

# The Reality of Primitive Norms

Hannah Ginsborg

Forthcoming in *Normative Realism*, edited by Paul Boghossian and Christopher Peacocke (Oxford: Oxford University Press)

Preprint (version of August 2023)

Abstract: Debates about normative realism typically take for granted a contrast between claims that do, and those that do not, state objective facts. Fact-stating claims, for example empirical or mathematical claims, are seen as a privileged class; the question about realism is whether normative claims, for example ethical claims or claims about reasons, can be assimilated to that class. Drawing on examples from Wittgenstein’s rule-following considerations and from developmental psychology, I introduce a kind of normativity—“primitive” normativity—which, I argue, complicates this dialectical framework. The capacity to make and understand fact-stating claims depends on our recognition of primitive norms which govern our behavior in continuing patterns, sorting objects and learning words. I argue that these norms cannot be explained in terms of the more widely-discussed norms of ethics, aesthetics, rational action and belief, or game-playing, and that they are not amenable to standard anti-realist treatment. They are, in an intuitive sense, real; and yet claims about these norms cannot be assimilated to fact-stating claims, since the possibility of fact-stating claims—and perhaps even the possibility of facts themselves—depends on the existence of these norms and on our recognition of them. This yields an argument against familiar forms of normative anti-realism in domains such as ethics. The normative anti-realist has to accept the reality of primitive norms in order to make sense of the privileged class of fact-stating claims presupposed by the realist/anti-realist debate; but this precludes her citing the normativity of, say, ethical claims as a ground for denying realism about ethics.

## I. Introduction

Debates about normative realism or realism about values are often framed in terms of the question whether there are properties or facts corresponding to normative or evaluative expressions and statements. These debates typically take for granted that we can draw a distinction between fact-stating discourse and discourse which does not state facts. They also typically take for granted that, within the facts stated, we can distinguish those that are objective from those which are merely subjective, in the sense of being grounded—in some philosophically distinctive way—on the beliefs, attitudes, or feelings of the people making the statements or of human beings more generally. Although there is disagreement

about which kinds of statements qualify as statements of objective fact, philosophers usually include claims of physical science and mathematics, for example that protons have positive charge and that  $2+2=4$ , as well as claims about particular objects or events that use the vocabulary of science and mathematics, for example that this body has a mass of one kilogram. At least in the context of debates about normative realism, as opposed to realism about say, colours, or the past, or social properties, they are also likely to include descriptive statements using ordinary vocabulary: for example that grass is green, that this particular object is a chair, that Anscombe was born in 1919, and that the Eiffel Tower is in Paris. In the context of debates about normative realism, statements like these—which I'll call paradigmatically fact-stating—constitute a privileged class; they exemplify a standard against which other statements can be measured and with respect to which they may be seen as falling short.<sup>1</sup>

On the face of it, normative statements—for example about behavior being morally right or wrong—seem to have a lot of important features in common with paradigmatically fact-stating statements. They have a similar grammatical structure, their contents allow of being embedded in more complex statements, and they are often accepted as true. But anti-realists about various areas of normative discourse deny that the normative statements in those respective areas can be assimilated to the privileged class. Typically they are motivated by metaphysical or epistemological concerns about the kinds of facts to which we seem to be committed if we take these statements at face value. The classic statement of these worries is J.L. Mackie's "argument from queerness."<sup>2</sup> Commitment to objective moral facts seems to commit us to the existence of objective moral properties like rightness or goodness, and this is problematic both because of the intrinsically motivating character of these properties and because there seems to be no accounting for how they supervene on non-moral properties. Relatedly, there seems to be no explanation, at least within an empiricist framework, of how we come to form ideas of these properties or to recognize their presence in individual cases. Knowledge of moral facts would seem to require a special faculty, distinct from the capacities which enable us to know the kinds of facts stated by

---

<sup>1</sup> This is implicit, for example, in David Enoch's claim that normative truths are "just as respectable as empirical or mathematical truths" (2011, 1).

