

Evil haunts philosophy. That is a double truth. It haunts *philosophy* by showing up as a problem in many areas. And it *haunts* it by philosophers tending to shy away from it: either boldly denying its existence, or seeking to tame it by ontological reduction. In either case, philosophy removes its hoofs, claws, and horns. Its ghost will then lurk around precisely because evil *is* something – something horrible.

Traditionally, philosophers have done this because evil, as a *sui generis* ontological category, was hard to reconcile with belief in the supreme perfection of the ultimate origin and cause of being. Shortly put: the positive existence of an imperfection like evil is in tension with the cause of all existence being perfect. One is led to choose between denying either evil or the ultimate cause.[2] Today people are perhaps less moved by such considerations. At least, doctrines of perfection no longer function as final arbiters for serious ontological discussions. Yet, the deprivationist-cum-reductionist attitude remains dominant. For despite the popular use of the term in certain contexts – politics, films, songs, literature, media, and so on – people overwhelmingly consider evil as an outdated concept not to be taken seriously. This goes for philosophy: evil is not considered a serious concept apt for describing things and events in the world. Perhaps we can talk about bad things, consequences, and people. But *being* evil? No way!

Today's evil-scepticism is not justified by reference to an alleged conflict between evil and the perfection of the world's origin; rather, it is cut short in *us* already. It is claimed

*Email: heine.a.holmen@uit.no

that there is a conflict between our nature as agents and being truly evil, so that no one can be evil in the sense of willing it. Of course, we can produce suffering, imperfection, and bad things by our actions. No one is denying that. The claim is, rather, that we cannot do so without, at the same time, act in light of the good. If this is right, we can no more intentionally be cruel – that is, cruel merely for cruelty's sake – than there could exist square circles and ironwood. That is an interesting philosophical development. For one thing, it shows that the problem of evil might go deeper than the traditional ontological debates assumed. After all, the traditional problem of evil was relational: it was the problem of combining the existence of evil with the perfection and divinity of the world's origin. Today's problem of evil, however, stems out of the idea that evil in itself is an *inherently* problematic notion. If this is right, the traditional problem might turn out to be superfluous, since there might be a real problem at hand with evil in itself and not merely between evil and perfection.

For the sake of clarity, I shall refer to the notion of willing evil and cruelty as *diabolical evil*. If such actions could obtain, it would involve the following: (1) some action, Φ, constituting a wicked act of evil (i.e. by yielding terrible suffering and malice); and (2) there being agents who Φ-es intentionally because of or for the reason that the act has the properties described in (1).[3] In what follows, I argue that properly understood the nature of agency leaves room for such evil. Denying its existence as absurd is therefore no consequence of our nature as agents; rather, it is a genuine possibility of our agency and something whose existence will need to be settled empirically.

Denying evil

In one sense, evil is clearly an empirical fact. The world has witnessed the *Holocaust* and the systematic extermination of millions of people. Add the atrocities of warfare in general: WWI and II, Korea, Indochina, Algeria, Congo, Middle East, and so on. Add ethnic-cum-religious cleansing: the former Yugoslavia, Rwanda, Sudan, Native Americans, and the Middle East of today. Add the cruelty of civic life: rape, murder, crime, and violence. To study this is, as Arne Johan Vetlesen says, 'to confront an abundance of empirical material'.[4] What can nevertheless be denied, though, is that any of these results from people *being* evil. That is, one can deny that it was evil people who did it. If one denies this on principled ground – and not by reference to mere empirical contingencies – one is an evil-sceptic. Evil-sceptics are opposed by evil-realists, who on the contrary think that: (1) there *is* a property of (being) evil in the diabolic sense, (2) this property is instantiated (at least, sometimes), and (3) this property has the metaphysical status that other garden variety properties have (i.e. causal influence, ontological irreducibility, and so on).[5]

The chief reason why some are evil-sceptics is that such actions violate what many take to be a central condition of intentional agency: namely, pursuing (subjective) goodness. The orthodoxy – what Michael Stocker labelled 'the philosophical view' – is overwhelmingly clear: 'the good or only the good attracts' in the sense that 'acts or features attract because or only because they are (believed) good'.[6] Given that diabolic evil, as defined above, 'would mean doing intentionally the opposite of what you take to be good' it would seem to follow, as Lars F. Svendsen puts it, that '[d]oing evil because it is evil is an incoherent idea'.[7] If this is right, there is action-theoretic motivation for evil-scepticism.

The doctrine that agents *must*, by necessity, act for the (perceived or subjective) good is traditionally characterised as the claim that agents, insofar as they act intentionally or

for reasons, *must* act under or in the light of a 'guise of the good' (or act *sub specie bonie*, as the scholastic term goes). This idea has ancient roots: 'Every soul pursues the good and does its utmost for its sake', Socrates says in the *Republic*, which is why he sometimes claims no one errs willingly.[8] Aristotle joins him when he holds that 'every action and choice is thought to aim at some good'.[9] Augustine declaring, 'the evil will itself is not effective but defective',[10] and Aquinas identifying the will as rational desire – viz. where '[r]eason comes before the will and directs its activity' – are variants of the same doctrine.[11] Finally, and closer to our age, we find Kant arguing that the form of volition in general is identical to the form of a practical moral law and thus that agents cannot fail to act in light of morality.[12]

These views may not correspond precisely to the doctrine I refer to as 'the guise of the good'; however, what I have in mind clearly descends from them and seems to inherit, as Kieran Setiya points out, 'whatever plausibility they have'.[13] In its modern take, the doctrine springs out of assumptions and principles guiding action theory. Following Anscombe, the fundamental mission – and, indeed, problem – for such theorising is to illuminate the nature of deeds we recognise as truly *ours*. According to her, these are the *intentional* actions or the class of actions 'to which a certain sense of the question "Why?" is given application': namely, the sense 'in which the answer, if positive, gives a reason for acting'.[14] The idea is that intentional actions are all and only those having (or allowing) a certain explanatory structure: viz. being explainable by citing one's own reasons for so doing. The central question for action theory then becomes: What am I doing when I act for reasons?[15]

The 'guise of the good', in one version or another, may here earn its privilege as a universal action theoretic principle if it turns out to be a universal mark of the reason-giving structure of intentional actions. According to its defenders, it is such a mark by determining how a reason *must* relate to the agent and action in order for it to count as intentional and done *for* a reason. The distinction between the reason why I am writing these words – say from duty, to gain insight, to impress someone – from the reason why a volcano erupts or my heart pumps blood, is, accordingly, that *my* reason also figures as a reason for me in the sense that I take to be a good reason for so doing. Agent act *for* reasons, then, insofar as the deed is done because it appeared in a favourable light and as good.[16]

