

Moral principles as generics

Ravi Thakral
ravithakral@icloud.com

Forthcoming in the *Journal of the American Philosophical Association*

Abstract: I argue that moral principles involve the same sort of generalization as ordinary yet elusive generic generalizations in natural language such as 'Tigers are striped' or 'Peppers are spicy'. A notable advantage of the generic view is that it simultaneously allows for pessimism and optimism about the role and status of moral principles in our lives. It provides a new perspective on the nature of moral principles on which principles are not apt for determining the moral status of particular actions while they may be apt, and even fundamental, to our acquisition of moral knowledge. A natural consequence of the view is variation amongst moral principles, with some regularly warranting exceptions, and some appearing arguably exceptionless. It is also argued that this generic conception of moral principles has significant advantages, as a normative model of moral reasoning, over the view of moral principles as defaults advanced in recent years.

Keywords: generics; moral principles; moral reasoning; moral particularism; moral epistemology

1 Introduction

Moral principles, in ordinary cases of moral deliberation, are usually expressed by sentences such as the following:

- (1) a. One ought not to steal.
- b. Lying is wrong.
- c. If you make a promise, you should keep it.
- d. It is wrong to murder.

In this paper, I advance a new understanding of moral principles. I argue that we should think of such principles as generic generalizations. Moral principles are generalizations but they are only quasi-general. They mirror common yet elusive constructions in natural language such as:

- (2) a. Koalas sleep all day.
- b. Americans drink coffee.
- c. Tigers are striped.
- d. Generation Z cares about the environment.

Generics express generalizations with a kind of quasi-universal flavor: they communicate generalities without a full-on commitment and remain true even in the face of counterinstances. These generalizations are conveyed without any lexical item responsible for telling us about the nature of the generalization in question. I argue that taking this perspective seriously reveals a novel argument that undermines our reliance on moral principles in ethical life and allows us to understand the moral realm in a new and intriguing way. If moral principles involve the same sort of generalization as ordinary generics, then there are two main consequences:

- The correct understanding of moral principles is then one on which there is an incredible amount of variability: moral principles remain true even without it being the case that they are successful in giving the right verdicts in particular situations.
- Moral principles do not come with any information about how many instances are required to conform to the generalization in order to be true: information about the prevalence of a particular property does not come from a generic generalization itself.

The status of the idea that moral principles admit of exceptions has gained serious appeal, particularly since the work of Ross (1930) on prima facie duties. However, there has been significant divergence over what this may mean for the role of moral principles in our lives. And there are also significant questions regarding how exactly to understand moral principles in the first place, if they are not universal and exceptionless.

It is commonly thought that moral principles serve a distinguished role in moral reasoning. We can use them in reasoning because we believe that their content is suitable enough to get the job done. Philosophers have understood moral principles in a broad variety of ways: from being exceptionless universally quantified generalizations to weaker generalizations allowing for defeasibility. No matter which route we take, the destination is the same: moral principles can be the source of our knowledge of the moral features of particular actions.

If my arguments are right, then it allows us to think of moral reasoning in a new way, particularly in cautious terms. The generic view therefore reveals a further angle in the controversy over the nature and role of moral principles in that it involves an account of moral reasoning that is neither deductive nor defeasible and instead gives us motivation for accepting a weaker theory of reasoning. The generic view also reveals an understanding of moral principles where there is considerable variation, especially the extent to which principles might allow exceptions. These ideas are introduced and defended in Section 2.

Accepting the generic view leaves plenty of room for a positive role for moral principles. In Section 3, I argue that the generic view offers an interesting and productive way to think about the role of principles in areas such as moral explanation, moral knowledge, and moral education.

In Section 4, I discuss one of the contemporary competitors to the generic view, namely the view that moral reasoning is a form of non-monotonic reasoning, which has been defended independently by Horty (2007, 2012) and Thomas (2007, 2011). I argue that conceiving of moral principles as defaults is compatible with the idea that such principles are not associated with information about prevalence, but that construing moral reasoning non-monotonically can make us miss some of the normative considerations surrounding moral reasoning. In Section 5, I conclude by briefly discussing some further questions for research.

2 The Generic View

I will build up the generic view of moral principles in three main steps. The first position defended in this section is that the most plausible view of the semantics of ordinary moral principles is that they express generic generalizations. In particular, I will argue that moral generalizations are generalizations in the sense that generics are generalizations: such generalizations exhibit the quasi-generality that we typically associate with ordinary generics. The second part takes a closer look at the nature of generic reasoning and its weaknesses. And in the third part, building on the aforementioned discussion, I advance the position that the generic understanding of moral principles provides a new line of opposition against their presumed role in moral reasoning.

2.1 Evidence for genericity

The project of providing adequate truth conditions for generics is challenging as our intuitions vary widely between different generics and whether or not a generic is true need not correspond with statistical regularities. ‘Ducks lay eggs’ is widely accepted as true yet it is only

female ducks of reproductive age that lay eggs. So, a majority of ducks do not need to lay eggs in order for the generic to be true. It is widely accepted as true that ‘Ticks carry Lyme disease’ yet only around 1% of ticks are infected. So, we accept the generic even at such a low prevalence level. At the same time, there are generics that are unacceptable even when the majority of a kind possess a property. We do not accept ‘Books are paperbacks’ as true despite the fact that the vast majority of books are in paperback form.

Examples of these various kinds abound. The amount of variation in the data is wide and our inability to account for the meaning of generics in terms of statistical regularity is only one example of how complicated matters can be. The landscape of options for semantic theories of generics is thus broad and diverse (Leslie and Lerner 2022; Sterken 2017). The majority of theorists hold that the underlying logical form of generics features a covert operator, Gen, which is responsible for the unique quasi-generality of generics. But there are also others, such as Liebesman (2011), who argue that generics involve kind predication at logical form; this problem has resulted in significant debate (Leslie 2015; Sterken 2016; Collins 2018). Despite these controversies, there is broad agreement that certain properties distinguish generics from other constructions. And it is such factors, not the semantics, that will be important in building the case that moral principles should be recognized as generics. For our purposes, we will work with minimal yet central distinguishing characteristics of generics:

Resistance: Generics can remain true even if there are counterinstances.

Non-Numerity: Generics do not carry any information about how many instances are required in order to be considered true.