<sup>2</sup> Mackie 1977, ch. 1.

the paradigmatic fact-stating claims, not only for detecting the moral facts but also for understanding their relationship to the non-moral facts.<sup>3</sup> More recently, further epistemological worries have been raised in the form of the “evolutionary debunking argument”: assuming that our capacities for normative judgment have emerged as a result of natural selection and hence for the purpose (loosely speaking) of advancing our evolutionary fitness, it would seem to be a sheer, and intuitively unacceptable, coincidence that they also put us in a position to grasp objective truths about the world.<sup>4</sup> But simply raising these worries is not sufficient to establish moral, or more generally normative, anti-realism; the anti-realist also has to explain away the apparent parallels between normative and non-normative statements, or in other words to justify rejecting the assimilation of normative statements to the paradigmatic fact-stating claims.

Three kinds of approach have emerged for defending anti-realism about normative statements. The first approach, often labelled non-cognitivism, is to deny that normative utterances—for example, utterances like “stealing is wrong”—aim to express beliefs or report facts. Rather, they express non-attitudes, prescriptions, or plans, for example disapproval of stealing, the command not to steal, or plans, under various hypothetical circumstances, to refrain from stealing. The second approach, represented by so-called error theories, allows that normative utterances express beliefs but denies that these beliefs are ever true or that the facts they report ever obtain. According to an error theory of morality, we are simply mistaken when we claim that stealing is wrong or make any other moral claim: claims that some action or type of action is right or wrong are, invariably, false. The third approach, which I will label subjectivism,<sup>5</sup> allows that normative utterances express true beliefs and, in doing so, report facts, but

---

<sup>3</sup> These worries can be addressed by identifying seemingly normative facts and properties with natural facts and properties: for example by identifying the moral goodness of an action with its potential to maximize pleasure and/or minimize pain. This approach, famously challenged by G.E. Moore’s “Open Question” argument, is usually described as a form of normative realism, specifically naturalistic realism. As I see it, naturalistic realism is more closely allied to anti-realist approaches than to non-naturalistic approaches to normative discourse, since, like these other approaches, it seeks to undermine the idea that the world, broadly construed, includes irreducibly normative properties. So, unless otherwise specified, when I speak of normative realism in what follows, I have in mind the more traditional, non-naturalistic variety.

<sup>4</sup> A very influential example is Street 2006.

<sup>5</sup> Joyce 2021 proposes “non-objectivism” rather than “subjectivism” as a name for this kind of position, but the traditional “subjectivism” has a somewhat more informative ring. Sharon Street’s “Humean constructivism” (see e.g. Street 2008) is an example of the kind of position I have in mind. Although J.L. Mackie is usually designated as an error theorist, I follow Berker 2019 in reading him, also, as a subjectivist.

denies that the facts they report are facts: there are normative facts of a kind, but these are, in a certain sense, mind-dependent; they hold in virtue of peoples' beliefs and attitudes about whatever the beliefs are about.<sup>6</sup> The utterance "stealing is wrong" does report a fact to the effect that stealing is wrong, but that fact is grounded in non-normative facts about psychological states relevant to the phenomenon of stealing, for example facts about the attitudes that the person, or a suitably chosen, group of people have about stealing.

Historically, most discussion about normative realism has focussed on norms of morality. But there are other varieties of normative or evaluative discourse. One important variety is that of reasons for action. To say that someone has a reason for performing a certain action is to make a normative claim: it is to say that, in the absence of any competing reasons or other considerations telling against the action, the person should or ought to perform that action.<sup>7</sup> Another is that of reasons for belief: to say that someone has reason to believe a proposition, say, in the form of evidence for its truth, is to say that, in the absence of countervailing evidence, she should or ought to believe it. A third variety is aesthetic discourse: in saying that a thing is beautiful or has some other kind of aesthetic property I am making a normative claim about what kind of affective response perceivers of the thing should have to it. These kinds of normativity are all subject to the same dialectic which has arisen for moral normativity: in all of them we can raise metaphysical and epistemological doubts about how we can make sense of the norms and our capacity to know what they are, and we can try to avoid those doubts by adopting non-cognitivist, error, or non-objectivist accounts of the corresponding claims. There are also further kinds of normative discourse which are more closely tied to specific cultural practices, and which do not on the face of it seem to raise the same problems. The putative norms which figure in these kinds of discourse include, for

---

<sup>6</sup> It is notoriously difficult to characterize the precise sense of mind-dependence at issue; the gloss I propose is approximate, but I think it will suffice for the purposes of this paper,