It is important that we notice the emphasis put on the agent's perspective. The claim is *not* that agents must in fact have good reasons for acting or in fact pursue goodness: only that it must appear as such and that s/he acts because of this appearance. A distinction among reasons for action is here implied: there may be good normative reasons to, say, donate money to charity, a fact that would count in favour of so doing. On the other hand, we have the agent's own reasons and motivation for acting, which are relevant for explaining why s/he did something. The latter need not correspond to considerations that in fact support the deed. If they do, the agent's reasons were in fact good reason for pursuing her/his course. If not, her/his reasons fail to be good in this way and so fail to have the action-justifying character that good normative reasons have. What the 'guise of the good' doctrine claims is that the agent's own reasons must nevertheless appear to be good for the agent so that *if* things had been as the agent supposed, the action would have been the right thing to do.[17]

Moreover, the doctrine may allow a wide range of diversity in the term *goodness*. '*Bonum est multiplex*: good is multiform', Anscombe remarks, before adding that what the requirement asks for is that 'a man should see what he wants under the aspect of *some* good'.[18] One man's subjective goodness need not be identical to or overlap with

another's. Thus, it may make perfect sense for me to pursue what you take to be bad, such as pleasure, as long as I pursue it as falling under the aspect of some good. However, even this multiplexibility has its limit, as it must still be an act under some aspect of *good*. Due to this limit, it is impossible to do something that from the agent's point of view fails to come under *any* aspect of good. An example of that would be killing someone violently *simply* in order to produce this mischief: that is, to do it without considering the act as conductive for some good or other. The importance of this is, of course, that such strictures contribute to rule out diabolic evil.

A final thing worth noting is that for something to operate as one's reason the agent must not only regard it as a good reason to act. After all, one can regard some *p* as a good reason and yet act for some other reason *q*, or not act at all. For *p* to become one's reason one must therefore also act *because* one takes *p* to bespeak the goodness of the action. Now there are deep controversies in action theory regarding the precise nature of this *because*-locution. Briefly put, one may follow Davidson and treat it as a causal–psychological relation between certain mental states of the agent – viz. those constituting her/his judgement that *p* is a good reason to act – and her/his bodily behaviour and that the main target of action theory is to clarify this relation with the aim of defining intentional actions with reference to its aetiology. Alternatively, one may follow Anscombe and consider it a constitutive relation so that when an agent takes some *p* as a good reason to act this is what makes the whole sequence an action in the first place.[19]

Important as this is, we shall side step the whole issue. The reason is that both camps endorse some version or another of the 'guise of the good' in order to specify the precise nature of what it means to take some *p* as one's reason to act and thus make *p* operate as one's reason. According to both, this happens only insofar as the action is done because the agent took (rightly or wrongly) *p* to bespeak the (or some) goodness in so acting. It is thought that we need to say this to distinguish intentional actions – things done *for* a reason – from things we do merely out of motivation. The fact that some behaviour is occasioned by one's *mere* wants and pricks of desire is not enough to render it an action in the more elevated sense we are talking about. Something more is at stake.

Let me illustrate with a case Sigmund Freud discussed: he broke his inkstand once in what stroke him as an unusually clumsy manner. Why, he asked, did this happen? After some analysis, he thought he had the answer: his body moved to destroy the inkstand because of a hidden, unrecognised desire for a new and better one. If Freud's self-analysis is correct, the mishap turns out to be a rather intelligent, goal-directed deed. It is something that was thoroughly guided by some of his mental states, and thus could have been intentional. Except, it was not: too much happened behind his own self-conscious back for this to count as intentional. Rather, we should say that Freud's accident was an instance of *merely motivated behaviour*. The difference between that and intentional action proper is thought to be that the latter and not the former involved an agent regarding *his desire as a good reason to act*. In this respect, the 'guise of the good' plays an important, action-theoretic role in making a distinction among different psychological categories of human action and behaviour.[20]

We may refer to this view as the *normative interpretation* of agency. It recommends itself insofar as the 'guise of the good' doctrine is vital for understanding the contrast between intentional agency and other things we do. A consequence of this normative view, however, is that diabolic evil becomes an absurdity. After all, if the 'guise of the good' is a universal mark of agency, no one could possibly pursue some deed merely for its wickedness or badness: that is, without taking it to be good in some respect or other. Since a truly diabolic agent, as I have defined him/her, does *not* pursue his/her deeds

because s/he (erroneously) perceives some goodness in the act – on the contrary, truly diabolic agents are not deluded or benighted about these terrible consequences of their own deeds – s/he will by necessity fail to act *sub specie bonie*. And, according to the normativists, this is just the same as opting out of human intentional agency altogether.

First stab against the 'guise of the good'

The road to Hell is paved with good intentions; or so the proverb says. In fact, it could not possibly be different, *if* the only way to act at all is by virtue of believing the act good. Even Satan would then aim at being a goody-goody. 'Evil be thou my good', he cries in Milton's *Paradise Lost*. This is now rendered a literal truth. Satan's demonic quality is nothing but an honest mistake: he has simply misidentified the most evil cruelty and malice as goodness. Alternatively, he is not acting intentionally at all and what he does is merely the behavioural products of some sub-intentional pathology or other.

This is clearly an inadequate and, indeed, sappy portrayal of demonic evil at its peak of wickedness. It analyses the evil agent as nothing but a well-intended fool or as mere pathology, thereby loosing sense of that which is so despicable and abhorrent in the motivational structure of such agents.[21] Now it might of course be true that these demonic qualities fail to be a part of the paradigm of contemporary evil.[22] *Our* paradigm might perhaps be closer to that thoughtless banality Arendt found so dangerously flourishing in Adolph Eichmann – viz. the 'architect of the Final Solution' and the bureaucratic organiser of the Holocaust.[23] But we should not let that contingent fact – if it is a fact[24] – springing out of conditions in today's society, determine the analysis of evil in general. In particular, we should not allow such contingency to rule out the possibility of evil agents of a more vibrant and demonic kind.

The problem of pointing to such cases, however, is that one often gets involved in discussing the specifics and details of it. Some might for instance dismiss Satan as fiction having little to do with the nature of human agency. Others may dispute that there is something wanting in the analysis, deny his status as an intentional agent, and so on. When we come up with real-life cases, the dialectical situation is more or less the same. Christine Korsgaard's discussion of evil is illustrative: she analyses the deeds of ruthless tyrants as psychosis or pathology and claims that the tyrant fails to be an agent because he 'can never separate himself from *one* of his impulses, and so consolidates himself into a mere force of nature, an object, a thing'.[25] In my view, this shows how the normative conceptions of agency, much like Plato in the *Republic*, succumb 'to the moralist's temptation to represent the bad person as a compulsive addict, an unenviable wreck'.[26] However, I also think such verdicts show why discussing specific cases is likely to lead to dialectical stalemate.