The features of Resistance and Non-Numerity help distinguish generic generalizations from generalizations involving overt quantifiers, as the latter lack these two features. Resistance and Non-Numerity are closely related in the sense that Resistance says that generics can tolerate exceptions while Non-Numerity says that generics do not tell us how far this toleration extends.

First, generics possess the property Resistance because they tolerate exceptions, yet we cannot express their truth conditions in the way we can with generalizations involving overt quantifiers. In the case of quantified generalizations, we know that, for instance, a universally quantified generalization would be false if there were a single counterinstance; a *most*-quantified generalization would be false if most of the instances are actually counterinstances; and we know that a *some*-quantified generalization would be false if there is no instance acting as a witness to the generalization. With generics, however, there is no way of similarly specifying how it is that a generic can be false.

Second, generics possess the property of Non-Numerity because they do not provide us with any information about how many instances of a generalization are required in order for it to be considered a true generalization (Carlson 1977). Suppose one is asked ‘How many tigers are striped?’ It is possible to reply by suggesting that *all/most/some* tigers are striped, but it would appear problematic if one were to reply by suggesting that tigers are striped. This second point is important. It not only cleanly distinguishes generics from ordinary quantified sentences, but it also highlights one of the very remarkable features of generics which makes them so intractable. The idea that generics are non-numerous reveals that they are independent of prevalence in the sense that the truth of a generic is consistent with varying levels of prevalence of cases which conform to the property specified by a generic.

The claim defended in this sub-section is that moral generalizations are generalizations in the sense that generics are generalizations: moral generalizations exhibit the variability and quasi-generality that we associate with generics. I will do this by arguing that moral generalizations possess the features of Resistance and Non-Numerity.

First, moral principles exhibit the property of Resistance because it is well-accepted that paradigmatic moral generalizations tolerate exceptions. At the same time, however, it is also understood that there are general difficulties in positing a proviso to handle the exceptions because there is no straightforward way to specify the content of such a proviso which can account for the various possible permissible exceptions. It is important to note that this reveals that a moral generic such as ‘Stealing is wrong’ is one that resists statistical explanation: stealing is still wrong even if it were to turn out that overwhelmingly many cases of stealing turned out to be somehow exceptional or perhaps blameless cases of stealing.

Second, moral principles pass the criterion for Non-Numerity: such principles do not wear any information on their sleeves concerning how many instances of a given moral generalization hold true. Suppose one is asked ‘How many instances of stealing are wrong’? Depending on what the present state of affairs is like, it may be appropriate to answer by suggesting that *all/most/some* acts of stealing are wrong; it would, however, be infelicitous to reply by putting forward the generic claim that stealing is wrong. This means that we get a version of the Non-Numerity principle in the moral case:

Moral Non-Numerity: Moral principles do not carry any information about prevalence, that is, they do not carry any information about how many cases need to appropriately conform to the principle.

It is worth noting that the examples of moral principles I have highlighted so far are grammatically different from the examples of generics I have discussed. The majority of research on generics focuses on the bare plural form, although there have been significant investigations of indefinite singular generics, and, to a lesser extent, definite singular generics. There are notable differences amongst these grammatical forms, particularly in what type of generalization appears to be expressed. Some generics are acceptable in the bare plural form, but not in the indefinite singular. The generic ‘Bears are wild’ sounds acceptable yet somehow ‘A bear is wild’ is not. One traditional suggestion is that the bare plural is associated with inductive generalizations and the indefinite singular is associated with normativity (Carlson 1995). In general, there are further questions left open about what type of genericity non-numerous moral principles might be associated with.

In addition, the features of Resistance and Non-Numerity help distinguish generic generalizations from explicitly quantified generalizations, but they do not completely distinguish generic generalizations from other types of non-quantified generalizations. There are non-generic generalizations that possess the features of Resistance and Non-Numerity (such as *ceteris paribus* generalizations). There are also views of moral principles that are compatible with these two features. Lance and Little (2004, 2006, 2008) treat moral principles as “defeasible generalizations” which can be genuinely explanatory; Robinson (2008, 2011) argues that moral principles should be treated as “moral dispositions” against the dominant view of moral

generalizations as true and law-like; Väyrynen (2006, 2009) defends a view of moral principles as “hedged generalizations”. Väyrynen, in particular, defends an account of hedged moral principles which allows for exceptions, including the feature that such principles can still be true “even if only few lies happened in fact to be wrong” (2009: 93) and also draws a connection to generics as part of his discussion.¹

One of the distinctive features of the generic view of moral principles defended here is that it arguably accommodates for a broader range of variability. Although there are frameworks, such as the ones just mentioned, which can accommodate the idea that moral principles can be subject to innumerable exceptions, my conjecture is that the generic view puts us in a better position to appreciate and accommodate the heterogeneity of moral principles. I will elaborate on this point further through discussion of a crucial test for genericity.

Canonical examples of moral principles appear to pass some standard tests associated with generics. For instance, moral principles do well with the standard adverbial quantifier insertion test (Krifka et al. 1995). According to this test, we have good reason to hold that a given sentence is a generic if the insertion of an adverb of quantification renders only the slightest change in meaning. For instance, inserting ‘usually’ into the generic ‘Birds fly’ gives us ‘Birds usually fly’ which only results in a slight change in meaning. To apply this test in the case of a moral principle, consider the following:

(3) Stealing is wrong.

Then combine with an adverb of quantification:

¹ I am grateful to an anonymous reviewer for encouraging me to address these important points.

- (4)
- a. Stealing is *usually* wrong.
 - b. Stealing is *normally* wrong.
 - c. Stealing is *typically* wrong.
 - d. Stealing is *generally* wrong.

These paraphrases sound appropriate, and also help bring out the prima facie exception-granting character of the moral principle in (3). Overall, it feels reasonable to say that the addition of the adverbs in (4) only results in a slight change in meaning. If this is right, then we can conclude, by the test, that there is good reason to consider (3) and other related principles, whether contributory or overall, to be generics.

I imagine that the majority of readers would agree with my linguistic judgments, but it is important to acknowledge that many readers will have some trouble in hearing even the ordinary, commonplace moral principles as generics. Some readers may not be making judgments in a more ‘neutral’ way as I would imagine. Sometimes metaethical commitments can readily influence some linguistic judgments. Strongly committed Kantians, for instance, would of course firmly disagree with the intuitions that I have tried to tease out, and may opt for some alternative explanation of the appearance of genericness. To those inclined in such directions, I would wish to emphasize the following: the evidence we are appealing to concern our ordinary practices, and the issue of whether the principles we deploy in reasoning are epistemically appropriate is one which is rather independent of the philosophical stances we might have within ethical theory.

