

2017

Can Moral Principles Change?

Joseph Raz
Columbia Law School, jr159@columbia.edu

Follow this and additional works at: https://scholarship.law.columbia.edu/faculty_scholarship



Part of the [Law and Philosophy Commons](#), and the [Public Law and Legal Theory Commons](#)

Recommended Citation

Joseph Raz, *Can Moral Principles Change?*, KING'S COLLEGE LONDON DICKSON POON SCHOOL OF LAW LEGAL STUDIES RESEARCH PAPER NO. 2017-40; OXFORD LEGAL STUDIES RESEARCH PAPER NO. 58/2017 (2017).
Available at: https://scholarship.law.columbia.edu/faculty_scholarship/2052

This Working Paper is brought to you for free and open access by the Faculty Publications at Scholarship Archive. It has been accepted for inclusion in Faculty Scholarship by an authorized administrator of Scholarship Archive. For more information, please contact scholarshiparchive@law.columbia.edu.

The Dickson Poon School of Law
SomersetHouseEastWing
Strand, London, WC2R 2LS



Can Moral Principles Change?

Joseph Raz

King's College London Dickson Poon School of Law
Legal Studies Research Paper Series: Paper No. 2017-40

Can Moral Principles Change?

Joseph Raz

The Colloquium

I didn't check dates, but in my memory, I started frequenting the Colloquium roughly when I started visiting New York regularly – in other words, a long time ago. So long ago that for a short time I saw others join Tom and Ronnie in conducting the colloquium, some officially, and some by the sheer weight of their presence. I am thinking especially of Donald Davidson, a regular attendee during his visit at NYU, and later of Barry Stroud, who had a similar position. The Colloquium changed in many ways over the years. After it was moved to its current location it lost the excitement of overcrowded class rooms and participants chasing chairs all over the building to bring into them, and after no more chairs could be squeezed in, just piling up, on top of one another on window sills, and in the open doorway to the room, spilling into the corridor outside. Like all human institutions, perhaps especially after they become hallowed, the Colloquium was not everyone's cup of tea. But apart from delighting and enlightening generations of students, and passing visitors, it also accumulated a bunch of loyal followers for whom it has become an important part of their philosophical diet.

What made the Colloquium special, what gave it the character it preserves through the various changes, was the incisiveness and insight of the criticism by Tom and Ronnie in their very different ways. True, the papers presented had to be of some quality to benefit, and to benefit the audience, by being criticized in these ways. Poor papers fell flat and left Tom and Ronnie with nothing of interest to say. But what I and others learnt most was not what the papers themselves brought to it, but how their strengths and weaknesses were illuminated by the two custodians of the Colloquium.

In offering the reflections that follow I availed myself of the liberty given by the organizers to deviate from the range of topics preferred for the Colloquium itself. But I trust that my remarks remain firmly within the concerns central to the thought of both Dworkin and Nagel, sharing their belief in the irreducibility of the normative, and struggling, as they did, to strive to understand how the normative aspects of reality, which are not secondary or derivative, relate to the non-normative aspects of reality.

Introduction

Among the many, often conflicting, images and ideas about morality that nourish our thinking, permeating it even when they do not lead to specific conclusions, are thoughts about morality as dependent on human nature, and on the social conditions of human existence, and as meant for humans, having no existence without them, and as created, or invented, by humans, individually or collectively. These and other thoughts tend to make us think of morality as changeable and changing.¹ But no less pervasive are thoughts that morality is independent of us, facing us with demands whether we want them or not, thoughts that we are subject to morality willy nilly, that sometimes, perhaps often, we are called upon to sacrifice the most cherished and desired things in our life in a moral cause. These and related ideas and images tend to make us see morality as unchanging, existing independently of us, so that we have to struggle to find out what it demands of us, and to struggle, sometimes against our nature, to comply with its demands, or to live up to its standards.

Ideas belonging to these two poles, or many versions of them, share important features. They all allow for a cognitivist view of morality, namely that there are morally correct and morally incorrect views. Even when they appear to conflict it is not obvious that they do, and theoretical attempts to reconcile them abound, the best known being Kant's idea of the moral law being both made by the self-legislation of rational agents, and binding on them in virtue of their rational nature, which leaves no room for choice as to which moral laws to make. This paper is another attempt to explore the two apparently conflicting strands in reflections about morality, prevalent in our culture, except that I will generalize and consider normative principles generally, taking moral principles to be among them. I will also simplify and discuss most of the time principles of duty only. The question will be can principles of duty change? The focus on them is not because they are more important or significant than other principles. I do not believe that they are. It is merely a simplifying device. The assumption is that if they can change so can other normative principles, and that by exploring the possibility of change and

¹ My interest is in the possibility of basic normative change, namely one which is not merely due to a change in non-normative facts yielding a change in the application of (independently existing) normative considerations. Much of the time I discuss the possibility of normative change generally, as a way of bringing out the problems involved with the possibility of basic change.

contingency in normative principles we will improve our understanding of the nature of normativity.

The difficulties in defending the possibility of change are considerable. That possibility runs contra to a common understanding of the nature of reality – roughly speaking affirmed by Hume, in the wake of the rise of modern physics, and the Gallilean revolution. That understanding has been challenged by Kant and others, the most interesting under Kantian influence. Needless to say, this is not the occasion for an extensive discussion of these issues. But their presence in the background dictates the structure of the paper. Its first part argues for the possibility of change by removing an objection to it, an objection that appears to be a compelling consequence of the nature of normativity. The second part of the paper surveys and comments on difficulties of accommodating the possibility of change of normative principles with various common views about the nature of reality.

Part One: Removing an Objection to the Possibility of Change

1. The Arbitrariness difficulty, and its background

A person's moral duties, rights and other normative conditions may change from time to time. Often, we invoke a general proposition in order to explain the changes. It says that one has that duty only under certain conditions and they existed at one point in time, but no longer exist at a later point, and so on. But is it possible for the explaining propositions to be true under some conditions and not others? Most societies hold that one has a moral duty to bury one's dead kin. Let us assume that they are right and we have such a duty. Sophocles' Antigone died for keeping that duty. That proposition: one has a duty to bury one's dead kin, both explains and justifies her actions. But what if one's kin died of Ebola? During the recent outbreak, many people knowingly put their lives at risk, and fought nurses and police, to bury their loved ones. Did they have this duty in these, changed, circumstances? Possibly not; in which case the explanation invoked over many generations to justify the duty did not apply to them, and perhaps does not apply anymore at all. How are we to understand that change? One common reaction is to say that the principle invoked was incompletely stated, the case of deadly infection from contact with the dead was omitted because it is rare, and there is no need to mention it in most cases. Another reaction is to suggest that the duty to bury may be overridden by conflicting principles, as is the case in the Ebola example. In other words, one

way or another these reactions suggest that there is always an unchanging explanation, and appearances to the contrary are due to incomplete statements of the explanation. People's normative conditions change, but the moral principles that explain their conditions cannot change.

Notice that confidence in that conclusion is often great, and is independent of actually having the complete explanation needed to establish the case. We are not deterred by ignorance of a complete explanation that is at least potentially unchanging. We feel confident that there is one even though we do not know what it is. Therefore, the way that we treat examples does not establish that moral principles cannot change. Rather it relies on that view. But why accept it in the first place?

I will continue to discuss the conclusion informally, but some clarification is required to avoid confusion: when referring to normative propositions I ignore propositions about people's actual or hypothetical or possible beliefs about the normative situation of this or that. I will also ignore propositions about whether or not, or to what degree people conform to, or enjoy etc. their normative conditions. Rather, the propositions to be discussed here are general propositions that express the application, actual or conditional, of some normative property, such as having a duty, or its implications. Typically, these are propositions saying that x has a duty, or a right etc. or has had it or will have it when C obtains, or they are propositions that entail some conditions for the application of a normative property. The conditions, C, can be non-normative or normative or a mix. Normally we say that such propositions apply when the conditions are met. My interest is in whether propositions that express principles are true, always or a-temporally, or not. We do not normally refer to all general propositions of this character as principles. I will use 'normative principles' informally to refer to true general normative propositions that have the power to explain or that non-trivially contribute to the explanation of the phenomena they are about. So, for example, that people should honour and care for their parents, truly states both how people should behave towards some people and that that those people are one's parents is part of the explanation why one should so behave. Or, schematically, that all humans have a duty to A is a principle if they have that duty because they are human. By way of contrast, that all males have a duty to A is not a principle, if they have the duty because they are human. The distinction is vague, but neither for our purpose

nor for any other I know of, does the precise delineation of principles matter. When in doubt count it as a principle.