<sup>7</sup> T.M. Scanlon (2014) is explicit that claims about reasons (for belief as well as action, see at 35) are normative, but he would not probably not put the point the way I just did, because he is inclined to characterize the normative domain in terms of reasons (34) rather than (as I do) in terms of the more general notion of *ought*. As will become clear in what follows, I think that the notion of the normative is broader and more basic than the notion of what one has reason to do, although I am inclined to agree with Scanlon that "concepts...such as good, value, and moral right and wrong, are best understood in terms of reasons" (34).

example, norms of etiquette, legal and political norms, and norms for speech acts like assertion, as well as the cultural or conventional norms governing more specialized activities like speaking English, composing music in the classical style, book-binding, fashion design, and games like chess. If there is such a thing as the normativity of meaning it would likely fall into this general category. To the extent that we think of these more localized norms as *sui generis* they do not seem to raise the same problems that arise for morality, the rationality of conduct and belief, or aesthetics. However some philosophers might be inclined to think of these as falling under the head of norms for rational action, in which case they are subject to the same dialectic that I described for moral norms.

In this paper I want to draw attention to a kind of normativity which I believe is different from any of those I have mentioned. This kind of normativity is worth discussing in the context of discussions of normative realism because, if it is accepted as a distinct kind of normativity, it complicates the dialectic I have described, and it does so in an interesting way. The normativity I have in mind is what I have elsewhere called “primitive normativity,” and it is manifested in claims—which might or might not be verbally expressed—about the appropriateness or correctness of one’s behavior under particular circumstances. That brief characterization might give the impression that they are claims either about what one has reason to do, or — assuming this is different — about what is required by cultural or conventional rules like those which arguably figure in etiquette, legal systems, games and so on. But I will argue that these claims are different, and, specifically that they involve the exercise of a more basic capacity, one that is presupposed by our capacity both to recognize reasons and to recognize conventional norms like the rules of a game. What is interesting about these normative claims, for the purposes of the normative realism debate, is that our capacity to make them is also a prior condition of our being able to make and understand the paradigmatically fact-stating claims that are taken as a point of reference in debates about normative realism. If we cannot recognize the primitive norms, I will argue, we cannot entertain propositions to the effect that grass is green or that  $2+2=4$ . This means that if we are not willing to grant, in some intuitive sense, the reality of these norms then we have to give up the idea that

we can know, or even believe, that grass is green or that  $2+2=4$ , which undermines the dialectical framework within which questions about normative realism are typically raised.

The remainder of the paper is structured as follows. In Section II, I explain what primitive norms are, and argue that we need to recognize them in order to grasp concepts or understand linguistic expressions. In Section III, I argue that these norms cannot be assimilated to any of the kinds of norms that are typically invoked in discussions of normative realism. In Section IV, I argue that primitive norms are not amenable to standard anti-realist treatments and that they qualify, in an intuitive sense, as objectively real; I suggest in conclusion that this yields an argument against more familiar forms of normative anti-realism, such as anti-realism about morality and reasons.

## II. Primitive norms as a condition of meaning and understanding

I want to explain what I mean by primitive norms by invoking three different examples, or rather groups of examples. One is from philosophical fiction, and the other two are from developmental psychology. The group of examples from philosophical fiction is offered by Wittgenstein, most famously in his case of the aberrant pupil who, after first being taught how to write numerals and then to develop the series of natural numbers, is now learning to continue the series  $0,2,4,6,8,\dots$ <sup>8</sup> We think he knows how to do it, but when he gets to 1000, he writes 1004, 1008, 1012 and expresses surprise when we tell him that what he has written is wrong. He thinks that what 1004 is what he *should* write or the *correct* thing to write once he has got to 1000. This brings into relief our own attitude, or that of a normal child who has received the same training as Wittgenstein's pupil, to writing 1002. We think, that in circumstances like the aberrant pupil's — having been shown examples of the series of natural numbers, then having been shown examples of the series of even numbers, and having written out the series up to 1000 — the correct thing to write, when told to “go on,” is 1002, not 1004. Wittgenstein presents this example in

---

<sup>8</sup> *Philosophical Investigations* §185; it is important to keep in mind, though, that the pupil is first introduced at §143. I discuss Wittgenstein's examples further in Ginsborg 2022.