This test, however, is merely suggestive and not conclusive. One significant complication is that generics can express generalizations of varying strength, with some generics expressing rather weak generalizations and some expressing rather strong generalizations (Sterken 2015;

Nguyen 2019; Almotahari 2022). We can add ‘typically’ to a strong generic such as ‘Tigers have stripes’ however it would not seem appropriate to paraphrase the meaning of weak generics such as ‘Mosquitos carry West Nile Virus’ in the same manner.²

I claim that such variation in the strength of the generalization expressed is an important part of revealing the heterogeneity of the moral domain. Some moral principles convey very strong generalizations, such as ‘Pleasure is good’, however many ordinary, commonplace moral principles seem easily exception granting and appear to express something weaker: it may be morally right to steal bread to feed your family and it may be morally right to lie if it saves an innocent life from danger. But not all moral principles appear so readily exception granting. It is far too difficult to find exceptions to certain moral injunctions, especially when they are worded in the right way. It is of course doubtful that one could find reasonable and realistic cases where torturing babies could ever be a morally right course of action. It is harder to find cases where it is morally right to harm the innocent and weak than it is to find cases where it is morally right to lie, cheat, or steal. It is therefore important to note that accepting the generic view does not necessarily mean that every single imaginable moral principle, whether foundational, or ordinary, has to be exception granting. In fact, this is something we should expect because paradigm cases of generics may also vary in their generic quantificational force.

If there are some general principles formulated in such a way that they really must be exceptionless, this is an important observation from the perspective of the generic view. My view is that the existence of such apparently and truly exceptionless principles reinforces the generic view; in fact, this is something we should be able to predict given that other sorts of generic

² I am grateful to an anonymous reviewer for raising this issue.

generalizations, whether about the natural or social world, involve considerable variations in prevalence. There are many ordinary generics, especially of the rules-and-regulations variety, such as ‘Bishops move diagonally’ or ‘Basketball players shoot free throws when fouled’ which do not easily have reading where they admit of exceptions. Such cases are generally tricky for most semantic theories of generics. But, from the point of view of evaluating the generic position on moral principles, the very existence of this sort of variation points to an important and interesting parallel: as ordinary and well-studied cases of generic generalizations exhibit broad ranges of quantificational force and prevalence, so do moral generalizations. A significant upshot, then, is that treating moral principles as generics does not assimilate them into a single category but helps reveal their broad diversity. It would take some additional argumentation to establish that competing views are unable to account for this diversity; however, I have argued that the generic view puts us in a very good position to appreciate these important nuances in our understanding of the moral realm.

2.2 Generic reasoning

I argue that the generic view of moral principles leads to a view of moral reasoning on which we should exercise more caution with moral reasoning than we would ordinarily appreciate. Here is a way of stating the argumentative strategy: if we understand the epistemology of moral reasoning in terms of the epistemology of generic reasoning, then we end up with a pessimistic view about the prospects of deriving justification for particular moral judgments from moral generalities. I will now comment on the epistemology of reasoning with generics and then pursue the consequences for moral reasoning in the next sub-section.

There is an important clarification and observation to start with. We need to distinguish between our justification/knowledge of generic generalizations and whether our justification/knowledge of such generalizations supports inferences concerning particular entities or individuals that are said to possess the property specified by a generic. Questions about such issues are some of the core problems at the heart of the epistemology of generics. The focus here will be on the latter question of the inferential support provided by generics and whether we can derive justification concerning whether or not some individual under the scope of a generic generalization possesses the property in question.³ In order to approach these problems, I will start by considering some issues from the literature on generics and defeasible reasoning.

Several researchers have connected their approaches to the semantics of generics with the literature on defeasible validity because it has been claimed that generics support certain defeasibly valid inferences (Asher and Morreau 1995; Pelletier and Asher 1997). Such inferences may take the following or related forms:

(5) **Generic Modus Ponens**

- a. Frogs jump.
- b. Kermit is a frog.
- c. So, Kermit jumps.

The conclusion that Kermit jumps is said to be defeasibly supported by the generic claim that frogs jump. The premises are said to support the claim that Kermit jumps, yet we might also find that the inference is defeated because we may learn, for instance, that Kermit does not jump

³ See Schiller (forthcoming) for a recent discussion of issues bearing on the former question and their connection to inductive inference.

due to some genetic defect. We have the intuition that the inference is reasonable, and thus should be considered defeasibly acceptable. If we take such intuitions seriously, then this is an important matter for semantic theorizing. Several semantic accounts of generics are a result of significant investment in ways to account for defeasible validity, especially through carefully crafted conditional analyses.

But it is worth questioning how much priority we need to give to our intuitions about the defeasible goodness of inferences involving generics. There is a simple argument from the variability of generics that can be given. There are several generics from which it appears as though we are unable to derive any reasonable consequences. Leslie (2007) uses the following low prevalence cases to illustrate the point:

(6) **Buzzy**

- a. Mosquitos carry West Nile Virus.
- b. Buzzy is a mosquito.
- c. So, Buzzy carries the West Nile Virus.

(7) **Beaky**

- a. Ducks lay eggs.
- b. Beaky is a duck.
- c. So, Beaky lays eggs.

In either case, it seems undesirable to draw conclusions about arbitrary members of a kind, given that we know that the generalizations deployed in the major premise of each inference are low prevalence generics. There might be other properties that we may feel

comfortable attributing to arbitrary ducks or mosquitoes, but not these ones. If we take cases like this seriously, then it appears as though we have part of a challenge to the notion that accounting for nonmonotonic inferences supported by generics should be a desideratum when theorizing about generics. At best, we should likely rethink the role we are assigning to the intuition that some patterns of reasoning such as Generic Modus Ponens are defeasibly good: this does not seem to be a pattern of reasoning that we want to associate with generics across the board.

There is a more general argument to be made here regarding the issue that there is a high degree of variability in the acceptability of nonmonotonic inferences involving generics. That is, there is more to say even about the paradigmatic cases of generic reasoning themselves. The inference to the claim that Kermit jumps seems almost uncontroversially reasonable, however, I claim that even this judgment is not as robust as it appears: closer examination reveals that we should find the inference questionable or at least marginally acceptable.