Note that principles are not rules or norms or standards, etc. These are made by people, individually or collectively. They may affect the duties, rights, privileges etc. that we have, and there are principles that state when they do. Binding rules, like binding promises, come and go, and they give rise to rights and duties. There are principles that determine when a rule or standard or promise etc. is binding.

The No change thesis says that normative principles cannot change. The challenge of arbitrariness may be the most powerful argument for the no change thesis. Consider the possibility that if people or their circumstances change, governments have duties that they do not have absent that change. Assume that governments do not have the right, let alone a duty, to place people in administrative detention; and that someone claims

C: if humans become subject to condition X governments have a duty to place those who have the condition in administrative detention.

We may well reserve judgement, because, we would say, it all depends on what that condition is. Compare the following

- (1) If humans become subject to deadly diseases that manifest themselves only after long periods of undetectable incubation governments have a duty to place those who had contact with people who have the disease in administrative detention.
- (2) If all people start losing the sight of their left eye during their 50th year governments have a duty to detain people on the day before their 50th birthday.

We are more likely to believe that (1) is correct than that (2) is. Why? Because we can imagine a normative principle that validates (1), e.g.

(3) Governments have a duty to prevent the spread of deadly diseases, but we find it hard to think of a normative principle that would establish that (2) is true. In the absence of a normative principle to explain and justify why the change makes a normative difference we reject the thought that it does. It renders the putative normative change resulting from the non-normative one arbitrary. And we do not believe that normative principles can be arbitrary.

But does the absence of a normative principle explaining it render the normative change stated in (2) arbitrary? Saying that it is arbitrary assumes that there ought to be an explanation of a certain kind and it does not exist. Without that assumption, the change is merely unexplained, not arbitrary. Why assume that there ought to be an explanation? And why assume, as I do, that there ought to be a normative explanation (i.e. an explanation where a semantically irreducible normative proposition is among the explanans)?

That second point is crucial, for obviously there can be other explanations. That the loss of sight in the left eye is due to a new mutation that has spread quickly and affects all humans may explain it, but it is not the kind of explanation that can explain the duty it is supposed to have generated. That if the government interns people before their 50th birthday there will be fewer people with limited eyesight at large is a consequence of observing the duty that (2) attributes to governments, but it can hardly be said to explain that duty. By way of contrast, that if governments observe the duty claimed in (1) the spread of deadly diseases will be halted, is a consequence of observing that duty, and it does seem to explain why governments have the duty. It explains that by connecting the duty asserted in (1) with the one asserted in (3), subsuming (1) under (3), and thus explaining (1) in the required way, it establishes that (1) provides a way of fulfilling the duty stated in (3).²

The example generalises, and suggests that the availability of normative explanations is in the nature of the normative. Furthermore, their availability is a feature of the normative for reasons that do not concern “our” problem of normative change. It is possible to explain the existence of any normative condition because it is in the nature of the normative that (a) it is intelligible to those to whom it applies (if it is a duty, it is intelligible to those subject to it, etc.), and (b) that intelligibility is communicable, in that normative conditions can be explained. Why?

2. The Normative Intelligibility thesis

The first condition – Intelligibility – first, beginning with what it is not: It should not be confused with intuitionism, at least not if intuiting that p is a transparent mental state (i.e. one of whose presence one cannot fail to be aware) whose presence makes it certain that p (or some variant of this condition). I will suggest that there is “a feel” as part of finding something intelligible, but

² I take a principle B to be subsumed under another principle A if either A entails B or following B makes it likely that A will be followed or that its goal will be realized.

the same experience is part of mistakenly thinking that one finds it intelligible. Nor is the intelligibility of duty-related principles, or the intelligibility of having the duties that they relate to, a feature of or a condition of believing that the principles are true or correct, or believing in the existence of the duties. One may believe that one has a certain duty because one was so advised by a reliable expert, or because that is the common view and, in the circumstances, it is unlikely that the common view is mistaken, and yet find it unintelligible that one has that duty.

Perhaps we must, if we believe that we have a duty, believe that it can be intelligible to us even though it is not currently intelligible to us?³ Not quite: there can be, and there are, views that take at least some domains of normative considerations, often a religious domain, to be beyond human understanding. We can know our duties because we were commanded what to do, but it is not for humans to comprehend the point of those commands, etc. Of course, people who have such beliefs are mistaken and by implication they have mistaken views of the nature of normativity. So, perhaps we could say of those who have a correct view of the nature of normativity that if they believe that they have a duty they believe that it can be intelligible to them, even if it is not currently intelligible. This may be right, but does it apply to (presumably) the vast majority of people who have no explicit view about the nature of normativity?

What the intelligibility thesis says is that the existence of a duty can be intelligible to those subject to it. We may assume that generally (but not that universally) those who entertain normative thoughts implicitly believe that. That is why we assumed that they are more likely to believe the duty stated in (1) than in (2) above – it is more likely to be intelligible, and they believe, at least implicitly, that normative conditions are intelligible.

Does it mean that false beliefs about duties are not intelligible to those who would have the duty if the beliefs were true? The thesis is about the intelligibility of the fact that we are subject to duties, and to other normative conditions. If it is right that generally people would not believe in a duty-related principle whose duties are not taken to be intelligible to those subject to them, when they do believe in such principles they implicitly believe that the duties, if indeed they exist, are intelligible to their subjects. Those believing in the principles may believe (rightly or wrongly) that they themselves are subject to them. In that case they generally also implicitly

³ Remember that the intelligibility thesis claims only that duties, and other normative conditions, are intelligible to those subject to them.

believe that the duties, if they apply to them, are intelligible to them (meaning, as we saw, that they can be intelligible to them). As they also believe that the duties apply to them they believe that they are intelligible. If the duties do not exist that belief is mistaken, though their implied belief that if they exist they are intelligible is correct.

But what is intelligibility and why assume that the normative domain is intelligible? First, I have to disambiguate two ways in which I have been using the term without drawing attention to the ambiguity in its use. In some contexts 'intelligibility' refers to (1) a property of a duty or principle. In others, it is (2) a relation between the person subject to the duty or principle and the duty to which he is subject. A duty is intelligible (1) when it can be intelligible (2) to the people to whom it applies. Hence, my duty may be intelligible to me, but I may find it unintelligible while knowing that it is intelligible, meaning that it may become intelligible to me. This is a natural way of referring to matters, though not when they are referred to in one and the same statement. Put it in other words: I may understand why I am subject to my duty, but I may also not understand that, while knowing that I can come to understand it (is intelligibility understanding? See below). So, a duty is intelligible to a person when it is intelligible to that person, or when it can become intelligible to him or her.

The normative domain is one in which we react, actively, using our rational powers, to situations in light of a (correct or mistaken) understanding of the proper way to react to them in thought, action or emotion. These reactions, let me underline, are not knee jerks, nor cases of being propelled (by desire, fear, anger etc.), having lost one's self-control. They are our active responses, ones that we take to be appropriate, and that means that we see their point or believe that they have a point. We react as we do because we think that that is the way to react – we understand (or think we do, for we may be prey to self-deception, rationalisations, etc.) that these are the ways to react. That means that normative considerations (the considerations we thus react to) are intelligible to us. We see their point, or believe that they have a point, and except in cases in which we succumb to self-hate, self-debasement, and other such psychological motives, we approve of how we react, we stand by it, as it were, or when convinced that we were mistaken, we regret our reactions.

