

Why ethics is hard

To discover the right style [in moral philosophy] is to discover what you are really trying to do... The aim is to sharpen perception, to make one more acutely and honestly aware of what one is saying, thinking and feeling.

If I could tell you what it meant, there would be no point in dancing it.¹

1. Three questions compared

One worthwhile project in ethics is to explore the possibilities of the notion of the *moral imagination*: to think about the possible shape of an imaginatively-enriched account of how we can know what to do, of how to make good decisions and act well. In pursuing this project we need to keep it in mind that to talk about “the moral imagination” can be, and should be, to talk about many things, not just one. One thing that talk about the moral imagination certainly does not exclude is the possibility of moral knowledge. In this paper I make some moves—perhaps small and tentative moves—towards a conception of moral knowledge by thinking about the place of experiential knowledge in ethics. (If that sounds like a different topic from “the moral imagination”, read on to see the connections.)

My way of approaching this begins with three cases of question and answer.

First question. I show you some liquid, water as it happens, and you ask: “What is the property in virtue of which this is water?” A good question, with one good simple answer which even philosophers know: it’s water because it’s H₂O.

Second question. I show you another liquid, an acid this time, and you ask: “What is the property in virtue of which this is an acid?” A good question, though as philosophers may or may not know, current chemical theory implies that there is no straightforward answer to it. Wiki tells us that there are three common definitions for acids (Wikipedia article “Acid”, <http://en.wikipedia.org/wiki/Acid>, accessed 30.3.11): this means that when informal talk about acids is tightened up by modern chemical theory, we find that we are talking about three different chemical properties which largely but not completely overlap in their extensions. So there isn’t really just *one* property in virtue of which anything is an acid (unless you count as one property the disjunctive property “fits the Arrhenius definition or the Brønsted-Lowry definition or the Lewis definition”). So the question “What is the one property in virtue of which anything is an acid?” has no straightforward answer, but at least it leads us in the right direction to see that things are a bit more complicated when it comes to defining “acid” than they are in the case of “water”.

¹ My first epigraph is a *mélange* from two different editions of Bernard Williams’ *Morality: an introduction to ethics* (Cambridge UP 1972)—respectively from p.11 in the first edition’s Preface, and from p.xv of the new Preface that Williams wrote for the 1993 reprint of *Morality* (also CUP). My second epigraph is a quip widely attributed—I haven’t found a source—to the American dancer Isadora Duncan.

Now for the third question. Here is Jan Brueghel the Elder's beautiful painting *Eden*:



In virtue of what property is this picture beautiful? (I hope the reader agrees with me that it *is* beautiful. If she doesn't, I invite her to pick a picture that she thinks is beautiful, and we can talk about that. If she thinks *no* pictures are beautiful, I don't know what to say.) The thing about this third in-virtue-of-what-property question is that any good answer to it will have to be interestingly disanalogous to good answers to either of my first two questions.

A good answer to this question could not be analogous to the good answer to the question about water, because there is no *one* property in virtue of which *Eden* is beautiful. (Or at any rate there is no one *interesting* property in virtue of which it is beautiful. You could say that it's beautiful in virtue of having the property of beauty, but why bother?)

Nor could a good answer be analogous to the good answer to the question about acid. While modern chemical theory strictly speaking recognises no one property corresponding to the folk's and earlier chemists' talk of "acid", it does strictly speaking recognise three properties the disjunction of which corresponds, and so far as we know exactly corresponds, to the folk extension of "acid" (any inexactitude will be the folk's, not the chemists'). Now of course there are philosophers who have said that it is the finer-grained aesthetic properties that we should focus on: in J.L. Austin's famous phrase, the dumpy and the dainty rather than the beautiful. That might look like an analogy with the case of acid, but it isn't really, for at least two reasons. First, in the acid case we have discipline-wide agreement among chemists that there are

exactly three properties in the disjunction, whereas in the case of beauty there is very little consensus at all on what properties should be in the corresponding disjunction, on how many properties there should be in this disjunction, or indeed whether we should be analysing beauty by way of a disjunction (or by way of anything else) at all.

Secondly, the three chemical properties in the disjunction that analyses “acid” are quite different in character from anything we might put into a parallel disjunction of aesthetic properties to analyse “beauty”. The trouble with “acid” as a scientific term was that, unlike “water”, it corresponded to no one precisely-defined chemical property. The three-way disjunction for “acid” fixes this problem, and (according to the present best theory of chemistry) fixes it definitively, by tying the extension of “acid” down to the extension of three chemical properties which are precisely-defined. Any likely disjunction that could be offered for “beautiful” may tie its meaning down to a variety of more particular properties. But even if (as seems very unlikely) this disjunction has a definite and finite length, these more particular properties will never be *precise*, in the way that the three properties in the chemical disjunction are precise. In fact, the more particular properties will be imprecise in just the same way as the original overarching property of beauty was imprecise.

Three popular responses to this puzzle about beauty are dogmatism, obscurantism, and subjectivism. Dogmatists find a formula, disjunctive or not, that gives a finite definition of beauty, and insist on it, explaining away all counter-examples with rather too much abandon to convince. Obscurantists tell us that there *is* no formula, that beauty is real but indefinable, end up (very often) with little else to say, and also fail to convince. And subjectivists agree that there is no formula for beauty, offer a reductive explanation of why not that makes beauty consist in our reaction to certain things, not in any property of those things, simple or disjunctive—and convince all too often.

I reject all three of these responses to the puzzle, and want to suggest a fourth. The key to the response that I shall offer is the notion of a *phenomenal property*.