Consider the following: it is known that frogs jump, that Kermit is a frog, and no other facts about Kermit are known. Although we know the generic information that frogs jump, by Non-Numerity, it follows that we are not in a position to know any facts pertaining to the prevalence of jumping frogs or anything about the prevalence of jumping amongst various subsets of frogs. Such information about prevalence is not given by, nor follows from, the meaning of the generic. Suppose, however, that it is known that Kermit belongs to a group of frogs which have a high prevalence of jumpers. If this is the case, the inference that Kermit jumps would appear well-supported. But in the case where information regarding the prevalence of a generic in a given situation is known, it would appear likely that we should attribute any

resulting sense of inferential goodness to the facts about prevalence themselves and not the generic in question.

Furthermore, absent any information about prevalence, the inferences we draw from generics would not then appear well-supported. Knowing the generic generalization that frogs jump should not make us feel entitled to infer anything about particular frogs on the basis of the generic alone. And this point easily generalizes across the whole range of generics, regardless of whether there is witness to sufficient prevalence.

The property of Non-Numerity reveals an important lesson regarding the relationship between generics and defeasible reasoning. Without sufficient prevalence as a given, it is hard to get a handle on what the notion of defeasible validity involving generics is meant to capture; it is also not clear whether the various proposed formal principles would turn out to be helpful if our judgments can vary so easily across a range of situations. The general epistemological lesson we can extract here is that knowledge of generic generalizations do not themselves support our beliefs about particular cases under the scope of such generalizations. However, it is a controversial matter, in general, whether we should think of propositions as supporting conclusions on their own or only together with supplementary information, and many will agree that what conclusion a generic supports should depend on the information it is coupled with. The core point, however, to appreciate in this context is that generic generalizations do not support their instances in the way we would expect generalizations to do so.⁴

At the very least, what we get is an important point of caution in the epistemology of generics: we should accept a position of caution when relying on generic generalizations in

⁴ I am grateful to an anonymous reviewer for encouraging me to clarify this.

reasoning because they are unable to confer the appropriate kind of justificatory support required when reasoning about particular cases. Even very little support provided by knowledge of a generic generalization is not enough to provide robust justificatory support for our beliefs about particular cases. Accepting the truth of a generic does not entail that there is a sufficient prevalence, in the current circumstances, of whatever property is specified by the generic. We know this point well from the cases where we know we are dealing with low prevalence generics. But the point is a general one. Suppose we consider a supposedly high prevalence generic such as ‘Tigers are striped’. We know that it is true that, in the actual world, a vast majority of tigers are striped. However, even if there is a felt presumption in favor of cases of generic reasoning, this presumption is itself unstable and potentially influenced one way or the other by extraneous factors—perhaps our beliefs about prevalence or some other information. In other words, it would be a mistake to think that any apparent epistemic goodness of generic reasoning comes directly from generics themselves. Therefore, any felt presumption could not serve the relevant justificatory role.

It is also important to distinguish the descriptive from the normative here. Recent work in psychology has shown that once subjects accept a generic, they are disposed to treat the generic as inferentially powerful in the sense that they will believe that an arbitrary member of a kind possesses the property in question regardless of their beliefs concerning prevalence (Khemlani et al. 2009; Prasada et al. 2013). However, in characterizing the nature of reasoning with generics it is important to note that the account of reasoning we are interested in is deliberately not fully descriptive. The concern is not only with capturing how it is that we, as a matter of actual practice, reason, but, more importantly, the concern is with how we *should* reason. The epistemic question of whether there is justificatory support for beliefs about particular cases on the basis of

generics is itself an inherently and thoroughly normative question: it is a question either about how we are permitted to reason or how we ought to reason.

2.3 Generic moral reasoning

The core lesson from considering the epistemology of generics carries over to the case of reasoning with moral principles. This is so not only because of the plausibility of the idea that moral principles are generic generalizations, but also because a closer examination of moral reasoning itself reveals similar features to generic reasoning. The overall consequence for moral reasoning given the generic view depends on two claims:

- That moral principles are generic generalizations.
- That generic generalizations do not provide the appropriate inferential support in reasoning.

My view is that there is strong evidence behind both of these claims; if we accept them, then this naturally leads to an understanding of moral reasoning on which we do not ultimately derive robust epistemic support for our reasoning in particular cases.

The generic view of moral principles yields a perspective which points to Moral Non-Numerity. And if we accept Moral Non-Numerity, then this means accepting a view where moral reasoning is associated with significant caution and does not ultimately bestow us with support for our particular moral judgments. The support that we do get for our particular moral judgments may derive from elsewhere.

The theoretical landscape for understanding the nature of moral reasoning is largely divided into deductive and defeasible approaches. Perhaps the most prominent view of moral reasoning throughout the history of Western philosophy is that inferences drawn from moral principles are deductive inferences. The most orthodox theory of moral reasoning arising out of the generalist tradition is the theory on which moral reasoning is a form of deductive reasoning. We might characterize the core of this view as follows: the correct psychologically descriptive and normative account of moral reasoning is one which conforms to the structure of deductive arguments. The deductive model of moral reasoning enjoys a form of historical prominence in that it was arguably held throughout the history of philosophy up to recent times. Demos (1958), who was, relatively speaking, an early detractor of the idea that moral reasoning is deductive, notes: “The fact that moral principles are stated in the form of universal propositions has made it appear that they function as premises in deductive inference” (153). The deductive model of moral reasoning enjoyed prominence even in the twentieth century, with, for instance, Hare (1953, 1963) explicitly endorsing and advocating this position; for instance, Hare says of moral reasoning that “the only inferences which take place in it are deductive” (Hare 1963: 88).

The position that moral principles are generic generalizations is clearly not compatible with the stance that moral reasoning should be characterized in terms of deductive reasoning. My contention is that the generic view leads to a view of moral reasoning that is neither deductive nor defeasible. Furthermore, my position is that the theory of moral reasoning implied by the generic view fits well within the particularist tradition in moral philosophy. It is hard to accept the understanding of moral principles as generic generalizations without accepting these particularist implications. This is because Non-Numerity reveals that the weakness of moral principles as generalizations means that they are not strong enough to recommend defeasible

support in reasoning. However, the evidence for the generic view of moral principles does not imply some of the more radical implications about the status of moral principles, such as the idea that there are no true overall or contributory moral principles.