Talk of intelligibility and of seeing the point of a duty may sound mysterious. Does it mean more than knowing or believing that one has a duty? I think that it does. Intelligibility includes a measure of understanding, which consists to a large degree in ability to place the knowledge or belief in context, seeing its place in relation to much else, and an ability to extrapolate, to derive more beliefs from it (all being a matter of degree). But intelligibility also comes with a certain feeling that accompanies or colours the belief or knowledge that there is a duty (or right or normative reason, etc.). It is indicated by the common choices of terminology: a feeling that acting as one has a duty to do is a fitting act in the circumstances, a fitting or appropriate response to the situation, and other similar expressions. We can understand something without realising that we do, while continuing to feel puzzled by it. This may be due, e.g., to a misunderstanding of what understanding consists in, wishing for some miraculous answer which does not make sense. When a principle or the having of a duty is intelligible to us our understanding is coloured by that feeling that things are fitting, or something like that. When we explain why we have a duty it is this feeling that we explain by justifying it. We explain why we are justified to feel that acting as we have a duty to is fitting, and or appropriate, and why our belief has that colour. But the feeling is there independently of our ability to explain it. Perhaps it is easiest to perceive it when our conviction is strong even though we find it difficult to explain. Many people firmly believe that incest is wrong, but are less sure about why. The feeling of fittingness or appropriateness accompanying the belief is there to see, often rising to indignation when their belief is challenged, or to discomfort when they sense a doubt in it. That kind of colour or feeling accompanies any securely held normative beliefs (whether right or mistaken), though it is possible to know that one has a duty without that feeling, without seeing its point.⁴

3. The Explainability thesis

This does not entail that we can communicate, that is explain, why our reactions make sense. We may see the point of having a duty without being able to explain it. For example, some versions of intuitionism have it that while the normative is intelligible it cannot always be explained. It can only be intuited. We “see” why there is a duty in such circumstances, and the intuition can be shared, but those who do not have it cannot be rationally convinced that this is

⁴ I am not suggesting that it is exclusive to normative beliefs.

so by providing them with an explanation of why the duty is there. If an explanation does not compel (on pain of irrationality) acceptance, it is no explanation – they contend.

I will not try to argue that the normative can always be explained. But I want to offer some observations that may make it more plausible that it can be, by pointing out some of the things that even good explanations cannot, or at any rate need not, accomplish.

For example, the thesis that the normative can be explained is not at odds with intuitionism. Explanations convince people qua explanations only if they understand them and realise that they do. That understanding (and the point is general, not being confined to explanations of the normative) presupposes both that the people to whom the explanation is addressed have certain capacities, rational, perceptual, emotional, and that they have certain experiences. Little can be explained to the very young. And that is not because of the limits on what can be explained, but the limits of their abilities and experiences required to understand the explanation. Given the required capacities and experiences, and the strength of mind to question beliefs that have been instilled in one by those one has been brought up to respect, a good explanation, once it is understood, convinces. It convinces because one sees the point. It makes one see why there is a duty in such circumstances, and so on. In other words, explanations, to work, require the capacity to see how they explain, to realise that they explain, and that may be the power to intuit, or at least a power to intuit.

Second, the availability of explanations does not guarantee agreement, not even if everyone is given an explanation. Explanations can be based on false assumptions, and there can be disputes regarding which explanations are good ones. In any case, many mistaken beliefs about duties etc. may be, up to a point, intelligible. We cannot expect that if there are explanations they will, if offered, eliminate disagreement, not even in the long run.

Perhaps this indicates that the practical importance of explanations is not as great as is sometimes supposed. All the same the theoretical question whether there are always explanations for the existence of normative conditions, such as that someone has some duty, is significant.

Perhaps we can make progress by comparing the explainability thesis with a view fast becoming very popular, according to which there is a special kind of relationship, the grounding relationship, such that every true normative proposition is true in virtue of there being a

ground for the condition it expresses or describes. Call this the **Grounding Thesis**. It may be thought to vindicate the explainability of the normative: these explanations, it may be said, consist simply in pointing to the ground of the condition explained. If one has a duty then there is a ground in virtue of which one has that duty, and that ground explains the duty.

I remain neutral about the grounding thesis. Even if true it does not establish that explainability is true. In part this is because there are likely to be true normative propositions whose truth does not depend on a ground. For example, I doubt that there are grounds that can explain why propositions of the following forms are true:

- (1) Necessarily, if x is of value then it is possible for there to be some circumstances in which x 's value is a reason to protect x 's existence or to bring it about that x .
- (2) Necessarily, if there is a reason to φ then it is possible⁵ to φ for that reason.

More to our point, however, is the fact that the grounding thesis does not establish that the grounds are either knowable or explainable. Moreover, one may know that G is the ground of some duty D , without understanding why D , and without understanding why or how G is its ground. Example: the ground of my duty to φ may be my having promised to φ . I may know that, and still not understand why when I promise I have a duty. Given that knowing the ground for a duty does not entail either having an explanation of why the duty exists, nor that one understands why it exists, the grounding thesis does not guarantee explainability without the support of an additional thesis about the explainability of grounds.

The two theses have different intellectual sources, though the theses are compatible and may turn out to converge. The grounding thesis as commonly presented is based on the thought that whatever is normatively binding must be so in virtue of something. That *something* is, or is related to, the essence of the normative condition. This is an objective pull, drawing on the constraints or the appeal of metaphysical theorising. It contrasts with the subjective pull drawing us to the explainability thesis, subjective because it draws on the way normative considerations function in the life of those capable of appreciating them. This is the kind of consideration that argued for the intelligibility thesis: The normative is the domain of considerations by which we guide ourselves, intentionally, because we see the point of those

⁵ Meaning: It is possible for there to be circumstances such that when they obtain it is possible to φ for that reason.

considerations, because we understand, or think we do or can do, their point. The case for the explainability thesis derives from the same source.

The explanations the thesis is about are, needless to say, good explanations. Explanations are good if the person to whom they are addressed can come to understand what they explain by considering the explanation and being minded to follow it if it is sound. In other words, explanations are good explanations if they succeed with a qualified and willing addressee.

Various arguments converge to show that what we can understand can be understood by others, provided they have the required capacities and experiences. Given that explanation consists in bringing the other to understand, it is plausible to think that often Mary can explain to John why there is a duty (or whatever) by pointing out to him the features that made her understand why there is a duty (or whatever). Sometimes the differences between John and Mary will undermine that possibility. But then Mary could adjust her explanation to those differences, and point to features that would have made her understand had she been in John's position. She can explain this to John, provided he has the required capacities and experiences. The pointing that explanations consist in can be purely verbal. But it need not be verbal or exclusively verbal.

These reflections do not, of course, establish that everything normative can be explained. They even suggest that not everything can be explained to everyone, because people's capacities and experiences, including their potential to expand, may be limited. But given the intelligibility thesis, and the connection between intelligibility and explainability, it is plausible to think that an appropriately qualified explainability thesis is true.

This is not the occasion to try a relatively precise formulation of the thesis. But one aspect of it is important to our purpose. Our capacity to understand and communicate is enhanced and limited by the concepts, both normative and non-normative, we master. There is no necessary sequence by which we acquire them and other concepts that we use to explain and to understand them. I mean that we acquire both normative and non-normative concepts at an early stage of growing up, and we explain concepts of one category with the help of concepts of the other. There is no priority, let alone independence, of one of those categories relative to the other.

4. Does Explainability require subsumption?

The arbitrariness argument against change is:

- (1) The coming into being of a new normative principle constitutes an arbitrary change unless it can be explained.
- (2) An explanation must include reference to a principle that is already in force, which, together with additional factors, explains the new principle.
- (3) Such an explanation shows the new principle to be an application of the existing principle to some of the circumstances that fall under it. It explains the new principle by subsuming it.

Subsumptive Explanation is a capacious form of explanation. It includes, of course, explanation by generalisation as, for example, explaining why Americans have a duty not to kill by invoking the principle that all humans have that duty and Americans are humans. But it also includes constitutive explanations, namely explaining why A is valuable by the fact that A, which is a constituent part of B, contributes to the value of B, and since B is valuable, all its constituents part that contribute to its value are valuable.

Accepting the first premise I will challenge the others. But before I do that two important concessions:

The first is that there are forms of intuitionism that escape this argument, by allowing that the new principle is intelligible but cannot be explained. Its validity is simply intuited. I will not consider that possibility.