2. *Phenomenal properties*

The most familiar example of a phenomenal property is *redness*:

Mary is a brilliant scientist who is, for whatever reason, forced to investigate the world from a black and white room via a black and white television monitor. She specializes in the neurophysiology of vision and acquires, let us suppose, all the physical information there is to obtain about what goes on when we see ripe tomatoes, or the sky, and use terms like ‘red’, ‘blue’, and so on.... What will happen when Mary is released from her black and white room or is given a colour television monitor? Will she *learn* anything or not? It seems just obvious that she will learn something about the world and our visual experience of it. But then is it inescapable that her previous knowledge was incomplete. But she had *all* the physical information. *Ergo* there is more to have than that, and Physicalism is false. (Frank Jackson, “Epiphenomenal Qualia”, *Philosophical Quarterly* 1982, 127-136, at p.130)

I don't think that "the knowledge argument" here presented is a sound argument against physicalism. (Neither does Jackson any more, though I and others find the reasons he now gives for rejecting it both puzzling and obscure.) The knowledge argument attacks a formulation of physicalism that says that "All the information there is to know about anything is physical information". But if "physical information" is narrowly interpreted, as "information of the type that physicists are professionally concerned with", then we do not need the knowledge argument to tell us that physicalism is false; a bus timetable will do the trick, or indeed a book on chemistry. Whereas if "physical information" is broadly interpreted, say as "information about the material world in general", then it is unclear why the information that "Experiencing redness feels like *this*" could not be "physical information". A better formulation of physicalism could say something like "All the facts there are either are facts of physics, or supervene on the facts of physics". But what Mary discovers when she leaves her monochrome room still need be no counter-example to this: perhaps what she discovers supervenes in an *a posteriori* way on the facts of physics, or perhaps (as some will argue, e.g. proponents of the Nemirow-Lewis "ability argument") it is not a fact of any sort. Physicalism so formulated is (I think) clearly false for a quite different reason, namely that at least some facts are neither physical facts nor supervenient on physical facts: $2 + 2 = 4$, for instance. Maybe physicalism is still false or unpersuasive even if we except mathematical facts from its scope: for example, maybe the thesis of supervenience can only be true, if it is true at all, at such a global level that its assertion is more like a statement of faith than an empirical discovery (Dupré 2010: 292-3). But however many ways there may be to show the falsity of physicalism, the knowledge argument isn't one of them.

The knowledge argument, then, does not achieve what it was originally supposed to achieve. Yet it does achieve *something*. What the knowledge argument brings out is a distinction between two different kinds of content. One of these is the sort of things that a scientific theory of colour could tell you; the sort of things that, where you are correctly informed, count as factual knowledge, knowledge that, propositional knowledge. The other is the kind of new information (or rather content) that Mary acquires on release from her room: new experience; acquaintance with new objects; knowledge what it's like to see red; what I call phenomenal content. Non-phenomenal content is the kind of thing we get by knowing *about* the Taj Mahal, or the colour red; phenomenal content is the kind of thing we get by knowing *the Taj Mahal*, or *red*.

Clearly the distinction between phenomenal and non-phenomenal consciousness, or content, bears at least some comparison with two other historically famous philosophical distinctions: Bertrand Russell's very explicit distinction between propositional knowledge and knowledge by acquaintance, and Thomas Nagel's more or less explicit distinction between other kinds of knowledge and knowledge "what it's like". Jackson himself wants to put some distance between the knowledge argument and Nagel's kind of argument (Jackson 1982: 131-2); it's a pity he doesn't also discuss Russellian acquaintance—the comparison would have been interesting and instructive. Anyway, it looks like all three philosophers are pointing towards a fairly intuitive distinction between knowledge of facts and knowledge by experience, between non-phenomenal and phenomenal content. It is that intuitive distinction that is my real concern here.

As Jackson says, you can get a full non-phenomenal understanding of redness just by learning the complete physics of colour. This won't give you any phenomenal understanding or knowledge of redness. A phenomenal understanding or knowledge of redness can only come from experiencing redness. The basic mechanism in explaining or teaching the concept of (phenomenal) redness is not definition or theoretical explanation. It is ostension: "redness is like *this*". The content of propositional knowledge is indeed, as the name suggests, basically propositional: it is a matter of logical relations between propositions that are essentially expressible in words. By contrast, the content that phenomenal experience gives you, although we usually can talk about it in words, is not essentially expressible in words. Propositional knowledge *must* be expressible in propositions, because propositions are—so to speak—what propositional knowledge is made of. By contrast, phenomenal knowledge *may* be expressible in propositions, but for one thing they will be a different sort of proposition—more about that in a moment—and for another thing, it does not *have* to be. There is a contradiction in the very idea of propositionally-inexpressible propositional knowledge; that would mean, roughly, that there is some sentence that can be framed that can't be framed, which ~~is would be~~ absurd. There is no corresponding contradiction in the very idea of propositionally-inexpressible phenomenal knowledge; that would mean that there is some experience that we can't frame any (informative) sentences about, which is not only not absurd, but may even happen. (Sometimes: I'm not suggesting this is typical.)

With propositional knowledge, the correct form of words, once you arrive at it, gives you complete and definitively correct understanding of the subject-matter: in a sense, once you have got there, there is nothing more to be said. The knowledge *is* the words.

With phenomenal knowledge, by contrast, complete and definitive understanding of the subject-matter can only come from having the relevant experience (or experiences). We can use words to ostend those experiences: "*this* is red here". Or we can use words to construct analogies with other experiences that are interestingly similar: if our audience have had these other experiences, they may grasp the analogy ("scarlet is a red, a very bright and pure red with no trace of either pink or purple in it"; or think of how we may get from other blues to Hume's missing shade). Again, we can use words to do more complicated metaphorical or otherwise evocative work (like the blind man in Locke, who said that red must be like the sound of trumpets). What we cannot do is use words not ostensively but discursively, to capture the essence of the phenomenal in just the same way as we can use them to capture the essence of the propositional. The knowledge is essentially *not* the words; it is non-verbal.