Smith (1988) helpfully distinguishes between the role of moral principles in serving either our theoretical or practical aims. This is important for our purposes here. Broadly speaking, moral principles can be said to serve either an explanatory or action-guiding role. But there are several other recently proposed theoretical objectives associated with moral principles, including their role in grounding (Väyrynen 2013; Enoch 2019) and supervenience (Enoch 2011). These various aims have had a significant influence on the theoretical directions in philosophical discussions of moral principles, especially against the backdrop of the generalism-particularism dispute over the previous decades.

A core component of the traditional stance is that moral principles provide robust support for reasoning about particular cases. One of the widely recognized aims of moral theory is to elucidate the principles which characterize and govern our actions. Accounts of the structure of morality may differ not only according to which principles are the foundational ones, but also in terms of whether there is a single foundational principle at all. One way of putting the thought is that if moral principles are general enough to be applicable to all persons, at all times, and in all circumstances, then they should have a very significant role in moral reasoning. There are, of course, a number of notions of ‘reasoning’, some which are particularly concerned with, for example, unconscious thinking, including biases and heuristics which influence moral judgment (Sunstein 2005). The focus here, however, is on the philosophically more paradigmatic notion of moral reasoning which is reflective, and which can be carefully reconstructed.

The debate between generalists and particularists is an intriguing one involving multiple angles, some to do with metaphysics, and others to do with epistemology. However, one of the major issues concerns the very nature of moral principles and what they ultimately recommend in particular instances. Broadly speaking, those who take moral principles as exceptionless will see such principles as specifying properties that always count in favor or against the rightness or wrongness of a particular action. Dancy (1993, 2004) famously disagrees that this is how moral reasons work and advocates for variability in the theory of reasons, although it is highly contentious whether Dancy's appeal to variability is successful at undermining the status of general moral principles (McKeever and Ridge 2006).

Many intriguing and influential competing accounts of moral principles have then emerged as a result of these debates, such as the ones defended by McKeever and Ridge (2006) and other views I previously discussed in relation to Non-Numerity (by Lance and Little 2004, 2006, 2008; Robinson 2008, 2011; Väyrynen 2006, 2009). There is much to be said about the details of these various accounts and others. The main point I wish to highlight here is that while there have been many attempts to make sense of variability associated with moral principles, the generic view goes a step further and says that the variability is more than what we appreciate. In addition, accepting the argument from Non-Numerity allows us to rethink the nature of moral reasoning. It is important to distinguish between the question of whether epistemic support exists in reasoning and the normative question of whether we should be relying on principles in reasoning, given the evidence that has been unearthed about their underlying nature.

3 The generic moral landscape

I have argued for evidence that ordinary moral principles are best understood as generic generalizations and that accepting this position can lead to a novel form of skepticism concerning the nature of moral reasoning. Thus far, I have defended the broad position that the generic view of moral principles has greater affinity to the particularist tradition in moral philosophy. The generic view offers something distinctive to particularists as they face the challenge of offering an understanding of principles in a non-substantive manner (Stangl 2006). Although particularism is largely associated with the work of philosophers such as Dancy and Lance and Little, it is best to think of particularism as a broader family of views, perhaps weakly unified by a shared opposition to some or several aspects of the place of exceptionless moral principles, or perhaps as a research program (Leibowitz 2009). Moral particularists, in this inclusive sense, wish to say something enlightening about the role of moral principles in various domains such as moral education. However, it would seem difficult to say anything about such topics while denying that moral principles have a prominent standing.

Thus, considering the generic view can also be valuable for those who are broadly sympathetic to aspects of moral generalism in that there is something substantial to be said about the role of moral principles in our lives, despite the pull of arguments that their role in reasoning is rather limited. At this juncture, I want to comment on some implications on the generic view in this spirit. I will suggest that this view can open further avenues of inquiry for particularists and sympathetic generalists with respect to explanation, knowledge, and moral education.

When we take the generic route, there are going to be some important complexities to consider, given the various difficulties in theorizing about generics, not only in terms of the

intractability in determining their truth conditions, but also in terms of the types of generalizations that are invoked by different generic noun phrases. All of these issues I described in the preceding discussion transfer to the moral case: the truth of a moral principle cannot be explained in ordinary quantificational terms or a statistical regularity. This then leaves us with difficult questions about what makes a moral generic true and so on.⁵

My view is that it is in fact very desirable to understand moral principles in ways that make them seem comparatively elusive; I think that this is more revealing of their heterogeneous character compared to other approaches. Depending on what stance we take on the semantic interpretation of Gen, there may be different ways to interpret generic moral principles. This leaves open an interesting further question to investigate which approaches to genericity are more or less suitable in the moral case.

One promising route is to regard generic moral principles as capturing our basic knowledge of right and wrong. This path is naturally suggested by the understanding of Gen as a default mechanism of generalization (Leslie 2008). Relatedly, Lerner and Leslie (2013) suggest, on the basis of various psychological and linguistic tests, that moral generics express characteristic properties. If this is right, then the connection between, say, acts of stealing and being wrong is one that is causal and explanatory. Moral generics, construed in this way, have exceptions, but there is a difference between exceptions and counterexamples. The familiar moral generalizations admit exceptions, but it is possible to argue that a characteristic generic does not hold. Lerner and Leslie provide the example of the false moral generic: “Sexual acts outside of religiously consecrated marriage are wrong” (2013: 392). The cases that do not

⁵ I am grateful to an anonymous reviewer for raising this important objection.

conform to the generic are not exceptions but rather can be used as evidence that the characteristic generic is false because it does not provide the best explanation of particular cases.

In my presentation of the generic view, I emphasized the heterogeneity in the strength of the generalization that appears to be expressed. If we combine this broad variability with the idea that generics express characteristic properties, we can then consider the intriguing implication that there are different kinds of characteristic properties expressed by moral generics. The kind of characteristic property expressed by a generic such as ‘Eating animal products is wrong’ may be very different from the kind of characteristic property expressed by ‘Murder is wrong’ in part because many would override the former and very few could imagine overriding the latter. If this is right, then it raises an interesting potential for a further classification of moral principles via a diverse range of characteristic properties.