An alternative is to follow Setiya and try to settle the score by reference to 'general constraints on a theory of intentional action' and ask whether 'these constraints rule out a normative interpretation of the recognition and adoption of reasons involved in acting on them'.[27] When we do we see that the normative conception fails to address a fundamental desiderata on a theory of intentional action: namely, the fact that an agent *qua* acting intentionally or for reason knows – indeed *must* know – *what* s/he is doing and his/her reason *why*. This necessary self-knowledge is something that ought to be explained by a theory of intentional actions: after all, sound methodological principles suggest that such a necessary connection springs out of the nature of intentional agency.[28] If one cannot explain this connection, one has either a wrong or incomplete conception of agency. Now the normativists are *not* explaining this epistemic necessity by merely saying that agents

act in virtue of taking one's action to be good. To illustrate: suppose we grant that Pierre takes some action, say donating money, to be appropriate and good. The question then is: how does Pierre get from that – viz. from knowledge of what he values – to know that he is in fact donating money? The answer is that he does not, for then he would also need to know that he will do what he values, which requires additional premises. Merely knowing that one takes an action to be good is therefore not enough to explain what agents know in acting. This shows that whatever its merit the normative conception cannot be the full story.

However, it gets worse: in explaining this epistemic necessity, it turns out that the normative conception is not even called for. As I mentioned before, Anscombe aimed to classify the range of intentional actions and she claimed that these are the ones 'to which a certain sense of the question "Why?" is given application': the sense being a request for the agent's reasons.[29] An action fails to be intentional, then, insofar as this request is inappropriate. Anscombe knows that acting for a reason and acting intentionally means more or less the same thing, so she needs to be more explicit about the sense of the question in order to break into this conceptual circle. What she suggests is that the question is 'refused application by the answer: "I was not aware that I was doing that"' and that intentional actions must be subject to a kind of non-observational, direct, and spontaneous knowledge of what one is doing and why.[30] Baking a cake, for instance, is not something you would be doing intentionally insofar as you failed to know it, or if you knew about it through the testimony of others, your senses, or in virtue of an inference to the best explanation. If you really are doing this intentionally, knowing it should spring out spontaneously and directly from the very fact that you are you doing it. This sort of knowledge is a necessary condition for the action being intentional, and Anscombe thought we could use it to break into the conceptual circle of reasons and intentions.

Anscombe is here operating according to the very same sound methodological principles mentioned above: in general, if being of a kind F goes by necessity along with also being G, then this necessary connection must stem from the fact that Fs and Gs either overlap completely in their nature, or one is an essential component or aspect in the nature of the other.[31] Knowledge and intentional actions clearly fail to overlap completely and knowledge can occur without intentional actions, so Anscombe's observation, if true, would show that knowing is an integral element or component of acting. That is, this epistemic necessity suggests that acting *just is* an epistemic relation, a way of knowing.

In contemporary literature, people have sought to explain and warrant this insight by prodding more deeply into the metaphysics of action. In particular, philosophers have sought to explain the necessity by defining intentional actions by something like the ensuing steps: (1) defining intentional action as behaviour brought about by a complex mental state (viz. the intention), (2) defining the intention cognitively as being a belief or cognitive attitude about *what* one is doing and the reason *why*, and (3) showing how these intention-beliefs normally are justified to such an extent that they count as knowledge.[32] According to this model, an agent acts intentionally or for a reason just in case she acts in virtue of believing that she is (or will be) so acting. If (1) through (3) are true, however, there will be no action theoretic motivation for making 'the guise of the good' a universal criterion of intentional agency. In turn, this makes diabolic evil a metaphysical possibility. Such an agent will simply be someone who acts in order to produce suffering in virtue of believing that s/he is (or will be) so acting – which satisfies conditions (1) through (3).

I am sympathetic to the form of this argument against normativism;[33] however, the details of (1) through (3) are problematic on several points. A minor problem is that it is a mistake to locate the defining characteristics of intentional actions in their effective causal

antecedents and thus read the 'brought-about' locution in (1) as psychological–causal. As I mentioned above, this is the disagreement we found between Davidsonians and Anscombeans on how we are to read the 'because-of' locution in action explanations – that is, whether it is to be understood as a causal or constitutive relation – and contemporary cognitivists side, mistakenly in my view, with Davidson on this one.[34] However, this is *not* something they must do and they could possibly revise (1) to suit better with the Anscombean conception. Given that – and since I will *not* argue decisively for the Anscombean model in this paper – I will simply bypass this problem.

Let us move on to graver difficulties. Starting with (2): by defining the intention cognitively, one will make agency answerable to distinctively epistemic requirements. A cognitive state is a state that purports to represent the world *as it in fact is*. It has a representative aim, or a world-to-mind direction of fit. This is in contrast to distinctively practical or conative attitudes, whose contents rather aim to amend the world and have a mind-to-world direction of fit. Intention-beliefs would thus make agents answerable to *how the world is*. Agents forming a belief about what s/he is (or will be) doing may consequently be asked for her/his epistemic reasons for so doing. However, agents do not normally have this at their disposal. On the contrary, what we decide to do is something that often goes contrary to or beyond our body of evidence, so that we do things that fail to be likely or probably given what we know about ourselves. Accordingly, agents would be epistemically at fault for doing things that go beyond their evidence – an absurd result.

In response, some cognitivists have argued that agents have unusual epistemic reasons at their disposal. Setiya has for instance argued that one's know-how provides non-evidential ground for the intention-belief, whereas Velleman has emphasised a notion of incoming evidence.[35] However, that will not take cognitivists off the hook. For one thing, agents do not form their intention for those reasons. On the contrary, agents intend to act for distinctively practical reasons – such as, because s/he wants it, values it, or so on – so even if it were true that s/he has some peculiar sort of epistemic justification at her disposal, s/he will nevertheless be epistemically culpable for forming the belief for the wrong sorts of reasons. Moreover, we can challenge the assumption that these things constitute epistemic reasons in the first place. An epistemic reason for belief in the truth of some *p* is something its agent *must* be able to have as his/her reason for so believing. The problem is that an agent cannot do that without at the same time render the intention-belief false. For suppose that an agent deciding to walk in order to be healthy were to form her intention-belief that she will be so doing in light of some other, epistemic reason. Then her intention-belief would turn false, because it would no longer be true that her decision is based on her health-concern – rather, it would be based on the epistemic reason. So, if agents were to form their intention-beliefs for the epistemically speaking right sort of reason, this would make their beliefs false – an odd sort of epistemic reason, indeed.