Here, I suggest, lies the solution to our puzzle about the question "What is the property or properties in virtue of which Brueghel the Elder's *Eden* is a beautiful painting?". The difficulty raised by that question is not that dogmatism is right, and there is some simple formula to answer it that we are just missing (Clive Bell's "significant form", perhaps, or the "golden proportion" discussed in architectural theory). Nor is the difficulty that obscurantism is right, and we can't say at all what makes it beautiful—it just *is*. Nor again is the difficulty that subjectivism is right, and what we call the beauty of the picture is really just a fact about us. The difficulty is that the properties in virtue of which *Eden* is beautiful are (mostly) its phenomenal

properties, and the best way to get at these is not to try and catch them in a definition, but to ostend them. The beauty of the picture is something that is there to be experienced. You can say lots of right things about the beauty of Brueghel's painting, but none of these is *the one* right thing to say, in the way that H₂O is the right thing to mention if you are trying to give the essence of water. Unlike water, and indeed acid, there is something inexhaustible about *Eden*: there is no end of good and interesting things you could go on saying about it to explain its value, all of which might be worthwhile and enlightening, but none of which would have the conclusiveness and definitiveness of a chemical formula or a scientific definition.

3. *Broadening the notion of phenomenal content*

I began with a puzzle about trying to state the value-grounding property of a picture in the same way as we state properties in science: a puzzle to which I have offered a resolution that depends on developing a notion of phenomenal content, the kind of content that experiences characteristically have. For phenomenal content, the order of analysis that applies with propositional content is reversed: with propositional content we find a formula which gives the essential and underlying explanation (insofar as they have one) of the various things that we might incidentally ostend, but with phenomenal content we find something to ostend, and this ostended thing turns out to be the essential and underlying explanation (insofar as they have one) of the various formulae that we might incidentally formulate.

How wide does the notion of phenomenal content go? Some examples of phenomenal content are uncontroversial. Jackson's own "what it is like to see red" is one; Mark Johnston's cases of seeing canary yellow and feeling nausea are two more (Johnston, "How to speak of the colours", *Philosophical Studies* 1992); the smell of almonds and the sound of car doors slamming and the pain of a broken leg and the sound of a major-seventh chord and the sense of vertigo you get when you look down from the viewing platform of the Aiguille du Midi... any sort of sensory experience will do. All of these are cases where you can, to echo Jackson, know all the scientific or theoretical facts in advance about what it means, how it works, and what it is, and still not know anything much at all about the experience itself.

Phenomenal contents occur in the typical operation of each of our five basic perceptual senses. As I am equally happy to say: for each of our senses, there is something it is like for us to perceive through that sense. Or rather, for each perceptual sense there is a whole genus of "somethings it is like", of phenomenal contents. There are all sorts of things we can see, but there is a unity to our experiences of seeing which is difficult to give a name to beyond saying, rather lamely, that they are all visual experiences. Likewise with each of the other senses: our particular experiences of e.g. smell and touch can of course vary enormously, but they are all generically unified by being, well, olfactory or tactile. (Presumably a bat's particular echolocatory experiences are generically unified in the same sort of way.) Phenomenal contents not only have particularities on which we can reflect in various ways, for example by comparing them ("this taste is the same as that taste"; "I've never smelt this smell before", "this is a bit like the feel of silk, but also a bit rougher", and so on indefinitely); they also come in modalities.

Do the phenomenal contents that we experience come in just five modalities, the familiar five senses of smell, taste, touch, sight, and hearing? Since I've already given vertigo as one of my examples of a phenomenal content, my answer is obviously No. It is now a familiar point in modern psychology that the list of the five senses is incomplete anyway. Think about the variety of other ways in which we are phenomenally aware of our own bodies: the loose bundle of sensory abilities that get lumped together under the name "proprioception", such as our sense of where our limbs are and what they are doing relative to each other and to our bodies and environments, our sense of which way is down, and our sense of whether we are moving or still—all of these have their own particular what-it-is-like-nesses, their own phenomenal feels. Again, physical pain and physical pleasure are obvious cases of phenomenal content that do not fit under the heading of any of the five senses (in particular they do not fit under touch); whatever else ethical hedonists are wrong about—and there's plenty—at least they see the ethical importance of at least some sorts of phenomenal content. Then there are moods like euphoria or depression, and bodily states like feeling sleepy or lively, hungry or bloated, thirsty or restless: all of these forms of bodily awareness, as we might call them, have their own particular phenomenal contents, none of which seems straightforwardly reducible to the phenomenal content of any one or any combination of the traditional five bodily senses.

Moving outwards from our basic senses, via our various forms of bodily awareness, we thus quickly come to our moods and then to our emotions. And in these cases too we uncover more and more examples of what-it-is-like-ness that can't be reduced to anything simpler, or even to any combination of simpler elements. Fear, anger, joy, sadness, anxiety, depression, elation, tiredness, and relief all have their own distinctive phenomenal contents; none of these contents has any particular tie to any one of the basic five or six senses.

Notice the importance of the conceptual loading of phenomenal feels. My first examples of phenomenal contents were the smell of almonds, the sound of car doors slamming, the pain of a broken leg, the sound of a major-seventh chord, and the sense of vertigo you get when you look down from the viewing platform of the Aiguille du Midi. Phenomenal contents can be non-conceptual, but don't have to be; it is an interesting question how conceptual awareness affects or transforms these phenomenal contents. Do almonds smell different if you know it is almonds you are smelling? In a sense, no—it is, after all, "the same smell" that you had before that you are now labelling as the smell of almonds. In another sense, yes. It is hard to say how, but it does seem to be true that labelling and conceptually categorising even this rather simple and basic experience brings it into all sorts of connections that it did not stand in before; after all, that's the whole point of labelling it. (And then there are the contrastives: if you know it's almonds that you're smelling, then you know it's not cyanide. Circumstances in which this would be a nice thing to know~~comfort~~ are readily imaginable.) With time and training, it seems, rather mysteriously perhaps, that the connections imported into our experience by categorisation and conceptualisation can come to be built into the phenomenal content itself.