Secondly, given that by their nature normative considerations are capable of guiding and evaluating conduct, character, etc., and that therefore they are accessible to the people whose conduct etc. they can guide and evaluate, it is reasonable to assume that no normative change can retrospectively reverse the standing of a person's actions, emotions, character etc. For example, no change of normative principles happening today can make something I did yesterday wrong, if it was not wrong when I did it. In retrospect, it can turn out to have been an unfortunate thing to do, etc. But it cannot become wrong or dutiful or virtuous if it was not so at the time. This is consistent with changes to our understanding of our moral conduct, and indeed of moral and other normative principles, changes in their understanding that occur due

to the emergence of new normative principles. The new principles may provide new explanations of the older principles, relocating them in a deeper context.⁶

Back to the argument against basic normative change: Its third premise was that normative explanations are subsumptive explanations. But not all of them are. I will mention only one other kind of explanation: explanation based on argument by analogy. I treat it as a distinctive kind of argument. Possibly it should be treated as a family of related arguments; that makes no difference to our purpose. I take analogical arguments to be defeasible (non-monotonic) arguments consisting of four types of premises or groups of premises.

- (a) *The Target*: Some *a* has a feature *F*.
- (b) *The Similarity*: both *a* and something else, *b*, have a feature *G*.
- (c) *The Relevance*: *a* and *b* having *G* is relevant to *b*'s having *F*: Roughly, given the background, that they both have *G* makes it likely that if *a* has *F* so does *b*.
- (d) *The Closure*: There is no reason to think that the relevance is undermined in the instant case.

The conclusion is then drawn that *b* has *F*.

There are also negative analogies: given that *a* is *F* and that while *a* is *G*, *b* is not-*G*, *b* is not-*F* either. Their additional premises are inevitably different, but these differences do not affect our concerns and I will not discuss them further.

The key to the success of an argument by analogy is in the relevance premise: does it support the conclusion, and does it do so to the degree required given the nature of the conclusion, and the circumstances in which we draw it?⁷ Naturally, there have been various attempts to formalise that premise, or at least to make it much more detailed⁸, and I have no doubt that various kinds of analogical arguments can be made precise, and some can be formalised to make them suitable for use by deterministic robots. But I doubt that analogical reasoning itself is amenable to these kinds of explanations.

⁶ See my "Moral Change and Social Relativism", CULTURAL PLURALISM AND MORAL KNOWLEDGE, eds. E.F. Paul, F.D. Miller, and J Paul (Cambridge: C.U.P., 1994).

⁷ See my paper 'Normativity: The Role of Reasoning' (PHILOSOPHICAL ISSUES, 25 Normativity 2015, ed. Ram Neta) for an elaboration.

⁸ For a helpful survey and bibliography see Bartha, Paul, "Analogy and Analogical Reasoning", The Stanford Encyclopedia of Philosophy (Fall 2013 Edition), Edward N. Zalta (ed.), url = <http://plato.stanford.edu/archives/fall2013/entries/reasoning-analogy/>.

We are familiar with the fact that people's knowledge, including their "knowledge how", extends beyond their ability to articulate what they know, let alone to explain it. This ability to exceed one's power to articulate enables people to reason on the basis of deductive arguments. Many who rely in their reasoning on deductive arguments do not have explicit knowledge of what deduction is, nor what rules of inference they rely on. Probably the vast majority of people do not have explicit knowledge of all the rules of inference they rely on, not even all the rules of deductive inference in virtue of which their reasoning is valid. The ability to reason well does not depend on an ability to articulate, let alone to defend, the (sound) rules of inference one relies on. It is sufficient that one relies on them, and that it was not by chance that one did. Needless to say, the same is true of analogical arguments. Analogical reasoning may be special in that its rules of inference cannot be completely articulated and formalised. That is, it seems to me plausible that our knowledge extends beyond our ability to articulate it and that it always will. It is a necessary feature of knowledge that those who possess it cannot articulate everything that they know. And it may be a necessary feature of knowledge that regarding the knowledge possessed by any living organism, it is impossible to articulate it exhaustively.⁹

If that is so then it is plausible to assume that there is a form of argument, or a family of arguments, regarding which our mastery of them exceeds our ability to articulate in detail, let alone with robotic detail, their rules of inference. It is plausible that analogical arguments are of that kind, and therefore have special importance as a way of vindicating and explaining knowledge claims, including explaining normative knowledge claims. They are arguments our mastery of which exceeds the possibility of exhaustively articulating their rules. For it appears that, while implicit knowledge, as we can call knowledge that we have but cannot articulate, may be relied on whenever we employ other forms of arguments too, analogical arguments are particularly suited to reasoning that relies on implicit knowledge, that is reasoning when our limited ability to articulate what we know does not enable us to describe in detail the inferences we rely on. Even though various special categories of analogical argument may yield their secrets and be formalised, there are others that will not. Our ability to employ analogical arguments will remain dependent on our implicit knowledge.

⁹ None of this implies that there is anything we know but cannot articulate. The claim is merely that we cannot articulate everything we know.

It also appears that we rely heavily on analogical arguments in political and ethical thought and discussion, and that makes them relevant to our concern here today. For analogical arguments do not “work” by subsumption. The explanations they provide do not presuppose a more general existing principle. They explain not vertically but horizontally, metaphorically speaking. They thus remove the arbitrariness argument against the possibility of basic normative change.

5. The Possibility of Change (a first go)

Removing this one objection to basic normative change does not establish that it occurs. Nor is it my aim to show that it does. Rather, mine is the more modest ambition: to illustrate that it can happen. Those who will be convinced of that will not, I predict, doubt that it has happened and will happen again.

The intelligibility and explainability theses point the way. We are comfortable with the thought that there are limits to our normative understanding and that changes in the world can change them. But, can changes in the world not only improve our normative understanding, but change the normative principles themselves? They can, seems to be the answer.

Suppose that research into dark matter reveals to us the existence of a condition we were unaware of, and indeed were unable to comprehend before the possible existence of dark matter could be thought of, and its mysteries cracked open by science. These scientific developments came along with mastery of new concepts. One of them, let’s designate it ‘akatem’, refers to a condition that is dangerous to humans. So now we know that we should not expose people to akatem. This is new knowledge, but not a new principle. It is a simple application of the principle against endangerment. Morality did not change. Our understanding of what conforming to it involves has improved.

Now imagine that many things change in the dark matter age. Conditions of life, the economy and technology change radically in unimaginable ways. Among other things a new form of human association, which they call ‘demte’, emerges, radically different from any form of friendship or other associations we can imagine today. It defies the distinction between public and private interactions, between a one-on-one relationship and group relations, etc. Demte-related reasons are, naturally, entirely new: when it is desirable to form demtes, when it is better to avoid them, how to conduct oneself toward demte associates and how to conduct oneself towards others, in matters that affect demte associates, and much else.

By hypothesis, demte is (at least potentially) a good form of association among people. And again, by hypothesis, the reasons or duties involved in it are no mere application of very general principles that we know, or can know. Rather, they constitute independent principles. Their emergence constitutes a basic moral change.

The considerations canvassed so far present no objection to the possibility of moral change of this kind. It can meet the conditions of intelligibility and explainability so long as its explanation is not by subsumption. And it can meet the condition that there should not be a retroactive change in the valence of actions and events, because it only applies to events and actions occurring since the arrival of the age of dark matter, for only then did demte associations become possible.

6. Does Explainability require a pre-existing principle?

This concludes my outline of an argument for the possibility of basic moral change, based on rejecting the third premise of the argument from arbitrariness, the premise that took all normative explanations to be subsumptive explanations. As indicated above, there is another way to challenge the argument from arbitrariness: its second premise too is misguided. It stated that a normative explanation of a new principle must include reference to a principle that is already in force, or at any rate, one that is in force before that new principle becomes valid.

To establish that it is mistaken I will explore an objection to the preceding analogy-based argument. The objection fails, but it reveals the mistake underlying the second premise.