Last, but not least, by defining intentions in cognitive terms one will reserve intentional agency for creatures who represent the world cognitively and in terms of representational contents. It is not clear to what extent small children or non-human animals have such complex capacities. It is not even clear whether adult human beings always operate in light of such representational contents, given that the phenomenological richness of our senses in all probability can only be accounted for in virtue of non-representational contents. On the other hand, it also turns the very idea of omniscient agency into an absurdity. Such an agent would not need to aim for a correct representation and thus form beliefs about the world at all: after all, s/he knows from eternity all there is to know and need not discover the world and try to represent it accurately, no more than s/he needs to make conjectures about things. By defining

action in terms of belief-states, we thus rule out that such an agent could act at all. Omniscience would come along with necessary omni-*im*potence, as it were. A general problem with cognitivism is therefore that it turns agency into an exclusive gift for knowledge-seeking and concept-mongering animals, such as us. That smacks of anthropomorphism and there is no motivation for restricting the notion of intentional agency to such a narrow cast of mind as ours.

This shows that cognitivism fails as a theory of intentional agency. Nevertheless, we could learn something important from both its ambitions and its failures. It correctly aimed at a theory of intentional agency that explains the epistemic necessity. In what follows, I will defend a conception of agency closer to Anscombe's original proposal that does explain this at the same time as it avoids the pitfalls of cognitivism – viz. anthropocentrism and saddling the agent with specifically epistemic burdens.

Second stab against the 'guise of the good'

As I said, Anscombe characterised the knowledge occurring in intentional agency as non-observational, direct and spontaneous – viz. coming about directly, merely in virtue of acting and not given via one's senses, inference, and so on. She later sought to explain these claims by calling it *practical knowledge*, which is defined, with reference to Aquinas, as knowledge coming about by virtue of being 'the cause of what it understands'. This is in contrast with contemplative or theoretical knowledge, where the fact known operates as the causal antecedent of the knowledge. There are two major advantages of taking Anscombe at her words: first, she talks about knowledge and this is, as I will argue, a more basic – indeed primitive – relation than talk about cognitive attitudes and representational-cum-conceptual attitudes. Second, she emphasises its practical dimension, which means that this knowledge (and *eo ipso* this action) comes about in virtue of distinctively practical thinking. In combination, these claims, if we can support and integrate them in a theory of intentional agency, avoid the epistemic- and the anthropomorphism worry we found marring cognitivism. I shall therefore argue that they can be substantiated and take them as a starting point for a theory of action that explains the necessity of knowledge in action at the same time as it leaves room for diabolic evil.

What we need to do is define practical thought – the process in virtue of which one decides what to do and why – as an epistemic faculty and as 'a power of the mind to cause itself to know'.[36] In general, it holds for all epistemic faculties that *if* they operate normally and the context of knowing is favourable, a fact can be known by virtue of the faculty causing one to know. Perception is such a faculty. Thanks to perception the fact that some wall is white or square can be known. Briefly put, features of the wall will, along with other factors, causally impinge on one's sensory apparatus, which under favourable circumstances produces knowledge of the walls whiteness or squareness. In this way, the known fact combines with normal faculty functioning in a favourable context to bring about knowledge. The question is whether we could say something similar about practical thought – viz. that this faculty produces knowledge of one's action by operating normally under favourable circumstances. If we could, we would have an explanation for how and why acting intentionally yields knowledge. Moreover, it would explain why such knowledge is necessary. After all, this faculty operates normally by yielding a decision about what to do and this, in turn, is what leads one to act intentionally under favourable circumstances. It would follow, that such actions and one's knowledge of them necessarily come about in the same situations.

One reason to think that practical thought is an epistemic faculty is the fact that it would explain Anscombe's observations. It is difficult to account for the knowledge one has in action solely as the work of the more traditional epistemic faculties. As I argued above, an action would fail to be an intentional action of mine if I had to know about it through perception, inference to the best explanation, the testimony of others, and so on. Moreover, if an agent were nevertheless to establish the knowledge s/he has in acting by means of these other faculties, s/he would in many cases fail to find rich enough evidence for what s/he knows through acting. After all, an agent often acts and does things whose occurrence runs counter to or beyond one's body of evidence.

To illustrate: a man may be on his way to somewhere, but his evidence – provided by his senses and so on – will not necessarily tell where he is going and so this is not the basis on which this agent knows what he is doing. He might, after all, not take the most efficient route to his destination. He may take all sorts of wrong turns and stride long detours just for the fun of it, and so on. Moreover, if the church and the pub lies at the same place, there may be no telling – on the basis of outer evidence – which place he is going to. But he, the agent, knows all along where he is going and why. Since his available evidence cannot provide basis for this knowledge, it must arise elsewhere.[37]

Note that this does not change by throwing introspection or the spontaneous and direct awareness of our own conscious, mental states into the pot. The action is, after all, an event out there in the external world and not something we know about merely in virtue of the inward gaze into our own minds. Even if we know the content of our intentions and decisions in this way, this will not yield knowledge of what we are in fact doing since these mental items can be had without being acted upon. The same reasoning holds for the bodily senses: even if there is a way for us to know how our body is behaving at all times – which is unlikely – this would typically not yield knowledge of what we are doing, since our actions more often than not involve elements in the world external to our bodies. Shortly put: we are – and know we are – baking cakes and erecting buildings, not just that we are moving our bodies while hoping for these effects to occur out there in the world. Anscombe's observations thus give at least *pro tanto* support for the idea that the practical faculty is an irreducible, *sui generis* and direct source of knowledge since the other faculties will fail to yield this knowledge.

Another line of support is that the relation *doing A intentionally* – and thus, thinking with a view to action under favourable circumstances – is factive in the same sense that *hearing that p* or *seeing that p* is. If you see that the cat is on the mat, it must be true. This is because *seeing that*... is a factive locution and a relation we have to all and only facts. In the same vein, *doing ... intentionally* is a locution that holds all and only if it is a fact that this action gets done. If you are baking a cake intentionally, then surely you must be baking a cake. In general, factive mental relations, such as hearing that *p*, seeing that *p*, remembering that *p*, witnessing that *p*, and so on, implies that the subject also knows the fact in question. If you see that the cat is on the mat, then you know that the cat is on the mat. The reason for this is that seeing that *p* just is a way of knowing that *p*. Now, we have already argued, following Anscombe, that baking cake intentionally implies knowing that you are. Analogously to how we treat seeing that *p* it would therefore be plausible to assume that intentionally doing *A* yields knowledge because intentionally doing something just is a way of knowing what one is doing.

The importance of this is that it requires an explanation. One option is to say that the relevant faculty yields direct knowledge in and of itself. This is what perception does. Another is that it yields knowledge only via other sources of knowledge. Remembering

that *p*, for instance, is factive, but not a direct source of knowledge in and of itself since memory depends on being fed by other epistemic faculties, such as perception.