With less simple and basic experiences, it is often obvious that the accompanying conceptual content does not merely *alter* the quality of the phenomenal content: it creates it. In order to have vertigo, you have to have the thought, though perhaps not

the belief², that you are at risk of falling; in order to experience a sound *as* a major-seventh chord, you have to have the concept of a major-seventh chord, a concept which typically presupposes quite a lot of musical training. With driving and skiing too, understanding what you are doing completely transforms your experience of what it *is* that you are doing: whereas inexperienced drivers have to listen to the engine-note and look at the rev counter and the speedometer to know when to change gear, for experienced drivers knowing when to change gear is—literally—a matter of feel. What the experienced driver has got that the inexperienced driver lacks is a new form of phenomenal content.

We can *learn to experience* vertigo, a major-seventh chord (as such), and the right moment to change gear. And in each of these cases, “learn to experience” can be taken quite literally: alongside our acquisition of new concepts and a new vocabulary, we also learn to access a new phenomenology. There is a distinctive phenomenal content (or better, a large family of phenomenal contents) in each case which is only available to those with a certain training or acculturation.

4. *Paradigm experiences and moral experiences*

What is true of driving, vertigo, and listening to music is also true for a wide range of what we may call the paradigm events or experiences of human life. Being a child, being a parent, being a friend, being a lover; feeling affection or jealousy or hope or suspicion; seeing someone born, or seeing someone die; hurting someone else, or being hurt; achieving an ambition, or failing to; giving or receiving sympathy, advice, admonition, warning, or a scolding; working, or being idle; keeping or betraying someone’s trust; protecting someone from injury or death—or injuring or killing someone... I call these the paradigm events or experiences of life, because the basic shape and the basic possibilities of human life can be derived from an open-ended list something like this one (only much longer and, for much of its length, more specific.)

~~W~~And as before, when I call these sem paradigm *experiences*, I mean it literally. My thesispoint is that life’s paradigm experiences or events have distinctive phenomenal contents (more accurately, phenomenal-content *types*) associated with them. (Not necessarily one-one, of course: there can be lots of different phenomenal contents for one paradigm experience, and different paradigm experiences can share a phenomenal content.) Understanding these paradigm experiences or events means understanding their distinctive phenomenal contents. Now understanding life’s paradigm experiences and events is exactly what ethics is centrally about. Therefore, ethics is centrally about understanding the distinctive phenomenal contents of life’s paradigm experiences and events. Ethics is all about knowing what it’s like to be a child or a parent, a friend or a lover, to hurt or be hurt, to succeed or to fail, to sympathise or be sympathised with, to betray or be betrayed, to protect or to kill. So no one who lacks a grip on phenomenal contents like these can hope to understand what ethics is all about; any more than someone who has never experienced the colours can hope to understand what scarlet or sky-blue or canary yellow is all about.

² On the place of belief in emotion see e.g. Robert Roberts, “What an emotion is: a sketch”, *Philosophical Review* 97 (1988), 183-209.

Is there really such a thing as a distinctive phenomenal content or set of contents that goes with, say, becoming a parent? Don't people respond (or fail to respond) in such a variety of ways to life's experiences and events that there is no way of regimenting all this diversity into anything like a shape? I admit the diversities, indeed I am happy to do so: life is (of course) complicated, and I'm not trying to simplify it. But complicated though life may be, it does have recognisable themes or groups of themes. (In the theory of emotions something similar to my list of "paradigm events of life" has often been called a list of the emotional themes; however, what I am doing here is not the theory of emotions.) A recurring phrase in Raimond Gaita's work is "what it means" or "what it is": what it means to be a racist or to kill someone, and again what it means to love someone, or simply what it means to be human. Gaita's phrase presupposes commonalities, and I do not think his presupposition is a mistake. Some generalisations about human life are true, and the themes that fit them are not the less recognisable for being subject to infinite variation. Both the themes, and the ways in which those themes can be varied, are and always have been the subject-matter of all human narrative art.

It is also true that there is such a thing as *failing to be properly responsive* to life's paradigm experiences. Just as there is such a thing as a "functionally correct" experience of redness, or at any rate a range of "functionally correct" experiences, so in human life there is such a thing as the way you *ought* to feel in situation X—or at any rate a range of ways. People who miss out on the full phenomenal content of some paradigm experience are indeed missing out.

But are there really *objective* phenomenal contents of this sort—aren't they *just our feelings*, our reactions? I believe in the transparency of experience to the event experienced—that's why I call them "paradigm experiences *or* events". These things are genuinely experiences, but that doesn't make them private: they are both experiences in us, and also events in the world. In that sense they're both subjective (in the subject) and objective (in the object). It is because of this duality that I think our awareness of phenomenal content is more like perception than reaction.

What about the familiar relativity worry, given that I've admitted that there is an element of learnedness and cultural conditionedness to most of our more sophisticated phenomenal contents? Well, learnedness and conditionedness do not equal unreality: the notion of a "trained observer" is a key one in the philosophy of science, precisely because there are important physical phenomena that one needs to be taught and conditioned to detect at all ("there goes a proton"). Likewise in the ethical life, there is moral expertise in what one detects as much as in what one does.