The objection is that I was wrong to assume, as the preceding argument does, that analogical arguments are an independent kind of argument. In fact they are valid only in cases where both sides of the analogy are subsumed under a more general principle. Admittedly in one way analogical arguments can be independent of a subsumptive explanation. One can master them and reliably follow them without knowing the more general principle on which their validity depends. This makes them helpful and important. But *au fond*, analogical arguments are valid only if there are valid subsumptive arguments supporting the same conclusion. Therefore, their reach is no greater than that of subsumptive arguments, and therefore they cannot vindicate the possibility of basic normative change if subsumptive arguments cannot do so.

We can examine the objection by applying it to the case of demte. The previous argument assumed that demte reasons are real, binding reasons for they are validated by analogy. Perhaps

that analogy is cogent only if there is also an argument by subsumption validating demte reasons. For example: If demte can be explained by analogy with other phenomena, let's say by partial analogy with C on the one hand complemented by partial analogy with E on the other hand, then there is a higher principle of which C, E, and demte are instances, and that vindicates the analogies. This higher principle explains demte by subsumption. Therefore, either the emergence of demte-related reasons is not a basic normative change, because the more general principle is valid prior to demte's emergence, or demte-related reasons are not valid. Possibly, the more general principle was not valid prior to the emergence of demte-related reasons, therefore we have no sufficient reason to think that they are valid.

Why doubt that the more general principle was valid prior to the emergence of demte-related reasons? On its face it is already valid, for it applies not only to demte-reasons (which have no instances yet) but also to C and to E conditions, and they can already obtain. However, it is possible, and the argument does not deny, that the more general principle can be known only to people who have concepts that do not yet exist, concepts that will emerge in the dark matter age.¹⁰ In that case there is as yet no such valid principle; it will be a valid principle only once the concepts required to state and to understand it have come into being. That much is established by the intelligibility and explainability theses.¹¹ True, we are not subject to demte-reasons, but we are subject to reasons relating to C and to E, and the putative higher principle governs them too. Therefore, if valid now it applies to us now. But if it is not intelligible and explainable to us now then it does not apply to us now. Given that we do not know that it is intelligible to us now we are not warranted to conclude that it is valid now.

So far – the objection, or a version of the objection. What are we to make of it? Not much, I am afraid. It ignores the fact that while we debate and argue about the validity of various analogical arguments, we do not rely on ignorance of a subsumptive explanation to doubt them. Rather we use some analogies to undermine others, expressing our acceptance of them as an

¹⁰ Even if the principle can be stated using concepts already available to us, possibly its validity cannot be explained at the moment for its explanation requires concepts that we do not and cannot have, concepts that will evolve only in the dark matter age.

¹¹ I am aware of course of a convention in some areas of philosophy to think of concepts as a-temporal, i.e. to deny that they have a history. This may be a useful fiction, or if you prefer, a useful stipulative concept, which is a close cousin of concepts, as we know them. There is no need to dispute the legitimacy of either concept of concepts.

independent form of argument. We may of course be mistaken. But the inescapability of implicit knowledge and the intelligibility and explainability theses point to an unavoidable gap between what can be explained by subsumption, which relies on knowledge that can be explicit, and what can be explained by analogical arguments that rely knowledge that is inescapably implicit. The objection has to be rejected. But it draws attention to the possibility that as demte relations emerge so does a principle (which at that point is intelligible and explainable), which explains demte-reasons by subsumption. Suppose that people of the dark matter era do not have an analogical argument to explain demte reasons, but they do have a subsumptive explanation of those reasons, though that explanation depends on a new principle, one that employs new concepts. Can a new principle that does not refer back to prior principles be valid, and thus validate demte reasons? We know that it must be explainable. But possibly its explanation can relate only to other new reasons and principles. It will have to be consistent with all other valid principles, including those previously valid (though it may conflict with them). To that extent its vindication relates to other principles, including those that are already valid. But must its vindication relate to all of them? Some people think so. In recent times, for example, Dworkin made that claim.¹² I know of no reason to believe that that is so. Absent such reasons we may be able to endorse a form of radical value pluralism, namely one that allows for the possibility of groups of valid normative principles that are intelligible and explainable and relatively independent of (some) other groups of principles. If that is right then the second premise may also have to be rejected.

Admittedly, even given radical value pluralism, the likelihood that there will be no explanation of the supposed new principle explaining demte-reasons that refers to existing principles or considerations is small. It is likely that it will enjoy supportive analogies with existing normative conditions. What is unlikely may still be possible.

7. Can the Possibility of Change be explained?

Radical normative change is possible. My illustration of its possibility derives from a specific, though very common and extremely important, social phenomenon that I will call 'normative

¹² R.M. Dworkin, JUSTICE FOR HEDGEHOGS (Harvard U.P., 2011). In "A Hedgehog's Unity of Value" THE LEGACY OF RONALD DWORKIN, eds. W. Waluchaw & S. Sciaraffa (Oxford: O. U.P. 2016) I explained how he failed to make good that claim.

social forms'. Some social actions, relationships or conditions cannot exist without their participants having appropriate concepts to mark them. That is, people can perform those actions (give a gift) or be in those conditions (be married) only if they know what they are doing or what condition they are in. They need not have our concepts, but they must have some appropriate concepts. For example, "gift" may be a very general concept that does not exist and is not available in all cultures. But people may have specific gift-related concepts and they give gifts only when they intend to perform an action that instantiates one of them (e.g. a thank-you gift, a wedding-gift). Typically, these concepts are normative in that by their nature, their instances imply certain reasons towards their occurrence: gifts are appropriate or inappropriate, whereas chairs, watches, earthquakes etc. do not essentially imply such reasons. Moreover, and that explains their normativity, these concepts are used to refer to, and to describe, aspects of cultural values, that is values that are the products of culture, values whose existence is made possible by social practices that include actions done in the knowledge of the normative significance of those cultural practices (though not necessarily as practices). Marriages, friendships, money are examples. Some of these normative institutions (perhaps money) have only instrumental value. Others bring with them new values. There was no marriage before the social institution of marriage emerged. Without it the value of marriage could not be realised. Moreover, it is plausible that one could not know of or understand these new values before they came into being, that is before instances of them became possible. Only once they are in existence could we think of them and realise either that they can be subsumed under a principle that explains them, or that they can be explained through analogies. These social forms are both normative, in that their existence generates new reasons for conduct for those who engage in them (and for everyone, regarding them and the people who engage in them). They are, when valid, intelligible and explainable in the usual way. Normally they will be explained either by subsumption or in some other way, e.g. analogy, involving reference to already valid principles or other social forms. This leads to a puzzle: The case for the possibility of basic normative change depends on the possible emergence of new normative principles that are not instances of existing normative principles. The intelligibility and explainability theses require that these new principles can be explained, and almost invariably they are explained by reference to existing normative

principles. However, if the explanation is by analogy, or in any way other than subsumption, the argument goes, their emergence constitutes a basic normative change.

So here is my puzzle: normally, both if they are instances of existing principles and if they are not, they are explained by **reference to existing** normative principles. Given that fundamental similarity, how could such a momentous difference, between mere application of existing principles and an emergence of basic new principles, rest on the difference between subsumption and, say, analogy?

Is it because there is – in some sense – a greater distance between the explaining principle and the explained one, so that the explained one is a greater departure from the existing normative principles when it is explained by analogy than when it is explained by subsumption? This might have been a good explanation had it been true, but it is not, as is obvious once we consider a case of a basic change, explained by analogy with existing cases, but where there is also a new principle emerging that covers both it and the existing cases to which it is analogous. In such a case the new and old normative cases, and their attached principles, are both instances of a single principle. Hence the distance between them is no greater than can be expected among instances of a single principle. But the new case constitutes a basic normative change because that principle did not exist before the new case came into being. This kind of situation shows that we have no reason to think that explanations by analogy indicate a greater distance than explanations by subsumption.

Perhaps the difference between subsumption and analogy is that once we can understand two principles, A and B, then if B is subsumable under A that will be plain for all to see. It will be, in other words, impossible to understand both A and B without realising, when thinking about both at the same time, that B is subsumable under A. This condition is not true of analogies. It is possible to understand both C and D without realising the relevant similarities between them that establish the validity of an analogical argument from one to the other.