Of course, a good test for whether a faculty yields direct knowledge is whether we can explain its epistemic functions and merits by means of the other faculties. The fact that one cannot do this for perception, but can do it for memory, indicates that the former is a direct source of knowledge whereas the latter is not. The fact that knowledge in intentional action cannot be accounted for by means of the other faculties is therefore a reason for thinking that this faculty really is epistemic and a *sui generis* direct source of knowledge. In this respect, this second claim is related to the first claim.

Moreover, as a working hypothesis, one should begin with the assumption that the faculty in question is a direct source of knowledge and only retract this in light of reasons for thinking otherwise. Absent any such reasons for considering practical thought a direct source of knowledge, we should assume that it really is epistemic in this sense given that it does give rise to knowledge through acting.

One reason many people have for doubting this initial hypothesis for practical thought, however, is that they assume that the nature of knowledge is cognitivist and thus something that cannot be constituted by purely practical thinking and action. Traditionally, knowledge is simply reduced to and conceived of as some sort of belief or cognitive state. Orthodox epistemology implicitly endorses this sort of metaphysics of knowledge when it seeks to define knowledge as a true belief in conjunction with conditions like justification, reliability, counterfactual robustness, and so on. As such, knowledge is thought to be a representational state with a world-to-mind direction of fit. Distinctively practical knowledge – viz. a product of distinctively practical thought alone – is then not something that could constitute knowledge all on its own, since these sorts of states have features that are incompatible with essential features of knowledge.

Fortunately, there is reason to resist this overtly cognitivist image of knowledge in favour of a neutral conception. For knowledge is a relation that is better understood as a metaphysical primitive rather than a composite or metaphysical hybrid composed by a belief (or cognitive state) in conjunction with certain other conditions. There is surely negative inductive evidence to be found for this thesis in the fact that for over fifty years now – more precisely, since Edmund Gettier published his famous counterexamples to the definition of knowledge as justified true belief[38] – intensive, philosophical research has failed to yield a satisfactory definition or analysis of knowledge along these lines. The result of this massive research is only a futile succession of increasingly complex analyses falling prey to increasingly complex counterexamples – a clear example of 'a degenerating research programme', as Timothy Williamson characterises it.[39]

However, one can also argue that such decomposition is ill founded even in principle. For the basic motivation for this reductionist programme is a couple of metaphysical assumptions about the nature of knowledge and how it relates to the mind of the epistemic subject: in particular, that knowledge as truth-entailing must be a metaphysical hybrid relation constituted by internal mental conditions of the subject in combination with conditions of his/her external and non-mental surroundings.[40] Knowing that Kennedy was shot, for instance, implies its truth and thus relates the subject to this non-mental, external fact. At the same time, the subject is thought to be in certain internal, mental condition in virtue of which s/he knows this fact. But if knowledge is constituted by such divergent factors, it would surely be a hybrid metaphysical category crying out for analysis into its internal and external components.

The problem with this assumption, however, is that it assumes wrongly that knowing is a *composite* state – that is, a state analysable as a conjunction of some internal, mental

conditions and some external, environmental conditions. If knowledge was composite in this way, knowing that Kennedy was shot would for instance be reducible to a conjunction citing the subject's internal conditions – such that s/he believes this, trusts her/his informant, and so on – in conjunction with external facts – such as the fact that Kennedy was shot, the informant being trustworthy, and so on. The thing is that we can prove this wrong. For on the assumption that knowing some p through, say testimony, really is a composite state, there should be a sufficiently rich internal condition, C, that stands as the purely mental reality underlying this knowledge – for example, one's belief that p, one's trust in the informant, and so on. Moreover, there will also be a sufficiently rich external, environmental condition, E, that stands as the non-mental reality underlying this knowledge – for example, the fact that p is true, that the speaker is trustworthy, and so on. Then, on the condition that knowing p through testimony is a composite state, knowledge should obtain merely in virtue of the fact that a conjunction of such sufficiently rich external and internal conditions obtains. It would fail to be composite, though, if such a conjunction were to obtain without knowledge obtaining. Trouble is, this can easily happen.

All we need is a triple of epistemic situations: suppose that knowing through testimony that Kennedy was shot is a state that obtains in two different epistemic situations, *S1* and *S2*. Suppose further that the difference between them is only this: the fact that Kennedy was shot remains fixed, but in *S1* one knows this through believing John when he tells that Kennedy was shot (one trusts him and he really is trustworthy), at the same time as one disbeliefs Paul's testimony to the same effect (one fails to trust him and he really fails to be trustworthy). In *S2*, this gets reversed: here one knows it by believing Paul (whom one trusts and who is trustworthy), while disbelieving John (who is neither trusted nor trustworthy). However, now we can design a third situation, *S3*, which is internally like *S1* but externally like *S2*. Here, then one does not trust Paul but John (as in *S1*); however, John is the trustworthy one while Paul is not. Consequently, one fails to know that Kennedy was shot in *S3* despite the fact that in *S3* there obtains a conjunction that combines the (allegedly) sufficiently rich internal and external conditions from *S1* and *S2*. This shows that knowing through testimony is not a composite state. Since this is a set-up we can use on any piece of knowledge, we may in fact conclude that knowledge in general fails to be a state composed of internal mental and external non-mental conditions.[41] It follows that knowledge is not a troublesome hybrid composite of internal and external components crying out for analysis. Together with the negative induction above, this result will thus provide strong support for the thesis that knowledge is rather a metaphysically basic relation.

A better thesis is therefore that knowledge is metaphysically basic: a primitive relation binding a subject to some fact. Of course, such a relation comes about in virtue of the fact that the subject is in some condition, C, and the world is in some condition, E; but there is a way to accept such claims without implying that knowledge is thereby a complex constituted by these conditions. The way to go is to follow Keith Hossack and distinguish more clearly between knowledge itself and its provenance or ground. Hossack suggests that the combination of factors yielding knowledge is better understood as its causal provenance rather than its constitutive ground. In other words, we should say that conditions of the world combine with conditions of the subject to *cause* knowledge of some p, rather than comprising it. In this way, we avoid commitment to analysis and the idea that knowledge is a composite, while at the same time allowing for the plausible claim that knowledge comes about in virtue of the fact that certain non-mental and mental conditions arise. Another thing worth mentioning is that this assumption squares better

with the fact that knowledge comes about in a vast plurality of ways: by representation, by beliefs, by the use of concepts, or by radically different means. If these are viewed merely as the causes of knowledge rather than its constitutive components, we may grant significant plurality in the provenance of knowledge while keeping the knowledge relation itself uniform and one. The same state, such as a fire, may after all come about from a wide range of causes. By treating the relationship as causal, we may thus grant that some knows some by virtue of conceptual representation and beliefs, while others know through, say, non-conceptual representation.[42]