What if other people just claim to see different phenomenal contents from the ones I think I see? Then the best way forward is to keep looking, and to try and look more closely, and to see if their descriptions of what they think they are seeing work better than my own. Careful attention to the phenomenology can often make people who up to now have been quite sure what they are seeing, come to realise that actually they haven't been seeing that at all—they've only thought they were seeing it. George Orwell famously modified his view about what it meant to be a soldier when he saw a Spanish fascist running away holding up his trousers as he ran; however evil one takes Franco's fascism to be, one cannot, Orwell found, see something like that and go on

seeing the servants of that fascism as mere ants to be crushed or obstacles to be blown out of the way.³

5. *Capturing moral experiences through art*

If you want to understand what it is like to see red, you need either (and ideally) to see red for yourself, or else to find some way of imaginatively summoning red up. Just likewise, if you want to understand what it is like to be a child or a parent, to hurt or be hurt, to succeed or to fail, to protect or to kill, or to have any of the other paradigm experiences of the moral life, then you need either to have these experiences for yourself, so as to grasp the phenomenal contents in question directly, or else find a way of getting hold of them by way of the imagination. And just this, as Shelley famously pointed out, is what poetry does for us:

Man, to be greatly good, must imagine intensely and comprehensively; he must put himself in the place of another and of many others. The pains and pleasures of his species must become his own. The great instrument of moral good is the imagination; and poetry administers to the effect by acting upon the cause... Poetry strengthens the faculty which is the organ of the moral nature of man, in the same manner as exercise strengthens a limb. (Percy Bysshe Shelley, *A Defence of Poetry*)

In a true understanding of the colours, what has to come first is not definitions of the colours but samples of the colours; the first thing to understand about red is what it is like to see red, and the only way to get hold of that is by experience. Similarly, in a true understanding of the paradigm events of the moral life, what has to come first is not definitions of those events or their moral value, but samples of those events. The only way to get hold of such events is by experience, either actual or imaginative. And one thing poetry does, along with other forms of art, is precisely to ostend and explore just such samples of experience.

...you write because you have to. If you rationalise it, it seems as if you've seen this sight, felt this feeling, had this vision, and have got to find a combination of words that will preserve it by setting it off in other people. The duty is to the original experience... [a poet's writing must] be born of the tension between what he non-verbally feels and what can be got over in common word-usage to someone who hasn't had his experience. (Philip Larkin, *Required Writing* 58, 82)

The poet's duty is to the original experience, says Larkin. Ethics as much as poetry is about making sense of experience, of phenomenal content. In ethics too, we might say, experience comes first, in the strict sense that the foundation of ethical thought—insofar as it has any *one* foundation—lies in the paradigm events of life, and hence not in definitions or other forms of words, but in the ostension and imaginative exploration of phenomenal contents.

³ George Orwell, 'Looking Back on the Spanish War' (1942) in *George Orwell: Essays*, London: Penguin Books, 2000. I am grateful to Raymond Boyce for reminding me of the passage.

This observation has, I think, rather widely subversive consequences for philosophical ethics as commonly done today. From Socrates on, ethicists have usually sought to base ethics on verbal formulae: their idea is that once we know the definition of justice, or the reason why killing is wrong, or why persons are morally significant, or... , then we will have got hold of a form of words which conclusively and definitively settles matters, which explains the theoretical shape and structure of ethics, or of this fragment of ethics, so fully and lucidly as to complete our inquiry, in just the sort of way that H₂O completes our inquiry into the essence of water. But if the business of ethics is more like the business of aesthetics—of understanding, for example, what makes Eden such a great picture—then there will be something wrong with all questions in moral philosophy that ask something like “What is the property in virtue of which” murder is wrong, or human life valuable, or being in love with someone a reasonable response to her reality, or ancient British oak woods worth protecting. All such “Socratic”—we might also say sophistical—questions will betray a fundamental misconception of what certain sorts of basic ethical understanding could possibly be. What a pity it is, then, that we philosophers keep asking them.

I think Plato's *Republic* and Aristotle's *Nicomachean Ethics* are both, in their extremely different ways, all about rejecting this Socratic conviction that the basis of ethics lies in definitions rather than in experience—though to develop that interpretation of those two *chefs d'oeuvre* in the detail it deserves would be another story, and a long one. Despite Plato's and Aristotle's campaigns against it, the Socratic conviction is deep, tenacious, and pervasive even to this day: it explains a very great deal in the way contemporary moral philosophy is usually done. I invite readers to supply their own analogues in contemporary moral philosophy of the positions that Section I called dogmatism, obscurantism, and subjectivism. TAgain, the Socratic conviction explains a good deal about why so much of moral philosophy today is a contest of the theories, a contest in which there can be only one winner with all other theories condemned as wrong: “Is the Kantian's explanation of XYZ right, or is it the consequentialist's? Is ABC a counter-example to contractualism, or can contractualists find a non-*ad hoc* way to handle ABC? Which of the ethical theories, overall, is *right*?”. It shows us the roots of what we may call the curse of the definite article: it is because moral theorists are looking for *one* theory of ethics, which is to consist in a definitive and conclusive form of words or sets of forms of words that ultimately will cover everything and exclude all competitor theories, that they look in each case for *one* explanation, *one* reason why, *one* value-basing feature, *one* fundamental property, and so on. But, as we have seen, not even chemists always find that.

What we actually have to deal with in the ethical life, it seems to me, is both richer and more interesting—but also, much harder to talk about. The paradigm experiences of the moral life present themselves to us, not as things that we can tie down in one neat definition, but as inexhaustibly redescrivable: saying this or that may bring out some aspect of those experiences' significance, but none of the things we might say about them is the definitive and conclusive formula that just wraps them up, in the way that, say, the chemical formula H₂O does just wrap up the real chemical nature of water.