One important feature of generics, emphasized by Nickel (2016), among others, is that they are important in the context of explanation. The reason that Tigger is striped can be explained by the fact that Tigger is a tiger and tigers are striped. Taking on the generic view of moral principles can provide us with a novel perspective on how principles have an explanatory role. The idea that exception-granting generalizations can be explanatory in ethics is not without precedent. Väyrynen (2009) argues that hedged moral principles can be genuinely explanatory despite their exception-granting character (cf. Lance and Little 2008; Leibowitz 2011; Väyrynen 2018). But even the idea that universal moral principles are explanatory is controversial. Berker (2019) argues that such moral principles do not themselves do the work of making it the case that specific moral facts are the way they are; principles are instead mere summaries of grounding relations. When we conceive of moral principles as generics, we can hold that moral generics explain their instances in the sense that they are capable of telling us why a particular action

possesses the moral property it has. But generic generalizations are not grounded in their instances. Generic moral principles can explain why an act of stealing is wrong, even though the explanation can be defeated. Explanation in this sense does not require a prevalence of conforming instances subsumed under a generalization nor does it require that it is probable that a particular instance is a conforming one.⁶

The next point is that the generic view promises to give us a fresh perspective on moral knowledge. There are robust empirical findings that point to the fundamental role that generalizations play in our cognitive development. Recent work in psychology has shown, *inter alia*, that generics underpin the theories about both the natural and social world that children develop by shaping a propensity to make causal inferences and support our schemes of classification (Cimpian et al. 2010). Generics are a valuable tool for acquiring and imparting information to developing minds: children master generics before they use explicitly quantified language and generics are ubiquitous in maternal speech (Rhodes et al. 2018). If we accept these empirical findings, then it would seem that makes the most sense of how it is that we acquire general moral knowledge at a young age. If moral generalizations were quantificational, then this would conflict with the idea that generic generalizations are learned relatively early compared to quantified generalizations. Therefore, if do not understand moral principles as generics, then we would be left in a weaker position to explain the way moral knowledge is acquired at a young age. This would also leave philosophical discussions in moral psychology less continuous with empirical findings in developmental psychology concerning the acquisition of moral knowledge.

⁶ I owe these points to an anonymous reviewer.

If generics can be used to teach morality to children, then this reinforces the potential for a less skeptical view on which there are contexts where it is appropriate to optimistically infer from moral generic generalizations. There is no factor, given the components of the generic view presented here, that rules out the idea that it can be useful to think of typical patterns of moral reasoning as acceptable in the context of moral education. If anything, it would seem a virtue of the generic view that it can accommodate the idea that there are contexts where we should be pessimistic about the prospects of moral reasoning and contexts where we should feel optimistic. Given the broad semantic flexibility exhibited by generics, it may appear unsurprising if we get uncontroversial contexts where the use of ordinary moral principles in reasoning are acceptable and warranted (in particular, not subject to issues raised by the argument from Non-Numerity). A consequence is that the skepticism associated with the generic view is therefore not a global skepticism concerning the role of moral principles in reasoning.

The generic view gives us powerful tools to understand the moral domain: it can deliver the idea that principles are explanatory and provides a novel perspective on moral knowledge. It also creates a helpful bridge to research in developmental psychology, especially in the study of how children acquire moral concepts. All of these points raise intriguing questions for further investigation. There are further matters to consider regarding the potential implications for ordinary moral reasoning, particularly the argument from Non-Numerity. This is pursued in the next section.

4 The non-monotonic view

One recent advance in logically-informed moral philosophy defended by Horta (2007, 2012) is the idea that moral principles specify “the defaults that underlie our reasoning” (Horta

2012: 154). Defaults have been well-studied in computer science and artificial intelligence ever since the development of default logic in Reiter (1980). Thomas (2007, 2011) has independently defended the view that the best way to conceive of moral reasoning is in non-monotonic terms. My discussion here will focus on details of Horty's view since it has more conflicts with the generic view.

The conception of moral principles as defaults is a competitor to related views which involve understanding moral principles as involving *ceteris paribus* provisos (and related ideas), such as those defended by Pietroski (1993) and Holton (2002). Defaults have a number of properties that make them especially valuable in characterizing human reasoning. This view competes with these other positions through greater emphasis on how exceptions work in the system: the core idea is that a moral principle such as 'Lying is wrong' allows us to establish, by default, that an act of lying is wrong. However, this is so "unless certain complicating factors interfere" (Horty 2012: 154). Defaults can be overridden by other defaults, provided there is some factor in force that prioritizes other defaults. And defaults can be overridden according to a number of different factors, all of which can be accounted for within a system of defaults with priority relations among them.

I argue that the generic view supports a view of moral reasoning that is neither deductive nor defeasible. I will add further support for the argument from Non-Numerity by engaging with the competing view that moral reasoning is a form of non-monotonic reasoning. One important factor to consider is that the default perspective on moral principles is consistent with Moral Non-Numerity. That there is a default which says we should conclude that an action is wrong if it involves stealing does not provide us with any information about the prevalence of cases of

stealing that ultimately turn out to be wrong. If this is right, then the proponent of the default approach should accept Moral Non-Numerity.

However, if Moral Non-Numerity is accepted by the default theorist, this leads to some issues from an epistemological perspective. The proponent of the default view has the burden of justifying why it is that moral principles are associated with the default entitlement or presumption that arbitrary instances confirm to the default principle. I will put forward two points against the default theory on this matter.

The first point is that there is not enough support for the idea that we should feel entitled to reason with moral principles by default. Suppose that the proponent of the default theory argued that the entitlement or presumption to infer according to moral principles is based on the idea that our place in the moral landscape sufficiently meets certain background expectations of normality in the sense that it should be considered a safe bet, so to speak, at least in our world, that a particular act of stealing counts as wrong, or that breaking a promise counts as wrong. This idea is very controversial, and there has been significant pushback against it, in different ways, due to Dancy (2004), Väyrynen (2004), and McKeever and Ridge (2007). One point that Dancy emphasizes is that if there is some metaphysical account of the world that ‘backs up’ the default entitlement or presumption on the default view, then this would have to seem to be a strange coincidence, if anything. The idea that moral thought and judgment depends on moral principles cannot be based on “a considerable preponderance of normal cases over abnormal ones” (Dancy 2004: 76). A factor to consider on top of this is that even if one were to grant a certain configuration of normal cases over abnormal cases as the basis for default entitlement or presumption, then it is still doubtful that this could be enough to justify a generalist epistemology

of our moral beliefs about the moral features of particular cases. Such a configuration would not show that moral principles which specify defaults are sufficient enough to justify beliefs about particular cases: it still does not show that a conception of moral principles consistent with Moral Non-Numerity would provide, on its own, the positive epistemic appraisal we are after because it would appear there are strong doubts that such positive support follows from the metaphysical story.