Again, if this were true it would establish an interesting difference between vindication by subsumption and vindication by analogy. But it is not true. It is possible for one principle to be subsumable under another without this being evident to people who understand both (to a degree), and think of both. A topical example is the difficulty that many people have to realize

that same sex marriage is an instance of a marriage relationship (even if it is not, say an instance of Catholic marriage, namely of a marriage recognized in Catholic doctrine).

All of this suggests that the difference between vindication by subsumption and by analogy is not very significant. Can it underpin the fundamental distinction between basic normative change and a novel application of existing principles?

That is my puzzle. I think that it has something important to tell us. It does not suggest that we should abandon or revise any of the arguments or conclusions so far reached. Furthermore, the distinction between application to new circumstances of an existing principle, and the emergence of a new principle, is unobjectionable. The puzzle shows, however, that some mere applications of existing principles resemble, in the relevant respects, the emergence of completely new principles. What the puzzle tells us, in other words, is that the distinction between mere application of the old and the emergence of the new is not all that important; that the newness of some mere applications is as great and as significant as the emergence of a completely new principle. And that is something worth learning, for it redirects our attention from wondering whether basic change is possible to trying to understand the significance of the emergence of principles, whether new or applications of existing ones, which could not have been known about or vindicated independently of the social changes that brought them about.

Part Two: Objections and Implications

The objections to be examined here are not so much objections to the arguments above as direct objections to the very possibility of basic normative change. But they are radically different from the objection of arbitrariness which was discussed in part one. That objection arose out of reflection on the nature of normativity which suggested that intelligibility and explainability seemingly imply that change, were it possible, would be arbitrary. The objections that follow are, if you like, external in origin. They aim to show that normative change is incompatible with fundamental features of reality or with well-established doctrines. I have no doubt that some will take this to make them particularly powerful objections. But the opposite view is also possible: they are weaker, because the argument can be reversed. One can take the possibility of normative change to constitute an objection to the views and doctrines that are incompatible with it. This role reversal was not available, at least not in simple form, in Part

One. The possibility of role reversal sets a limit to the discussion of this part. It cannot be developed into a full-scale examination of the doctrine that lead to the objections. All the paper can do is point to incompatibilities, where there are such, and hint at the implications of the possibility of normative change. It will have to remain a rough and informal discussion.

8. Propositions

To start with an objection to everything I have written here: I took normative principles to be true propositions with explanatory potential of the phenomena they are about. That makes it plausible to think of a radical change as a change in the truth value of a proposition that is a principle if true. But, the objection goes, propositions cannot change their truth value. They are a-temporal, or if one thinks of them as temporal that they exist and possess their truth value always.

It is convenient to take moral principles, and any other principles, if not to be propositions at any rate to be expressible in propositions, so that necessarily any principle is expressed by one proposition (or, to allow forms of vagueness: sets of propositions). The problem can be avoided, at least to some degree, if we understand principles to be true propositions that meet some additional conditions: I introduced one such condition explicitly: principles have explanatory relevance, and another condition implicitly: propositions can be known by those to whom they apply. Let's call putative principles, i.e. ones presented as principles, but which are not necessarily principles, valid principles if they are true, explanatorily relevant and knowable by their subjects. A proposition can be a valid principle at one time and not at another if it has explanatory relevance and is knowable by its subject at one time but not at the other. I assume that it is true at both points in time. Its status as a (valid) principle does vary, though its truth does not.

That may be part of the solution. But the considerations put forward above suggest that we should go further and deny that the truth value of propositions cannot change. True, some think that propositions are a-temporal, or that at any rate it is impossible for them to be true at one time and not true at another. That does not reflect the way propositions are commonly understood. We do understand the proposition that John was the best tennis player 5 years ago but is so no longer, as entailing that (the proposition that) John is the best tennis player is no longer true. It is best to avoid verbal disputes. The concept of a proposition whose truth can

change, meaning that the truth of some types of propositions can change, as well as the concept of a proposition whose truth cannot change have been explored in some detail, and there are accounts of the logical properties of each of them. Both have a reasonable claim to represent what one says when one asserts something, that is to express the content of what can be asserted. After all, that notion, of what one asserts or can assert, is itself subject of the same ambiguities and indeterminacies as the notion of a proposition.

Some would be happy to allow that there is a concept of a proposition that admits that there are propositions whose truth value is neither a-temporal, nor unchangeable, that they may have a truth-value at one point in time and not have it at another, but they claim that everything expressed in such propositions can also be expressed by propositions whose truth value is a-temporal or just necessarily unchangeable. For example, it can be argued that if there is a condition, C, on which the truth of a proposition, P, depends then there is a proposition: P only if C that is a-temporal or necessarily of an unchanging truth value. One way of understanding the thought that moral principles can change involves denying that all propositions with a changeable truth value can be expressed as, or are logically equivalent to, propositions with unchangeable truth value. So, if there is a case for thinking that normative principles can change there is a case for rejecting the thesis that all propositions can be expressed by a-temporal propositions.¹³

The temporal character of a proposition may manifest itself in being capable of being true at some times, and false or neither true nor false at other times. But the temporality of propositions may be due to other facts as well. One would be that the proposition does not exist at all times. If so then possibly its existence is time bound, though its truth need not be.

¹³ We have discussed matters using 'proposition' in the sense that allows for the possibility that some propositions can change their truth value. Normative principles are (among other things) truths about the relations that they express. The principle that Gratitude is a source of duty is a true proposition about the relations between gratitude and duty. That is, it is a principle if and only if and because the relationship holds. This leaves open the question what comes first, the principle or the relationship, or which is more fundamental, at least when that question is understood epistemically, or as a question about the order of understanding. E.g. if the only way to establish that the relationship exists is by *reductio* of the proposition that it does not, then the principle comes epistemically first. Similarly, if the only way to explain, to make sense of the existence of the relationship is through explaining the truth of the proposition then the principle comes first in the order of understanding.

Propositions that exist may be true (in a sense that is not time-relative) but that is relevant or assertible only when they do exist. Alternatively, and less problematically, it is true whenever the proposition exists, and when the proposition does not exist it has only conditional properties, properties it would have were it to exist. The fact that concepts have a history, and that propositions may have concepts among their essential constituents, shows that propositions have a temporal dimension, and that the content of temporal propositions cannot always be expressed by a-temporal propositions.¹⁴

That normative principles can change does not establish that moral principles can. But if the only reason to think that moral principles cannot change would entail that normative principles cannot either, then removing that reason would establish that moral principles can change. That marks one way in which the question about normative principles helps with answering the question about moral principles.

9. Conditionalization

Some may be tempted to argue that we can have our cake and eat it: what attracts us to the thought that the basic moral principles can change is the supposition that (a) the content of moral principles depends on human nature, and (b) human nature is contingent and changeable. Hence, it may appear that moral principles are also changeable. But the conclusion is a non-sequitur. Suppose that one fundamental aspect of human nature is H1 and that the moral principles we have, call them P1, are valid only if H1. Suppose that human nature changes, and that H1 is replaced by H2, and therefore the moral principles that bind humans are now different, call them P2. It does not follow that any proposition changed its truth value, nor that basic moral principles changed. Rather what follows is that the basic moral principles were (always, or a-temporally): If H1 then P1 and if H2 then P2. What changed is that previously when humans had an H1 nature P1 principles applied and now, with their H2 nature, they are subject to P2, but that is not due to a change of any principles, only that principles that were inapplicable before are applicable now and *vice versa*.

¹⁴ If normative principles can change then they are not necessarily true (or, if validity rather than truth makes them principles, not necessarily valid). But they can still be relatively necessary, e.g. relative to the laws of nature.

The difficulties with this approach can be brought to light by supposing that it is possible for humans to develop a sense perception they currently do not have, and indeed one that no animal currently has. We do not know what it will be, if it will be, and we cannot know what it will be, as it will be a new kind of sense. We lack a concept by which to refer to that sense, and can refer to it only by expressions such as 'the hitherto unknown sense', which tell us nothing about it other than that it is a perceptual sense. It is possible, though we do not know that for sure, that its emergence will make us subject to moral principles that we are not subject to at the moment. Needless to say, we have no clue, and can have no clue, about their content.