H₂O/ water, of course, is also responsible for a whole multitude of interesting phenomenal contents. H₂O can be wet and soft and shimmeringly reflective (when not

frozen), or fluffy and powdery and perfect for S-line skiing (in the form of snow), or toffee-thick and chewy and grand for slotting ice-tools into (in the form of best Scottish winter-gully névé). That suggests a question about the line of argument I'm developing here. Both in the case of water and in the case of red, there is undoubtedly something scientific to say: water is H₂O, red is a range of light-frequencies in the colour spectrum between such and such a number of measurement and such and such another number. And there is (presumably) some sort of explanatory relationship between what we can say at the scientific level, and what shows up at the level of phenomenal content where water has the properties I mention above. Mightn't something parallel be true in ethics? Mightn't the real business of ethics lie in finding and explaining the explanatory relationships, whatever they are, between the phenomenal contents that we know about, and the scientific subveners of those phenomenal contents—whatever they may be?

My answer to this question is simply “Bring it on”. If there is such a theory to be had, by all means let's have it. I only observe, first, that such a theory would be a descriptive psychology (and perhaps neuropsychology) of ethics, not a normative ethics: it would explain the physical basis for our having the ethical experiences that we do, not modify, except perhaps around the edges, the ethical experience itself. Secondly, I observe that we currently have no such theory, and no idea how to get one either. We don't even have a theory of how to get from the physical basis of colour experience to the phenomenal content of colour experience, though no doubt a Just-So story can be told about why it was useful for us to evolve a capacity to represent phenomenal contents as well as (and probably before) a capacity to represent propositional contents. Much less do we have any idea of what it would look like if we had a theory of how to get from the physical bases of our paradigm experiences to the phenomenal contents of those experiences. Anyway, even if we got hold of such a theory, that theory still wouldn't be the *only* business of ethics, any more than the theory of colour is only concerned with the scientific basis of colour experience, and not interested at all in colour experience itself. The scientist is naturally, and rightly, curious about both; she also takes the facts about colour experience that have always been available to us simply by careful introspection to be in at least some ways a limit on what a scientific theory of colour experience could be. (That is one reason why it might make sense, *pace* Jackson 1982, to treat what Mary learns as “physical information” too.) Just likewise, a theory that explained the scientific basis of the phenomenal contents that are central to our ethical lives would never be the full story; even if available, it would just be part of the story, and the phenomenal contents themselves would continue to supply crucial other parts.

Moral experience is, in the sense just explained, radically prior to moral theorising. What is foundational (if any one thing is foundational) is not a uniquely correct and definitive propositional formula, but the ostension of samples of moral experience. We can talk about these, either by ostending them or by making analogies or by metaphor and evocation, but what we cannot do is capture their essences in a definition. If that claim is right, then we should expect to find that moral experience too, at least in its central and most significant parts, is inexhaustibly redescribable in something very like the way our experience of Brueghel's *Eden* is inexhaustibly redescribable.

And this, I think, is exactly what we find. It is no more reasonable to look for the one uniquely correct defining formula to capture the moral importance, say, of human life than it is to look for the one “uniquely correct” poem about snow. There are many fine poems about snow, and no doubt an even greater number of bad poems about snow. But even if God is a poet, there is simply no such thing as the *one right* poem about snow. Neither is there any such thing as the one right thing to say about, for instance, what is valuable about human life, or what is wrong with murder, or why it can make perfect sense to love someone. Just as with Brueghel’s picture, the best answer to such questions—very commonly asked by philosophers—is often to refuse them, and look elsewhere. It is not to seek a single all-embracing definitional formula that is supposed to provide us with our one definitive way of capturing the phenomenal content that we are trying to talk about, but rather to ostend some examples of that phenomenal content.

~~Wi~~For th this thesis of mine about the basicity of ostension cp. Linda Zagzebski’s “exemplarism”:

A moral theory consists in part of a system of concepts. Some concepts in the theory are defined in terms of others. But... unless we are willing to accept conceptual circularity, some concept or concepts will either be undefined or will refer to something outside the domain. [For m]ost moral philosophers [the] basic evaluative concept in their theory is defined in terms of something allegedly non-evaluative, such as human flourishing in the biological sense. The alternative I am suggesting is to anchor each moral concept in an exemplar. Good persons are persons *like that*, just as gold is stuff like that. The function of an exemplar is to fix the reference of the term ‘good person’ or ‘practically wise person’ without the use of any concepts, whether descriptive or non-descriptive. Reference to an exemplar then allows the series of conceptual definitions to get started. So the system... is linked to the world the system is about by indexical reference to a paradigmatically good person. As with other theories based on direct reference, indexical reference is the hook that connects our theory to that part of the world with which the theory is concerned—in this case, the ethical domain. (Linda Zagzebski, “The admirable life and the desirable life”, in T.Chappell, ed., *Values and Virtues* p.59)

It is often an advantage for philosophy to be, not about our concepts of the world or about our language about the world, but about the world. Zagzebski’s exemplarism achieves exactly that directness: “The function of an exemplar is to fix the reference of the term ‘good person’ or ‘practically wise person’ without the use of *any* concepts”.⁴ The exemplars come first, the concepts come later. And the concepts are beholden to the exemplars, not the exemplars to the concepts: it is the exemplars that track truth, and the concepts that follow. If further deepening of our understanding of some exemplar leads to conflict between that new understanding and the concepts we use about the exemplar, then typically (not always but typically) it will be the concepts, not the exemplar, that need to be revised. We get our understanding of what

⁴ As I argued in Gill 2007, something like Zagzebski’s exemplarism not only gives us a better way of understanding what Aristotelianism in ethics might be than, for example, the biological naturalism of Philippa Foot. ~~I~~ It is also a genuinely *Aristotelian* position, one that Aristotle himself is at least close to.

it is to be a good person, not primarily from conceptual analysis of “good” and “person”, but from good persons: from exemplars, real or fictional, that—give us a descriptive richness and denseness of texture that feeds the moral imagination as no definition could. If you want to call such exemplars *icons*, there is no harm in that either, since the idea of icons in the religious sense of the word depends on the same point about the priority of direct experience over conceptualisation that I, in company with Zagzebski and (I believe) Aristotle, have been developing. In religion too, the point is that icons are meant to confront us with realities that outstrip any theorising we might do about those realities.