The second point is that there is a convincing alternative explanation for why we might feel a presumption in favor of reasoning with moral principles by default. It might be thought that one point in favor of the default theory is that it seems to do well from a descriptive point of view. It could be argued, for instance, that the appeal to defaults might do well at adequately characterizing the descriptive features of our psychologies by making good on the idea that we are prone to make such strong inferences on the basis of moral principles. For instance, as Cimpian et al. (2010) point out, there is intriguing evidence that people estimate prevalence very highly when presented with generics. However, there is an empirical point that can be made in response which allows us to look at the psychological evidence from a different point of view. Recent work by Leslie, Khemlani, and Glucksberg (2011) gives support for the idea that we are prone to a psychological effect whereby there is a tendency to generalize from the truth of a generic to the truth of the corresponding universal statement. Lerner and Leslie (2013) build on this and argue that it might appear that people accept rather strong moral generalizations, however the source of this is actually an overgeneralization based on the acceptance of moral generics. They observe that it is generally agreeable that universally quantified moral generalizations are false, because of familiar problems having to do with moral dilemmas. However, at the same time, many subjects report an intuitive attraction to such universally

quantified generalizations. Lerner and Leslie (2013) argue for the existence of what they call a ‘moral overgeneralization effect’: there is a psychological tendency whereby people end up endorsing the universally quantified moral generalizations because they mistake them for the true generic moral generalizations. It seems likely that the same point should apply in the case of accepting default principles: a reasonable and satisfying explanation of the tendency to reason with default principles could be that it is due to a similar sort of overgeneralization effect.

These arguments perform a double duty: they offer reasons to reconsider an alternative, competing conception of moral principles while further illustrating the ideas behind Moral Non-Numerity. The idea that moral principles do not wear any information on their sleeves concerning how many instances conform to the moral generalizations undermines the idea that there might be some metaphysical basis for the idea that there is a default entitlement to using moral principles in reasoning. And these points again emphasize that Moral Non-Numerity implies a claim with a distinctively normative flavor: because moral principles are independent of information about prevalence, they cannot generate sufficient epistemic support in moral reasoning. We therefore need to pay more attention to the question of how we should reason as opposed to how we in fact reason.

None of the foregoing discussion is meant to imply that a nonmonotonic view is in principle incompatible with particularist views of moral reasoning. Thomas (2011), for instance, argues that the nonmonotonicity of moral reasoning leads to particularism. The generic view is compatible with, and potentially enhances, such a position. We can derive inferences from moral generic generalizations, but only very cautiously. Examining the nature of reasoning with generics reveals that there are further complexities to consider in the moral case and that we

should be careful in thinking that patterns of reasoning which appear good should actually be regarded as good or adequate.

5 Concluding remarks

A notable advantage of the generic view is that it simultaneously allows for pessimism and optimism about the role and status of moral principles in our lives. It provides a new perspective on the nature of moral principles on which principles are not apt for determining the moral status of particular actions while they may be apt, and even fundamental, to our acquisition of moral knowledge. A natural consequence of the view is variation amongst moral principles, with some regularly warranting exceptions, and some appearing arguably exceptionless. There are several other further implications to explore, especially with respect to moral epistemology and the very nature of what it means to be guided by principles of varying moral statuses, whether obligatory, recommended, or permissive.

The status of the generic view in this paper does not rest on any particular semantic theory for generics, so the questions about the formal details would need to be pursued in further research. The generic view raises several further questions regarding how exactly we should interpret moral principles in terms of their logical form and their meaning. One natural consequence of this position would seem to be that many moral statements which feature overt modals should also involve underlying genericity. For instance, Saint-Croix and Thomason (2019) have recently discussed the idea that some ‘ought’-sentences are also generic. On their proposal, the underlying logical form of some ‘ought’-sentences involves two modals: a generic modal and a deontic modal. The defeasibility of such oughts is then captured by representing the generic modal as having wider scope.

Acknowledgements

For helpful discussion and comments, I thank Farbod Akhlaghi, Derek Ball, Joseph Bowen, Mark Bowker, Josh Dever, Miguel Egler, Mirela Fus, Eline Gerritsen, Olivier Lemeire, Bernhard Nickel, Justin Snedegar, Rachel Sterken, Alan Thomas, Clotilde Torregrossa, Caroline Touborg, and audiences at the University of Oslo, University of Sherbrooke, University of St Andrews, and University of York. I am also particularly grateful to anonymous reviewers for this paper.

References

- Almotahari, Mahrad (2022). Weak generics. *Analysis* 82(3), 405–409.
- Asher, Nicholas and Michael Morreau (1995). What some generic sentences mean. In *The Generic Book*, pp. 300–339. Chicago: University of Chicago Press.
- Berker, Selim (2019). The explanatory ambitions of moral principles. *Noûs* 53(4), 904–936.
- Carlson, Greg (1977). *Reference to Kinds in English*. Ph. D. thesis, University of Massachusetts at Amherst.
- Carlson, Greg (1995). Truth conditions of generic sentences: Two contrasting views. In Carlson and Pelletier (Eds.) *The Generic Book*, pp. 224–237. Chicago: University of Chicago Press.
- Cimpian, Andrei, Amanda C. Brandone, and Susan A. Gelman (2010). Generic statements require little evidence for acceptance but have powerful implications. *Cognitive Science* 34(8), 1452–1482.
- Collins, John (2018). Genericity sans gen. *Mind and Language* 33(1), 34–64.
- Dancy, Jonathan (1993). *Moral Reasons*. Oxford: Blackwell.
- Dancy, Jonathan (2004). *Ethics without Principles*. Oxford: Oxford University Press.
- Demos, Raphael (1958). Is moral reasoning deductive? *Journal of Philosophy* 55(4), 153–159.
- Enoch, David (2011). *Taking Morality Seriously: A Defense of Robust Realism*. Oxford: Oxford University Press.
- Enoch, David (2019). How principles ground. In R. Shafer-Landau (Ed.), *Oxford Studies in Metaethics*, Volume 14. Oxford: Oxford University Press.
- Hare, R. M. (1952). *The Language of Morals*. Oxford: Oxford University Press.
- Hare, R. M. (1963). *Freedom and Reason*. Oxford: Oxford University Press.