connection with numerals, which has led readers to think that it is about arithmetic; but in fact it has nothing specifically to do with arithmetic or even to do with numbers. It is part of a family of examples illustrating what Wittgenstein calls “knowing how to go on”: at a first approximation, knowing how to continue a pattern of behavior. This includes, for example the recognition that after drawing a pattern which consists of several dashes alternating with pairs of dots, the correct thing to draw after a dash is another pair of dots.<sup>9</sup> An even simpler example is that of knowing how to continue the pattern of behavior an adult might get into with a small child, where the adult claps her hands and the child imitates her: Wittgenstein suggests indirectly that the recognition of normativity is involved here as well: “If A claps his hands, B is always supposed to [sollen] do it too.”<sup>10</sup> The toddler who claps his hands when his mother claps hers shows that he “knows how to go on” in this kind of case too, where this involves recognizing his own clapping as the right thing to do in response to his mother’s clapping.

The second group of examples comes from a series of studies that were carried out, mostly in the 1980s and 1990s, about a kind of behavior in small children called “spontaneous sorting.”<sup>11</sup> It emerges from these studies that one- and two-year-olds who are presented with randomly arranged objects of two clearly demarcated kinds will spontaneously sort the objects by touching them in sequence or moving them around so that objects of the same kind are grouped together. The typical study involves giving the child eight objects comprised of two groups of four, where the objects in each group of four are qualitatively identical; for example they are given four yellow balls and four grey cubes, or four blue dolls and four green boats. But the results are the same even if the objects within each kind are different from one other, say four pencils of different lengths and colours, and four rings of different shapes and sizes. The children do not have to be given any special instructions to engage in the sorting activity; it is enough that they are presented with the objects and allowed to play with them. Starting around one year of age, children group together three of one kind of object, leaving the others untouched; starting around 18

---

<sup>9</sup> *Remarks on the Foundations of Mathematics (RFM) VI-17*

<sup>10</sup> *RFM VI-17*

<sup>11</sup> For references both to these studies, and to the false labelling studies described in the next paragraph, see Ginsborg 2021, where both sets of studies are described in more detail.

months they move objects of both kinds around to form two distinct same-kind groups, sometimes including all eight of the objects (“exhaustive sorting”); by the time they are two, most children engage in exhaustive sorting. Children around this age also correct classificatory mismatches: if the experimenter gives them three dolls and a boat, and three boats and a doll, they switch around the objects that don’t belong with the others. Significantly for our purposes, the children themselves seem to recognize this activity as having a normative dimension: they seem to recognize, when they put a doll with the other dolls, rather than with the boats, that this is the right thing to do. Starting around two and a half years of age, when normative language first enters their vocabulary, they use normative words to describe what they are doing. The experimenters report children at two and a half years of age saying things like “no belongs this way” or “no, they’re not on properly” when they are shown the mismatched objects. It seems clear that, when they put one doll with the others, they recognize themselves to be putting it where it belongs, or where it ought to go. And even if the younger children do not use this normative language, their behavior is similar enough to that of the children who do use the normative expressions that it seems reasonable to ascribe to them the same attitudes.

The third group of examples comes from studies of children’s response to “false labelling,” where small children are shown adults calling familiar objects by the wrong names. Starting around 16 months of age, when children are at the “one-word” stage — they can respond correctly to, and sometimes produce, words like “ball,” “dog,” or “shoe” — they react very strongly when they hear an adult mis-label an object using a word that is familiar to them. The experimenters in one study describe the children as showing “distress” when they hear the mis-labeling: they say “no,” they wave their hands, they shake their heads. (The children were sitting on their parents’ laps, and the parents were wearing sleep masks so that they could not inadvertently cue the child; several of the children turned around and tried to pull the sleep masks off their parent’s face.) The experimenters report that the children would try to produce the correct name, and that, if they were unable to do so, they would try to correct the adult in other ways: for example, when the experimenter said “shoe” while the experimenter and the child were looking at a cat, some of the children would point to their own shoe. Here, it seems clear that the children recognize the



mislabeled as wrong. And this recognition is related to the recognition of the mismatches in the previous example. The experimenter who says “shoe” while looking at the cat is in effect sorting the cat incorrectly: putting the cat with the shoes, that is with the objects that the child has previously connected with the word “shoe.” When the child points to her own shoe, she is in effect attempting to correct the mismatch by replacing the cat with the shoe: it is her shoe, not the cat, which goes with, or ought to be put with, the other shoes.