The question is whether given this situation there is something we are ignorant about, something we do not know. It is not something that can be known, but so are many truths. It may be impossible now to know whether there was a speck of dust on top of my desk lamp five minutes ago. Principles, however, are not like specks of dust. They are essentially knowable. They are guides to thought and action, and while they may be unknown, which leads us to look for them, it seems odd to think that they can be unknowable. In what way are they guides for thoughts and action if unknowable?

Perhaps, however, principles can be temporarily unknowable. Or, they can be unknowable under some conditions provided that they are knowable under others. If we turn our attention from the future to the past, we can entertain the thought that there were principles that applied to people who lived in very different circumstances and do not apply to us. Such principles may be unknowable to us. Perhaps they can be unknowable for the very same reasons we are considering: perhaps some people in some remote circumstances possessed mental capacities, perceptual or other, that we do not, and we have no idea what they were, or what they could have been, as they are not some variant of capacities we are familiar with in our or other species, but are entirely unlike anything we know of and in ways that we have no idea about. We cannot deny that there were such principles, after all they applied to the people who had those capacities. And we have no reason to think that these principles do not exist anymore. They do not apply anymore and cannot be known anymore, but that is, given our assumptions, possible.

However, there is, and not only in this case, an asymmetry between past and future. One way of thinking about it is to focus on the concepts involved. The arrival of a new sense will bring

with it a new concept, related to that sense as the concept of eyesight is to the sense of sight. Concepts exist once their instantiation is possible. They may apply to nothing anymore, but that does not mean that we ceased to have the concept. If it becomes unknowable this is contingent on loss of evidence that would enable us to understand it. However, a possible but not yet existing concept is necessarily unknowable, until it comes into existence, if it does.

10. Normative autonomy

I am not assuming a unity of explanations. That is, I am not assuming that all the sound explanations there are can be thought of as parts of one comprehensive explanation, using the same methodological principles, of everything that can be explained. I see no objection to the view that there are radically and irreducibly different forms of explanation, such that there is no explanation of how they all relate to one another, other than by pointing to what they are and are not. We can call this the assumption of radical explanatory pluralism. I assume that the truth or falsehood of that assumption can be known.

I am also assuming what I will call the autonomy of the normative. That is, the paper proceeds on the assumption that we have ways of referring to normative properties that make possible knowledge of the conditions under which they apply, and explanations of why they apply or fail to apply and of the normative consequences of their application, which essentially and non-redundantly employ normative concepts that refer to or apply to normative properties and are not reducible to any non-normative propositions or explanations.¹⁵

The autonomy thesis, if correct, both falls short of and reaches beyond what may be called epistemic autonomy. It falls short of it because it leaves open the possibility of knowledge of the conditions for the instantiation of normative properties based on non-normative evidence. It even allows for knowledge of normative propositions by beings who are incapable of understanding some or all normative properties, except possibly in partial and limited ways. Such people may acquire normative knowledge by imitation – possibly the way most young humans acquire some normative knowledge well before they reach any understanding, and the

¹⁵ Normative autonomy is consistent with some forms of so-called ontological reduction of normative properties. I will not consider such possibilities in this paper. Furthermore, the “normative concepts” I refer to should not be identified with so-called thick concepts or terms, a heterogeneous category, only some of which refer to normative properties.

way some people with depleted mental capacities retain moral knowledge even after losing the ability to understand it.

The autonomy thesis reaches beyond epistemic concerns. It is about explanations of the normative presupposing normative properties and deploying normative concepts expressing them. In this respect, some would regard it as having ontological implications. But here objections are possible. Why assume that explanations of the normative relate to normative properties? Take a familiar example:

- [Conclusion] John has a strong reason to jump out of the window!
- Why?
- [Premise 1] Because he will die by fire, if he does not.

This appears to be a good explanation. The autonomy thesis implies that it is incompletely stated; that anyone who takes it to be adequate implicitly includes in it an unstated reference to a normative property, unstated because it is obvious and needs no explicit mention. Perhaps something like:

- [Premise 2] It would be better for him to stay alive.

However, the objection goes, Premise 2 is not really an invocation of a normative property. It is tantamount to affirming that the explanation is complete, that the specified facts are indeed a reason for the indicated reaction, etc. In other words, some version of a buck-passing (or, perhaps, of an appropriate attitude) account of value would be invoked here. So, perhaps Premise 2 can be rephrased as something like –

- [Premise 2'] There are no other facts bearing on what he should do.

A familiar reply is that the fact that an action has some value properties does not only indicate that certain reasons follow, it explains why they follow. Buck-passers are inclined to say that “This action is good”¹⁶ is too thin a remark to have any explanatory power. It merely states that the action has some (non-normative) feature that is a reason to perform it. On this view, not only is it always true that if one has reason to perform an action then there is some good in the action, it is also true that whenever one has reason to perform an action that is because there

¹⁶ Where it means “this has some value”, and not all occurrences of “good” do mean that, not all are normative.

is some good in the action. But the second claim is false to our understanding of such claims. The fact that there is a reason to perform some action establishes that there is some value in that action, making the first claim true. The question is, however, in virtue of what is there such a reason (and therefore, what makes this action good)?¹⁷ It may be because the action is good in itself (and not merely because there is reason to perform it) or a constituent element in something good (e.g. a movement in a beautiful dance) or that it will produce something good. In such cases it is plausible to say that saying that the action is good is saying that it has some features (say, some consequences) that are reasons to perform it. But not all reasons are of this kind: that I promised to do something is a reason to do it, and that is consistent with the fact that giving the promise was bad and regrettable, and that the act of keeping it has no other merit than being the keeping of that promise. In this kind of case it would be wrong or misleading to say that one has a reason to perform the act because it is good. It follows that saying that one has a reason to ϕ because ϕ -ing is good is informative in a way that buck-passing does not allow: it tells us something about the character of the explanation of the reason we have, it has a character that not all fact-producing reasons have.

However, the most important failure of buck-passing accounts is in dealing with other value properties. Suppose the reason for some action, e.g. refusing to answer a question whether John was in Knox Street on the first of May, is that answering the question would be disloyal to a friend. That entails that there is some good in the action (the refusal). It does so by specifying what good that is: the good of loyalty to friends. If buck-passers take this to mean that my telling whether John was in Knox Street is bad are missing the gist of the matter, and ignore the explanatory power of the value statement. They cannot retreat to saying that they mean that answering the question is bad in the special way in which telling about a friend is bad. To say that is to give up on being a buck-passer. It admits that there are different values which differ in ways other than that they are reasons for action. Worse than that, this way of understanding my situation (a) employs irreducible value concepts, namely that of telling and of friendship, and (b) still does not get to the value that actually motivates me: loyalty, as not all telling on friends constitutes disloyalty, and there are other ways of being disloyal than telling.

¹⁷ Note that "action" refers to what is sometime called action-type, i.e. generally one can perform an action more than once, and others can perform the same action one does.

These reflections do not prove the autonomy thesis, but they lend it some support.

11. Queer properties

I have been defending normative autonomy in this paper because the paper presented an explanation of normative change in terms of changing normative properties in a way that implies normative autonomy, or at any rate a way that would encounter additional difficulties if there is no normative autonomy. However, the implied realism about the existence of normative properties encounters the charge that they are peculiar properties, the charge whose current form originates with John Mackie. There are various ways in which the charge can be understood and developed, and I will present my own understanding of it. The gist of it is that normative properties are not connected to reality because (1) their instances do not interact causally with the rest of reality; (2) their existence cannot be explained in the way that we explain other aspects of the reality; (3) if they are properties of objects or events it would be a mystery how they can move people to act as they are supposed to do; (4) they give the impression of real properties because within groups with shared culture there are conventions that govern discourse about them, and members of those groups by and large observe those conventions, with the result that it appears as if discourse purporting to refer to these properties can be true or false, whereas in fact it is merely conforming to the conventions or failing to do so.