If, however, the mental, intrinsic states of the agent – by virtue of which the agent knows something – stand only as the cause of knowledge in combination with some extrinsic context, the nature of these mental, intrinsic states will not impinge on – because not constitutive of – the nature of knowledge itself. The fact that the mental states in virtue of which we know something often are cognitive and representational, is therefore no reason for thinking that knowledge in itself has these cognitive features. It is the mental state in virtue of which we know that aims to fit its object, represent the world correctly, and so on that needs these properties: not the knowledge state itself. In this respect, this conception paves the ground for a neutral view of knowledge. In itself, knowledge is neither practical nor cognitive: these are features of its causal provenance, of our ways of knowing, and not of knowledge itself. With this conception in hand, we thus remove the major obstacles for considering practical thought as a possible *sui generis* provenance of knowledge. In light of Anscombe's observations and the factivity of intentional agency, we should therefore assume that practical thought really is a genuine epistemic faculty – a faculty yielding knowledge under favourable circumstances in virtue of distinctively practical thought.

By defining practical knowledge in this way, however, we are also in position to identify the class of intentional actions in an epistemic way: namely, as the deeds that are objects of such knowledge. As an epistemic faculty, practical thought yields knowledge of what one does and why if the faculty operates normally under favourable circumstances. Operating normally under favourable circumstances is also the conditions under which practical thinking issues in action and in doing what one decides to do. Thus, the conditions under which one has practical knowledge are exactly the same as the conditions under which one acts intentionally. Consequently, we can use our epistemological apparatus to define intentional action as a practical way of knowing.

A key advantage to this framework is thus that it promises to fulfil the project that Anscombe signalled by calling attention to the peculiar epistemic features of human intentional agency. Moreover, it does that in a way that avoids the pitfalls of cognitivism. On the one hand, its emphasis on knowledge rather than beliefs leaves open the possibility of acting intentionally without beliefs, or conceptual- and representational thought. On the other hand, it dispels the epistemic worry, since knowledge unlike belief is a state that warrants itself, as it were. If a subject believes that s/he will be going to act, s/he can be asked for her/his reasons for so doing and the belief will have to be justified by some means or other. If a subject knows that s/he will be so doing, asking for her/his reasons for knowing has a different, non-justificatory sense: the answer would be citing its causes, not justifying it. The only meaningful way one can ask for its justification is by challenging its status as genuine knowledge. Moreover, by emphasising the practical dimension of this knowledge, we also avoid saddling the agent with problematic cognitive and representational states standing in need of epistemic justification. The states and cognition taking part in such thinking are distinctively practical. As such, they may of course be asked for their practical warrant. But that burden – unlike the epistemic burden – is as it should be.

Last, but not least, since this account characterises agency in epistemic terms only, it shows that there is no need, on action theoretic ground, for committing universally to normative principles such as 'the guise of the good'. This epistemic conception can therefore function as our crucial lemma in an argument against the normativist conception of intentional agency. That, in turn, leaves diabolic evil as a genuine metaphysical possibility. Such an agent would act for no other reason than that the action in question realises or is conducive of all sorts of malice and suffering. This action is intentional, according to this epistemic conception, insofar as its performance produces practical knowledge in the agent: that is, whenever her practical thought and action yields knowledge of both what s/he does and the reason why s/he does it. In this regard, then, acting intentionally really can be identified as a way of knowing.

Conclusion: are there evil people?

It has been a long and winding road: we have gone through some of the deepest issues of action theory, plus made a couple of journeys into the realms of contemporary epistemology and the metaphysics of knowledge. What we got out of it, however, is significant. For at our disposal is now a strong action theoretic argument against the claim that diabolic evil is a metaphysical impossibility, given that action theory provides no leverage for the normativist conception of such actions. An agent may of course steer her/his life in light of what s/he deems good, valuable, pleasurable, worthy, and what have you. These may even be the dominant ways of acting for humans, thereby serving as a rough paradigm of our psychological profile. However, it is not of essence for acting intentionally. On the contrary, the essence of acting intentionally is to act while knowing in a practical way that one so acts. This can be achieved in service of the good, just as it can be done knowingly in service of what is bad and evil – or, as it were, when one acts merely for the badness-making features or cruelty of the act. Diabolic evil is therefore not something outlandish to human agency: rather, it forms a core element of it and is always there as one of the many looming possibilities of our agency.

The argument thus vindicates some of the key elements of evil-realism. However, a metaphysical possibility need not materialise at all in our world. Pegasus, for instance, has failed to do so thus far, despite being as metaphysically possible as any of us. I shall therefore end this paper by briefly saying why I think diabolic evil actually exist.

To prove this in a decisive manner would require a much more careful analysis of some case or other than I have room for now. The reason is that diabolic evil is a question of how best to explain the motivation and psychology of an agent responsible for cruelty and malice. The trouble with such agents is that they may often provide spurious motivation and information for doing what they did. Anders Behring Breivik – that is, the right-wing terrorist who on 22 July 2011 murdered 69 people, a lot of them teenagers, with hand guns at close range, after having blown up a bomb killing additional 8 people in the heart of the capital of Norway – explained for instance his deeds as countermeasures to avoid future civil war and casualties as a consequence of Muslim immigration. If this really was his motivation, he will not count as a diabolic evil agent. After all, he was then led by what he thought – erroneously, for sure – was good and right, and not motivated by cruelty and malice *in itself*. In short, he would be a goody-goody who had merely made some mistakes in judgement. Other cases follow a similar pattern: Theodore Kaczynski, or the UMA bomber, was allegedly led by environmental concerns, while Adolph Eichmann, the architect of the Holocaust, was motivated by a sense of duty and loyalty to his Führer.