- Holton, Richard (2002). Principles and particularisms. *Aristotelian Society Supplementary Volume* 76(1), 191–209.
- Horty, John (2007). Reasons as defaults. *Philosophers' Imprint* 7(3), 1–28.
- Horty, John (2012). *Reasons as Defaults*. Oxford: Oxford University Press.
- Khemlani, Sangeet, Sarah-Jane Leslie, and Sam Glucksberg (2009). Generics, prevalence, and default inferences. *Proceedings of the Cognitive Science Society*, 443–448.
- Krifka, Manfred, Francis Pelletier, Greg Carlson, Alice ter Meulen, Godehard Link, and Gennaro Chierchia (1995). Introduction. In G. Carlson and F. Pelletier (Eds.), *The Generic Book*, pp. 1–124. Chicago: University of Chicago Press.
- Lance, Mark and Margaret Little (2004). Defeasibility and the normative grasp of context. *Erkenntnis* 61(2-3), 435–455.
- Lance, Mark and Margaret Little (2006). Defending moral particularism. In J. Dreier (Ed.), *Contemporary Debates in Moral Theory*, pp. 305–321. Oxford: Blackwell.
- Lance, Mark and Margaret Little (2008). From particularism to defeasibility in ethics. In V. Strahovnik, M. Potrc, and M. Lance (Eds.) *Challenging Moral Particularism*, pp. 53–74. New York: Routledge.
- Leibowitz, Uri D. (2009). A defense of a particularist research program. *Ethical Theory and Moral Practice* 12(2), 181–199.
- Leibowitz, Uri D. (2011). Scientific explanation and moral explanation. *Noûs* 45(3), 472–503.
- Lerner, Adam and Sarah-Jane Leslie (2013). Generics, generalism, and reflective equilibrium: Implications for moral theorizing from the study of language. *Philosophical Perspectives* 27(1), 366–403.
- Leslie, Sarah-Jane (2007). Generics and the structure of the mind. *Philosophical Perspectives* 21(1), 375–403.
- Leslie, Sarah-Jane (2008). Generics: Cognition and acquisition. *Philosophical Review* 117(1), 1–47.
- Leslie, Sarah-Jane (2015). Generics oversimplified. *Noûs* 49(1), 28–54.
- Leslie, Sarah-Jane, Sangeet Khemlani, and Sam Glucksberg (2011). Do all ducks lay eggs? The generic overgeneralization effect. *Journal of Memory and Language* 65(1), 15–31.

Leslie, Sarah-Jane and Adam Lerner (2022). Generic generalizations. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Fall 2022 ed.). Metaphysics Research Lab, Stanford University.

McKeever, Sean and Michael Ridge (2005). What does holism have to do with moral particularism? *Ratio* 18(1), 93–103.

McKeever, Sean and Michael Ridge (2006). *Principled Ethics: Generalism as a Regulative Ideal*. Oxford: Oxford University Press.

McKeever, Sean and Michael Ridge (2007). Turning on default reasons. *Journal of Moral Philosophy* 4(1), 55–76.

Nguyen, Anthony (2020). The radical account of bare plural generics. *Philosophical Studies* 177(5), 1303–1331.

Nickel, Bernhard (2016). *Between Logic and the World*. New York: Oxford University Press.

Pelletier, Francis J. and Nicholas Asher (1997). Generics and defaults. In J. van Bentham and A. ter Meulen (Eds.), *Handbook of Logic and Language*, pp. 1125–1177. North Holland: Amsterdam.

Pietroski, Paul (1993). Prima facie obligations, ceteris paribus laws in moral theory. *Ethics* 103(3), 489–515.

Prasada, Sandeep, Sangeet Khemlani, Sarah-Jane Leslie, and Sam Glucksberg (2013). Conceptual distinctions amongst generics. *Cognition* 126(3), 405–422.

Reiter, Ray (1980). A logic for default reasoning. *Artificial Intelligence* 13(1-2), 81–132.

Rhodes, Marjorie, Sarah-Jane Leslie, Lydia Bianchi, and Lisa Chalik (2018). The role of generic language in the early development of social categorization. *Child Development* 89(1), 148–155.

Robinson, Luke (2008). Moral principles are not moral laws. *Journal of Ethics and Social Philosophy* 2(3), 1–22.

Robinson, Luke (2011). Moral principles as moral dispositions. *Philosophical Studies* 156(2), 289–309.

Ross, W. D. (1930). *The Right and the Good*. Oxford: Oxford University Press.

Saint-Croix, Catharine and Richmond Thomason (2019). Chisholm’s paradox and conditional oughts. *Journal of Logic and Computation* 29(3), 369–386.

Schiller, Henry (forthcoming). Genericity and inductive inference. *Philosophy of Science*.

- Smith, Holly (1988). Making moral decisions. *Noûs* 22(1), 89–108.
- Stangl, Rebecca (2006). Particularism and the point of moral principles. *Ethical Theory and Moral Practice* 9(2), 201–229.
- Sterken, Rachel Katharine (2015). Generics in context. *Philosophers' Imprint* 15(21), 1–30.
- Sterken, Rachel Katharine (2016). Generics, covert structure and logical form. *Mind and Language* 31(5), 503–529.
- Sterken, Rachel Katharine (2017). The meaning of generics. *Philosophy Compass* 12(8), e12431.
- Sunstein, Cass (2005). Moral heuristics. *Behavioral and Brain Sciences* 28(4), 531–541.
- Thomas, Alan (2007). Practical reasoning and normative relevance: A reply to McKeever and Ridge. *Journal of Moral Philosophy* 4(1), 77–84.
- Thomas, Alan (2011). Another particularism: Reasons, status, and defaults. *Ethical Theory and Moral Practice* 14(2), 151–167.
- Väyrynen, Pekka (2004). Particularism and default reasons. *Ethical Theory and Moral Practice* 7(1), 53–79.
- Väyrynen, Pekka (2006). Moral generalism: Enjoy in moderation. *Ethics* 116(4), 707–741.
- Väyrynen, Pekka (2009). A theory of hedged moral principles. In R. Shafer-Landau (Ed.), *Oxford Studies in Metaethics*, Volume 4. Oxford: Oxford University Press.
- Väyrynen, Pekka (2013). Grounding and normative explanation. *Proceedings of the Aristotelian Society Supplementary Volume* 87, 155–178.
- Väyrynen, Pekka (2018). Reasons and moral principles. In D. Star (Ed.), *The Oxford Handbook of Reasons and Normativity*, pp. 839–864. Oxford: Oxford University Press.