The now-familiar philosophers’ question “What’s wrong with murder?” brings out this point about the basicity of ostension to ethics particularly vividly. As Rai Gaita (*Good and Evil* 33-4) famously observes, there is something strikingly banal and bathetic about what most recent philosophers have found to say about the wrongness of murder: “‘My God, what have I done? I have been a traitor to reason. I have violated rational nature in another!’; ‘My God, what have I done? I have violated my freely chosen and universally prescribed principle that one shouldn’t kill people under circumstances such as these!’”. To date, on the wrongness of murder the real masters are not Kantians or Hareans, consequentialists or contractarians. They are Shakespeare and Dostoevsky:

Had I but died an hour before this chance,
I had lived a blessed time; for from this instant
There’s nothing serious in mortality—
All is but toys; renown and grace is dead;
The wine of life is drawn, and the mere lees
Is left this vault to brag of.
(William Shakespeare, *Macbeth* 2.3.89-94)

So he lay a very long while. Now and then he seemed to wake up, and at such moments he noticed that it was far into the night, but it did not occur to him to get up. At last he noticed that it was beginning to get light. He was lying on his back, still dazed from his recent oblivion. Fearful, despairing cries rose shrilly from the street, sounds which he heard every night, indeed, under his window after two o’clock. They woke him up now.

“Ah! The drunken men are coming out of the taverns,” he thought, “it’s past two o’clock,” and at once he leaped up, as though some one had pulled him from the sofa.

“What! Past two o’clock!”

He sat down on the sofa—and instantly recollected everything! All at once, in one flash, he recollected everything.

(Fyodor Dostoevsky, *Crime and Punishment*, tr. Constance Garnett, p.84)

Macbeth speaks hypocritically—the speech I quote is his first articulate attempt to do what he must always do from now on: be a “poor player”, keep “the mind’s construction” out of his face, present himself as the sincere mourner that he is not. Yet even as he pronounce-mouths these words, is sanctimonious hyperbole intending them to be false, we can feel him realising that actually terrible truth of his sanctimonious own-self-description has a horrible truth to it. His life *has* lost its meaning, or at any rate the meaning it had before; if it has a meaning now, it is the meaning of a horrible

parody of what a king's life should be. Because Macbeth has “broke ope the Lord's anointed temple” (2.3.65-66), turned Duncan's body into an empty corpse, his life too has become as empty and trivial as a corpse. In everything he will ever do from this time on Macbeth, like Raskolnikov, is *poisoned*: as Raskolnikov himself says, with some self-centredness but also an important kind of truth, “I killed myself, not that old creature” (402). There is nowhere to hide for either of them from the coarse and ugly fact of what they have, irreversibly, made of themselves. Macbeth “hath murdered sleep” (M 2.2.42): the murderer's sleep turns, as Raskolnikov's does, to a feverish mockery of restfulness. (Macbeth and Raskolnikov are both invited to murder by the bell, too (M 2.1.62, C&P 70): both are hurried, breathless, harried by what they feel as the inexorable demands of time.) Raskolnikov like Macbeth finds himself caught in a vain quest to for the forgetfulness of what he has done that is the only thing left to him that resembles the relative innocence—and how enviable it seems now to Raskolnikov—of the ordinary drunken prole.

These are the things—or some of the things—that we can say to spell out the wrongness of murder, what it means to be a murderer. Obviously they are not the only things, and that is not just because these are two specific artistic depictions of what it is to be a murderer; it is also—to say it again—because nothing in this direction counts as completeness or definitiveness. As you might put it, what we need—whether from novelists fulfilling their distinctive *métier*, or from philosophers fulfilling theirs—is not so much propositional knowledge of why murder is wrong, of why murder is wrong,—as the kind of knowledge that comes from imaginative understanding wisdom about of-why murder is wrong.

“All this about *what it's like* to murder. But that isn't the philosopher's business; the philosopher's business is *how to define* murder, which acts to *classify* as murder. And talking about the phenomenal content of actual or imaginary experiences of murder doesn't help at all with that.”

The philosopher's business is not one of these or the other. It is both. So I didn't need to have anything to say about which acts to classify as murder (or whatever) to have something philosophically worthwhile to say. But in fact I do have a thesis about which acts to classify as murder (or whatever). Implicitly I have already stated it, by talking about paradigmatic phenomenal contents. What we may say is that by getting a grip on the phenomenal content of certain paradigm human experiences—having or being a parent or a lover, friendship, old age, bereavement, and indeed murder—we get hold of a kind of first principle in ethical experience which is not a definition, but an icon—a picture. And this picture can serve, better than a definition ever could, as the source—a basis of a basic moral understanding on which we can build and from which we can extend into less basic forms of moral understanding.

Experience, as I said above, teaches us themes; and it teaches us variations on those themes. We learn what murder is from confrontation (we hope, in fiction or imagination only) with its paradigm instances; we learn what cases to count, or not to count, as murder by thinking about how they compare with those paradigms. And here too art and the imaginative experience of art are a rich ethical resource; for one of the key things that happens in art is, of course, complication and moral ambiguity, the exploration of life's most basic themes through variations, inversions, convolutions, juxtapositions.

6. *An ethics for zombies?*

It is interesting to compare the dialectic of my argument about phenomenal contents in ethics with the well-known debate between physicalists and their opponents about zombies.