However, I think, on careful analysis, one can challenge the sincerity of some or more of these motivations and explanations. Clearly, they did not need to choose the extreme measures they did to realise their goals. Nor did they think so themselves. Their choice of measures may thus reveal a pattern of preferences and an attraction towards cruel measures. A careful analysis might then demonstrate that what drove these people all along was hatred, evil. and a felt attraction for the badness-making features in and of itself. Perhaps such agency is rare. Perhaps modern society may have had the effect of making evil 'more to do with patterns of social interaction than with the character and motivation of the acting individual', as Vetlesen says.[43] However, we should take great care not to be caught off guard if – or, rather, when – such people surface on our social arenas. Whether their presence and existence on those arenas counts against the perfection of the ultimate origin of the world or not is something we can finally deal with now that we have a clear idea of what that involves.[44]

Notes

1. Chilcott, *The Platonic Theory of Evil*, 27.
2. Historically, Plato held that evil qua imperfection enjoys only a negative existence and is not something positive in and of itself. His reasons had theological resonance: all existence must be good, beautiful, and perfect, because its creator (viz. the supreme form of the good) was good, beautiful, and perfect. Evil only exists as a failure of the particular thing (such as the individual horse): a failure to participate in the Forms that constitute or cause its nature. Thus, evil is an absence of features and properties of the Form that failed to materialise fully in the particular. There is, accordingly, no more evil to this world than there are things constituting holes.
    Similar models and theories dominated the Christian tradition. Augustine, for instance, abandoned the Manichean doctrine of the material reality of evil to conceive evil as a failure of existence. Since God is 'existence in a supreme degree (...) the only contrary nature is the non-existent'. To seek its cause and origin is therefore 'like trying to see darkness or to hear silence'. See Augustine, *City of God*, 473.
3. Acting intentionally and acting for reasons are terms I will use more or less interchangeably in what follows. There are of course controversies regarding the relationship between these notions, but there is no need to settle these issues for the present purpose.
4. Vetlesen, *Evil and Human Agency*, 1.
5. See Barry, *Evil and Moral Psychology*, 7. Barry adds that realists also commit to the claims that: (a) attributing evil to x is to ascribe this property to x; and (b) asserting that 'x is evil' expresses a belief. This may hold for many realists; however, one need not commit to these claims about the nature of discourse employing the term 'evil' just because one is a realist about evil as a phenomenon. Since I am concerned only about the latter, I therefore keep these issues apart in what follows.
6. Stocker, 'Desiring the bad'.
7. Svendsen, *A Philosophy of Evil*, 108.
8. Plato, *Republic*, 505e. Also see, *Protagoras; Gorgias*.
9. Aristotle, 'Nichomachean Ethics', 1094a1–4. [My emphasis].
10. Augustine, *City of God*, 479.
11. Aquinas, *Summa Theologicæ*, II-I Q.13a.1.
12. See Reath, *Agency & Autonomy*; and Korsgaard, *The Sources of Normativity; Self-Constitution*.
13. Setiya, *Reasons without Rationalism*, 21.
14. Anscombe, *Intention*, 9.
15. By zooming in on this class of human events and behaviour, Anscombe is not necessarily committing to the claim that there is an exhaustive 'contrast between what an agent does and what merely happens to him' – that is, between actions and mere happenings – the illumination of which is what Harry Frankfurt defined as the ultimate task of action theory. See Frankfurt, *The Importance*, 69. She may for instance grant, as Frankfurt does in later writings,

the equally important distinction between actions and idly inattentive activity, such as a finger drumming occurring while listening to one's trouble-and-strife. Nevertheless, Anscombe clearly singles out intentional activity as the primary object of investigation.

16. Anscombe, for instance, thought of this as a mark of agency because it was the only way for the whole structure of reasons and actions to be *intelligible*. It is intelligible, she thought, when the structure has a point and terminates in a *desirability characterisation* where the further question 'What do you want that *for*?' ceases to arise. This happens when 'what the agent wants' is 'characterisable as *good by him*'. In other words, it happens when the action conforms to the 'guise of the good'. See Anscombe, *Intention*, 74ff. Donald Davidson puts this down as a fundamental requirement of reason-giving explanations: They must lead us 'to see something the agent saw, or thought he saw, in his action', which it does by displaying 'some feature, consequence, or aspect of the action the agent wanted, desired, prized, held dear, thought dutiful, beneficial, obligatory, or agreeable'. See Davidson, *Essays on Actions and* Events, 3.

17. 'We are active not when we are properly responsive to reasons but when we think that we are', as Joseph Raz says, 'We are active so long as we appear to ourselves to be conducting our life rationally, or as I shall say so long as we live *under semblance of rationality*'. See Raz, *Engaging Reason*, 16.

18. Anscombe, *Intention*, 75.

19. See Davidson, *Essays on Actions and Events*; and Anscombe, *Intention*.

20. Freud discusses this case in the beginning of his classic piece, *Psycopathologies of Everyday Life*. For further discussion and an interpretation of it along the lines I have put forward here, see Setiya, *Reasons without Rationalism*, ch. 1 and Velleman, *The Possibility of Practical Reason*, Introduction.

21. For a similar argument, see Velleman, *The Possibility of Practical Reason*, 119.

22. As Bernard Williams points out, it might be an 'achievement of the modern world to have made it impossible to rear that type, because it has made evil, like other things, a collective enterprise, a process that makes it more powerful but less interesting'. See Williams, *Ethics and the Limits of Philosophy*, 46.

23. Arendt, *Eichmann in Jerusalem*.

24. For an argument to the contrary, see Vetlesen, *Evil and Human Agency*, ch. 1.

25. Korsgaard, *Self-Constitution*, 173.

26. Williams, *Ethics and the Limits of Philosophy*, 44.

27. Setiya, *Reasons without Rationalism*, 22.

28. In particular, the principle that: If some x must be G solely in virtue of being an F, then, in order to understand the nature of Fs fully, we must explain why all Fs are G: either by appealing to the nature of being F, G, or both. See Fine, 'Essence and Modality', 'Semantics for the Logic of Essence'; and Dorr, 'Non-symmetric Relations'.

29. Anscombe, *Intention*, 25. Note that by this formulation Anscombe can allow that the question applies to an action even if there is no positive answer. That is, there could be intentional actions done for literally speaking no reason. Here too, she part ways with Davidson, who identifies acting intentionally with acting for a reason.

30. Anscombe, *Intention*, 11.

31. See note 26 above.

32. See Velleman, *Practical Reflection, The Possibility of Practical Reason* and Setiya, *Reasons without Rationalism*.

33. For the development of these frameworks and their employment as a lemma against normative conceptions of agency, see Velleman, *Practical Reflection, The Possibility of Practical Reason*, and Setiya, *Reasons without Rationalism*.

34. Setiya is explicit on this choice. See Setiya, *Reasons without Rationalism*, ch. 1.

35. Velleman, *Practical Reflection*, 56–7; Setiya, 'Practical Knowledge', 407.

36. Hossack, *The Metaphysics of Knowledge*, 13.

37. Keith Falvey tells a similar story in his 'Knowledge in Intention'.

38. Gettier, 'Is Justified True Belief Knowledge?'.

39. Williamson, *Knowledge and Its Limits*, 31.

40. For the notion of a metaphysical hybrid, see Williamson, *Knowledge and Its Limits*, 50.

41. This is, of course, Williamson's argument in favour of the claim that knowledge is a *prime* (i.e. non-composite) state. Williamson uses this line of reasoning to argue that primeness holds for

other mental states too. This he does in order to show both that knowledge cannot be reductively analysed as a combination of internal and external components, as well as support his thesis that knowledge is a mental state. In what follows, I need only the weaker claim – viz. that knowledge is not a composite – and I leave aside the question of whether it is a mental state or something other. See Williamson, *Knowledge and Its Limits*, ch. 3.