By “primitive norms,” I mean to pick out the norms that the children recognize (or in the case of Wittgenstein’s aberrant pupil, fail to recognize) in these examples. We can describe these norms by saying, for example, that, in the situations described, 1002 is the correct thing to write after 1000, or that one should clap one’s own hands after the other person claps theirs, or that the doll ought to go with the other dolls rather than the boats, or that it is the shoe rather than the cat that belongs with the other shoes and so should be labelled with the word “shoe.” But we need not be able to formulate or understand these descriptions in order to recognize the norms, and this is part of the reason why I have introduced them using examples of children who are clearly too young to articulate them. A two-year-old who is holding a doll and deciding where to put it can recognize that she should put it with the dolls rather than the boats before being able to use words like “doll” or “boat” or even to explicitly entertain thoughts like *this goes here*.<sup>12</sup> Moreover, the norms are specific to the particular context of the behavior which they govern. We cannot specify them by saying that, in general, 1002 is the correct thing to write after 1000, even in a series that begins with “0,2,4,6,8” There could be a context in which the correct thing to write is 1004, for example if the pupil has been told in advance to continue the series up to 1000 and then to add 4, or indeed if the pupil himself has formed the intention to provoke us by doing something different

---

<sup>12</sup> If, like Boghossian (2008, 475) we think of “norm” as short for “normative proposition,” so that grasping a norm would be grasping the truth of a proposition like *At the beginning of the game, White must make the first move*, then this proposal will seem incoherent. But I understand “norm” in a broader sense which allows us to speak of the existence of a norm in any context where we would also speak of what someone ought to do (or, for that matter, of how something ought to be, or what ought to happen, although in this paper I will not be concerned with these other contexts). Any recognition of what I ought to do, on this usage, whether or not it involves grasp of a proposition, is recognition of a norm. Here it might help to keep in mind the etymology of “norm” from the Latin *norma*, meaning carpenter’s square: in using a carpenter’s square to help you draw a right angle, you implicitly accept it as determining how you ought to move your pencil, or the correct placement of the lines, but this acceptance need not involve grasp of a proposition.

from what we expect. Similarly, it is not always the case that, when given four dolls and four boats, one should put the dolls together with the dolls and the boats together with the boats. Older children, instead of sorting the dolls and boats by kind, may well put each doll in one of the boats: in the case of a child who has been putting a doll in each boat, and has just one doll and one boat left, the right thing to do will typically be to put the remaining doll in the remaining boat. And even this is not generalizable, for the child may be enacting a story in which three people go off in boats and one person stays behind; in that context it would be wrong for the child to put the last doll in a boat.<sup>13</sup> So in saying that the children in our examples recognize norms I am not saying that they recognize general rules to the effect that, whenever you write the even numbers up to 1000 you should write 1002 next, or that whenever you are given a mix of four things of one kind and four of another you should sort the things by kind. Rather, I am asking you to imagine a child in a particular situation, say a two-year-old in one of these sorting studies who has put three dolls together and is now holding the fourth, and I am saying, of that child, that she recognizes, of the doll she is holding, that it should go with the other dolls. And again, this recognition need not require entertaining a thought that involves the concept *doll*; it could just be a matter of the child's recognizing, as she puts the doll with the other dolls, that this is the right thing to do.<sup>14</sup>

---

<sup>13</sup> I will say more about this kind of case in the next section. As I hope will become clearer in that section, the norms that apply in the more sophisticated cases are not themselves primitive: unlike the primitive norms, they are relative to an intention that the child has formed or a general rule which she has grasped. The norm recognized by a child who has, say, put three dolls in boats and now recognizes that the fourth doll should go in the fourth boat, is not on a par with the norm recognized by a younger child who has put three dolls in one place and now recognizes that the fourth doll should go in the same place. Rather than being alternatives, the more sophisticated norm is at a higher level than—in a sense presupposes—the more basic norm to which the younger child responds in the spontaneous sorting case. One way to see this is to note that the child who puts the fourth doll in the fourth boat is not only according with her intention of putting the dolls in the boats, but also, in a sense, “putting” the doll with the other dolls, and hence conforming to the primitive norm: not by placing the doll in the same location as the other dolls, but by treating it in the same way she treated the other dolls (i.e. by putting it in a boat).