Briefly, and admittedly inconclusively, a few comments addressing these points in reverse order. (4): As this paper argues for the possibility that normative principles can change, and illustrates such changes by appeal to cultural changes, the allegation that if normative properties exist they depend on social conventions is not exactly an objection. True, they are not, for the most part, the product of conventions¹⁸, but – and this is the whole point – they are dependent on social practices. Is there any objection to the possibility of properties whose existence and instantiation depend on social practices, properties without which we would not have institutions and their products such as money, intellectual property, novels, or poetry? (3): As an objection this point (it is a mystery how they move people to act) is out of place. It amounts to no more than denial of the existence of normative properties. That they provide normative

¹⁸ Though see: A. Marmor, *SOCIAL CONVENTIONS: FROM LANGUAGE TO LAW* (Princeton University Press, 2009), Chapter Three, on deep conventions.

reasons is the essential feature of normative properties. And normative reasons are essentially capable of moving people to action, thought, emotions or imaginings. So, we should take the point not as an objection but as a question: normative properties are peculiar and call for explanation. The explanation, or explanations, come in many parts. Some explain the relations between reasons and actions or beliefs or emotions or imaginings, others explain the specific character of each normative property.

The second part of the peculiar properties objection is, at first blush, not an objection. Why should these properties be explained in the same way, using the same patterns of explanations, etc., as others? If we assume that all explanations have the same character we are of course in trouble. It turns out that non-normative properties are peculiar because they cannot be explained in the way that normative properties are explained. So, the objection must be taken to be more radical: there are no possible explanations of normative properties. The issues raised are too numerous to be even mentioned here. But just to indicate the direction of travel let me divide them into two types: disputes regarding specific purported normative properties, and disputes about the more general possibility of there being any normative properties. The two feed into each other. Think of the explanation of specific properties first: Of course, not all purported normative properties exist. Some are illusory. The question is whether any can exist, and be intelligible. One possible argument proceeds by (a) criticizing all the explanations of all normative properties that have been so far offered, and (b) showing that their failure is not that they are incomplete, but that they are in principle no good, and (c) that their failure of principle is of the same kind in all cases. That, perhaps with one or two additions, would constitute an inductive argument to the conclusion that no explanation of any purported normative property can succeed. I detect nothing that gives one hope in the success of such an argument.

A more common alternative is to argue: (a) True, we know that (in certain contexts) people use arguments about some normative properties that carry conviction even though, needless to say, not always with their intended audiences. This may be thought to point to a tendency for views to converge when those who hold them are well informed. More importantly, a larger degree of convergence exists regarding an understanding of what constitutes a good argument. However, these contexts are ones where speaker and audience share a common culture and are simply relying on unexamined common assumptions and conventions. (b) Beyond such islands bound by local conventions there is no sign of agreement or of a tendency to converge on a common

understanding of what constitutes good argument. That means, the objection goes, that there are no objective standards of judgement regarding the existence or character of normative properties, and no possibility of such standards. Hence, they do not exist.

In assessing this and related arguments it is important to trace the causes of disagreement. For example, it seems to me that many of them derive from non-normative disagreements, especially from disagreements about matters derived from, or being part of, religious outlooks, but which are not themselves normative. Belief in miracles, and the use to which such belief has often been put, is an example. Non-normative beliefs that are not necessarily involved with any religious outlook are also at the root of some normative disagreements. Racist and homophobic views are often nourished by ignorance, entrenched by deep psychological disinclination to correct mistaken beliefs about these matters. When this is the source of normative disagreements it may be taken to cast doubt about the possibility of non-normative properties. More plausibly these sources of persistent disagreements show that the significance of disagreements may have been exaggerated.¹⁹ There is another common source of disagreement of interest here. Many disagreements are sustained by mistaken theoretical views about the nature of normative facts. For example, I have long argued that indeterminacy and incommensurability are very common within the normative domain. Many people think that that is necessarily false. This view inclines people who see reason to believe that their own way of life is legitimate to believe that ways of life which are incompatible with theirs are not legitimate, because they reject as necessarily false the correct view that both are legitimate. One ought, and it is possible, to explore in detail these sources of disagreement, and others like them – like them in that their explanation refutes the argument from disagreement by showing that its sources are not due to there being no normative properties, and do not cast doubt on their existence.²⁰

¹⁹ I am assuming that the claim that regarding objective matters there is a convergence towards agreement at the end of inquiry, either as a mark of objectivity or in any other way, has been effectively criticized, mostly as empty, and need not be considered in this context.

²⁰ Arguably, if people, perhaps the people of some society, share fundamental theoretical mistakes then we are in Error Theory terrain, for a shared understanding of the basic characteristics of normative claims at least partially determines their meaning. Be that as it may, the crucial point is that that does not mean that one is free to “invent” a new meaning. The correct theoretical understanding is not invented. It is determined by the same considerations that determine the theoretical underpinning of any area of common

Remembering that various doubts about normative properties ignore the possibility that the existence and instantiation of normative properties are the product of social practices, the preceding comments soften the force of many arguments against the possibility of such properties. But they leave the core doubts (as I understand them) untouched. Doubts about the ability to explain normative properties, if they are real non-naturalistic properties, may derive from the claim that they are not connected with the rest of reality in a way that would enable us to refer to them, to come to know them, or, therefore, to explain them.

12. The Connectedness of Normative Properties

At first blush the charge that normative properties are peculiar because they are unlike anything about the world that we understand, and that the explanations that enable us to understand the world around us, namely scientific explanations, do not apply and cannot be used for explaining normative properties, is appealing to contemporaries who came to see the natural sciences as the only source of knowledge and understanding. As a matter of fact, few people understand the natural sciences beyond a rudimentary level and few people understand aspects of the world surrounding them through the natural sciences. Instead people trust (quite sensibly) that other people do know science and that the natural sciences explain aspects of reality to the satisfaction of those who understand them. On the other hand, all those people whose knowledge of the sciences is poor know and understand much of the world around them in ways that are independent of the sciences, and that the natural sciences do not explain, and this includes their understanding of themselves and other living beings, and of social practices and institutions, and the culture and economic opportunities, or lack of them, that those social practices and institutions sustain.

This is not to cast doubt on the importance of the natural sciences, nor to deny that psychological and social and economic phenomena benefit greatly from specialist expertise, even though it is not widely shared in the population. I am only denying that they are the sole sources of knowledge and understanding for human beings.

discourse. There are issues here of the possibility of providing a correct account of the theoretical underpinning of a range of phenomena when many people hold mistaken views on their character.

Nor is knowledge and understanding of normative phenomena unconnected to our knowledge and understanding of other phenomena. We have a fair grasp of the way normative phenomena depend on the existence of certain physical, psychological, social and economic conditions, even though there is much more that remains to be discovered about these dependencies.

Normative explanations of the conditions for the application of normative properties typically rely on a mix of normative and non-normative phenomena. And the normative properties of how things are in the world have a considerable impact on the way matters develop in the world, even if this impact depends on people realizing how things are normatively speaking.²¹

Throughout this paper I have been relying on the assumption that there may well be distinct types of explanation and argumentation in different domains.

There is much that, so far as I can see, we do not understand. We do not know how far the explanatory methods we rely upon in various sciences and other domains can reach. We know that they have changed and were transformed during history, and expect that they will further develop in the future. But we do not know how far they can extend, and therefore we do not know whether their future development will lead towards greater integration of the different types of explanations known to us today, or towards their absorption in more comprehensive methods of explanation. That means that we are ignorant of the reasons for the existence of different types of explanation, and their inter-relations. As a result we are also in the dark as to the possibility and nature of conflicts between types of explanations. Various scientists are quick to declare that they have explained what was unexplained before about human conduct and more, and that their explanations replace the common, non-scientific explanations we rely on. Often enough such claims are exaggerated, and while the new scientific knowledge is valuable, it does not conflict nor make redundant the types of explanations we had relied on before. There is much still to discover about these matters, and until we do, many of our views, including those expressed in this paper, remain tentative and conjectural.

²¹ And, to repeat: the impact of the way things are normatively is not the impact of how people are brought up, socialized, etc. Needless to say, people can understand normative principles, just as they can understand physics, only if they are brought up and socialized in certain ways. But as most values are cultural products, this increases the reliability of people's views of normative matters, rather than undermining it.