42. The distinction between knowledge in itself and its causal provenance marks a subtle but significant difference between Hossack's and Williamson's models. Williamson still considers knowledge as constituted by a propositional attitude: namely, by the cognitive mental state of knowing. Hossack, on the other hand, regards such features as belonging to the provenance of knowledge, since this is the best way to allow for a wider plurality in ways of knowing. In particular, it counts against Williamson that he rules out the possibility of non-propositional knowledge. I have defended Hossack's position against Williamson in my *Ethics and the Nature of Action*, chs. 5 and 6, and in my 'Kunnskapens Metafysikk'.

43. Vetlesen, *Evil and Human Agency*, 7.

44. My innermost gratitude goes to many people for comments, suggestions and questions to this paper: in particular, to Fredrik Nilsen, Roe Fremstedal, Steffen Borge, Lars Svendsen, Peter Brian Barry, Trine Antonsen, Nadim Khoury, and Michael Morreau, as well as the audience to the conference *Evil and Guilt* hosted by the Ethics Research Group at the University of Tromsø 2014. Great thanks are also owed to the two anonymous reviewers in the *International Journal of Philosophy and Theology*.

Notes on contributor

Heine A. Holmen is currently associate professor at the University of Tromsø in Norway and has previously worked at the University of Oslo, both as associate professor, teacher, and Ph.D. fellow. His thesis, *Ethics and the Nature of Action*, was published in 2011 as a result of his fellowship at the *Centre for the Study of Mind in Nature (CSMN)* in Oslo. It focused on the relationship between action theory, the epistemology of action and issues in meta ethics – in particular, the question 'Why be moral?'. His research interests in general cover, among other things, topics like the philosophy of death and life, action theory, meta-ethics, virtue ethics, the metaphysics of knowledge, the episte-mology of action, and testimony, the nature of practical reason, the metaphysics of evil, philosophy of religion, the epistemology of faith, and the metaphysics of free will.

Bibliography

Anscombe, E. *Intention*. Harvard: Harvard University Press, 1957/2000.

Aquinas, T. *The Summa Theologicæ of Thomas Aquinas: Volume 1*. Scotts Valley, CA: Novantiqua, 2008.

Arendt, H. *Eichmann in Jerusalem. A Report on the Banality of Evil*. New York, NY: Viking Press, 1965.

Arendt, H. *The Life of the Mind: Volume I, Thinking*. London: Secker & Warburg, 1978.

Aristotle. "Nichomachean Ethics." In *The Complete Works of Aristotle – Volume Two*, edited by J. Barnes, 1729–1867. Princeton: Princeton University Press, 1984.

Augustine. *City of God*. Harmondsworth: Penguin, 1984.

Barry, P. B. *Evil and Moral Psychology*. New York, NY: Routledge, 2013.

Chilchott, C. M. "The Platonic Theory of Evil." *The Classical Quarterly* 17 (1923): 27–31. doi:10.1017/S0009838800014804.

Davidson, D. *Essays on Actions and Events*. Oxford: Oxford University Press, 2001.

Dorr, C. "Non-Symmetric Relations." In *Oxford Studies in Metaphysics: Volume 1*, edited by D. Zimmerman, 155–192. Oxford: Oxford University Press, 2004.

Falvey, K. "Knowledge in Intention." *Philosophical Studies* 99 (2000): 21–44. doi:10.1023/A:1018775307559.

Fine, K. "Essence and Modality: The Second Philosophical Perspectives Lecture." *Philosophical Perspectives* 8 (1994): 1–16. doi:10.2307/2214160.

Fine, K. "Semantics for the Logic of Essence." *Journal of Philosophical Logic* 29 (2000): 543–584. doi:10.1023/A:1026591900038.

Frankfurt, H. G. *The Importance of What We Care About*. Cambridge: Cambridge University Press, 1998.

Gettier, E. L. "Is Justified True Belief Knowledge?" *Analysis* 23 (1963): 121–123. doi:10.1093/analys/23.6.121.

Holmen, H. "Kunnskapens Metafysikk." *Norsk Filosofisk Tidsskrift* 3-4 (2014): 190–204.

Holmen, H. A. *Ethics and the Nature of Knowledge*. Oslo: Unipub, 2011.

Hossack, K. *The Metaphysics of Knowledge*. Oxford: Oxford University Press, 2007.

Korsgaard, C. M. *The Sources of Normativity*. Cambridge: Cambridge University Press, 1996.

Korsgaard, C. M. "Morality and the Logic of Caring." In *Taking Ourselves Seriously and Getting it Right – Harry Frankfurt's Tanner Lectures*, edited by D. Satz, 55–76. Stanford: Stanford University Press, 2006.

Korsgaard, C. M. *The Constitution of Agency*. Oxford: Oxford University Press, 2008.

Korsgaard, C. M. *Self-Constitution: Agency, Identity and Integrity*. Oxford: Oxford University Press, 2009.

Plato. "Gorgias." In *Plato – Complete Works*, edited by J. M. Cooper, 791–869. Indianapolis: Hackett Publishing Company, 1997.

Plato. "Protagoras." In *Plato – Complete Works*, edited by J. M. Cooper, 746–790. Indianapolis: Hackett Publishing Company, 1997.

Plato. "Republic." In *Plato – Complete Works*, edited by J. M. Cooper, 971–1223. Indianapolis: Hackett Publishing Company, 1997.

Raz, J. *Engaging Reason*. Oxford: Oxford University Press, 1999.

Reath, A. *Agency & Autonomy in Kant's Moral Theory*. Oxford: Oxford University Press, 2006.

Setiya, K. *Reasons without Rationalism*. Princeton: Princeton University Press, 2007.

Setiya, K. "Practical Knowledge." *Ethics* 118 (2008): 388–409. doi:10.1086/588126.

Stocker, M. "Desiring the Bad: An Essay in Moral Psychology." *The Journal of Philosophy* 76 (1979): 738–753.

Svendsen, L. F. *A Philosophy of Evil*. Champaign: Dalkey Archive Press, 2010.

Velleman, D. J. *The Possibility of Practical Reason*. Oxford: Oxford University Press, 2000.

Velleman, D. J. *Practical Reflection*. Stanford: CSLI Publications, 2007.

Vetlesen, A. J. *Evil and Human Agency*. Cambridge: Cambridge University Press, 2005.

Williams, B. *Ethics and the Limits of Philosophy*. Cambridge, MA: Harvard University Press, 1985.

Williamson, T. *Knowledge and Its Limits*. Oxford: Oxford University Press, 2000.