The physicalists argue that a physicalistic world is one in which, as a matter of necessity, the mental facts supervene on the physical facts. Supervenience is supposed to be a fixing relation; the theoretical point of making the mental supervene on the physical is to give a formal characterisation of the idea that the mental is completely determined by the physical. And the anti-physicalists contend that the supposition that the physical facts are as they are with us is equally consistent either with the mental facts that are true of us, with all our technicolour phenomenal life, or with the mental facts that are true of the zombie, the person with no phenomenal experience whatsoever. Here then is something mental that is *not* fixed by any physical detail. And this (says the anti-physicalist) shows the falsity and inadequacy of the physicalist thesis that the way things are at every other level, including the mental, is fully determined by how things are at the physical level.

A physicalist's best strategy to meet the zombie challenge seems to be to improve his theory in order to show how the physical facts do fix the mental facts after all. An adequate theory of the physical will then be one on which it is not only true that there cannot be zombies, but also true that we can see *why* there cannot be zombies. What such an upgraded physicalism might look like, whether it is even possible, are hotly disputed questions in contemporary philosophy of mind. What is widely agreed is that the zombies hypothesis sets a test that an adequate physicalism needs to pass; there can't be an adequate physicalism that doesn't in some way take account of our phenomenal contents.

On the other side of the debate, the anti-physicalist's most usual motivation in setting the zombie challenge is presumably not to concede that the mental is, to use Jackson's word, epiphenomenal. It's not normally going to be his point—far from it—that we can imagine the very same person *with or without* conscious experience, so that conscious experience becomes a causally inert nomological floater. The anti-physicalist's point is rather that phenomenal content *would be* subtractable without explanatory loss *if* supervenience were true. That, says the anti-physicalist, is a reason for thinking that supervenience isn't true, precisely because phenomenal content is *not* subtractable without explanatory loss in any adequate account of the place of the mental in the world. To put it another way, there can't be an adequate *anti*-physicalism that doesn't take account of our phenomenal contents, either.

Is there an analogous dialectic in ethics? I think there is. (So far as I know, I am the first to spell it out.) It goes like this.

The present paper insists on the crucial place in ethics of the paradigm phenomenal contents that I have described. One possible response is simply bluff denial. Someone might say that we can determine a complete theory of the good and the bad, the right

and the wrong, without the slightest attention to all this flaky stuff about what ethical experience is like. Such a theory would be an ethics for zombies, in the sense that it would make no difference whether or not those who lived by it had any conscious experience of the ethical kinds I have been describing (or indeed, quite possibly, of any other kind). The determinant of the rightness or wrongness of everything the agents of such a theory did would always be external, behavioural factors: perhaps certain specified kinds of consequences of their actions, or the fit of those actions with some principle. On such a view rightness could never be determined by internal factors of the kinds that arise from what a Humean might call “the force and vivacity of our impressions”, or from what I would the paradigm experiences of life. Perhaps that means—though it depends how we spell things out—that an external view of this sort could never say that what made an action right was not just that it fitted a principle, but that it was done *because* it fitted that principle. And surely it does mean—however we spell things out—that such an externalised normative ethics cannot say that sometimes what makes an act right is that it is not only the act that the *phronimos* does or would do, but also that it is done *in the way that the phronimos would do it.*

If we think (as I do) that any plausible normative ethics has got to be able to say this sort of thing, that suggests that there is an adequacy test for accounts of normative ethics that we can call the zombie test. Against any purely or close-to-purely external account of normative ethics we can object that nothing in the account would be any different if moral agents had no phenomenal experience at all. To be adequate, an externally-based ethics will have to explain, not only how to fix the external criteria of good action, but also how fixing those gets us the internal criteria too. So far—the objection will run—no externally based ethics has succeeded in doing that. (What consequentialists like J.J.C.Smart say about e.g. blame and atrocity seems a clear example of this sort of failure. The reason to blame, for them, is the purely external matter of the *effects* of blaming; the reason not to commit atrocities is the external matter of people’s *not liking it* when atrocities are committed. What we cannot do, if we are Smart-style consequentialists, is get inside the thinking involved in our natural practices of blaming and reacting strongly to atrocities.) We therefore need either a normative ethics which shows *how* the external fixes the internal; or, failing that, a normative ethics which gives up altogether on the ambition to fix the internal by fixing the external first, and instead, goes internal.

Whatever I may think about the analogous case of physicalism vs. anti-physicalism, I know what I want to say about “ethics from the outside in” vs. “ethics from the inside out”. I think we should go for the latter, for the position analogous to anti-physicalism. For it looks to me like the phenomenal contents that go with the paradigm ethical events are not *epiphenomenal* to our ethical lives, to our motivation, vision, and dispositions as moral agents. They are the heart of the matter if anything is. There could not be an “ethical zombie”, someone whose moral motivation, vision, and disposition was just like ours, except that that person had no ethical experience. If you take the ethical experience out of the picture, you take *everything* out of the picture. Without ethical experience, we are nothing. Literally, for if we are not ethical experiencers, there is no “us”.

Very briefly, ethics is hard because ethics is centrally about phenomenal contents, and it’s hard to talk well about phenomenal contents. Even more briefly: ethics is hard for

lots of reasons. One of them, the one I have talked about here, is because ethics is centrally about phenomenal contents, experience, and it's hard to talk about experience. Or most briefly of all: ethics is not a science, but a humanity. Which is not to say that ethics does not involve knowledge; the point here has been to develop a more detailed picture of what kinds of knowledge ethics might involve, and to move away from the usual idea that ethics is focally about propositional knowledge. If the present paper is right, it is about experiential knowledge too.⁵

⁵ Thanks for their comments to Peter Goldie, Simon Kirchin, Jimmy Lenman, Joel Rickard, and audiences in Milton Keynes, London, and Canterbury.