<sup>14</sup> I speak interchangeably of the child's recognizing that *she should* to put the doll with the other dolls, and of her recognizing that the doll *should be* put with the other dolls (or that it belongs with them, or ought to go with them). I believe that, depending on the child's level of development, we can capture the child's attitude with either of these formulations, but that the impersonal formulation is more basic. This is because, as I argue in Ginsborg (forthcoming b), the child's capacity for self-consciousness, in the sense of the capacity to entertain “I” thoughts, depends on the more capacity to recognize normative demands like those made salient in activities like spontaneous sorting. It is because the child can think, of some piece of behavior, that it ought to be done, or is normatively called for, that she can adopt something like an “agent's perspective” (in the sense of Moran 2001) towards her behavior; this in turn allows her to distinguish what she *does* from what merely *happens*, and so to think of herself as a subject of thought and behavior, and not just as something to (or in) which thoughts and behavior happen.

The fact that we cannot describe these norms in general terms might well lead us to doubt that are any such norms. (Such doubt would be over and above any concerns we might have about normative realism, in particular the metaphysical and epistemological worries that I described in the previous section.) Because, as we have just seen, there are circumstances in which it can be appropriate to write 1004 after 1000, or put a doll with three boats, or say the word “shoe” while looking at a cat, it might seem that we can never say definitively that, for example, a two-year-old holding a doll and looking, now at the other dolls, and now at the boats, ought to put the doll with the other dolls. Nor, a fortiori, would it seem that, when the two-year-old puts the doll with the other dolls, she recognizes that it is the right thing to do. At most, it might seem, she takes it that it is the right thing to do; but this is in fact an illusion, because it would be equally correct for her to put it with the boats, or indeed to do something completely different with it. But I think that this particular doubt is misplaced. The fact that a four-year-old might play a game in which the doll has been sent off by the other dolls to watch over the boats, in which case she would have reason to put the doll with the boats rather than with the other dolls, does not make the same behaviour correct for the two-year-old. We do not need to be able to invoke a general rule to the effect that a doll should always be put with other dolls, no matter what the circumstances, in order to be able to say, of some particular two-year-old that the right thing for her to do in that situation, is to put the doll with the other dolls. Similarly, we do not need to be able to say that it is always correct to write 1002 after 1000 in a series beginning with 0,2,4,6,8 in order to be able to say that the ordinary pupil is going on correctly in writing 1002 and that the aberrant pupil is going on wrong. If we assume, as seems clear from the way Wittgenstein himself presents the case, that the pupil is not yet in a position to understand or formulate rules like “add two up to 1000 and then add four,” and that he is doing his best to conform to what is expected of him, then it seems clear that this pupil should write 1002 rather than 1004. So, if he writes 1004, he is failing to recognize a norm which applies to him, as well as to any other pupil in the same situation.

I have offered a *prima facie* case for accepting that the cases described here exemplify norms governing certain kinds of behavior under certain circumstances, and that these norms can be recognized

to hold by the people—specifically small children—engaging in the behavior. I now want to strengthen the case by arguing that the recognition of these norms is a condition of understanding linguistic expressions and, relatedly, entertaining intentional content.<sup>15</sup> I can begin to make this argument in connection with Wittgenstein’s example. Wittgenstein presents the example as one in which we are attempting to teach the pupil the meaning of various arithmetical expressions, in particular expressions like “+1,” “+2,” and “+3.” As I suggested in the previous paragraph, it is assumed that the pupil does not yet have the concept of addition, nor, more specifically, of adding two. Showing the pupil sequences like 0,2,4,6,8, and encouraging him to continue them is supposed to be a way of teaching him the concept of adding two, perhaps with a view to getting him to grasp the concept of addition more generally. And when these sequences are presented in association with the expression “+2” — as they are in Wittgenstein’s example, where the pupil is taught to write them when he hears the order “add two!” — then learning to continue them is also a way of coming to understand that linguistic expression. At least one criterion for saying that he grasps the concept of adding two is that, as Wittgenstein puts it, he “knows how to go on” with the series; and if this “going on” is something he does in connection with the expression “+2” then this is also a matter of his knowing how to go on in his use of the expression. (Saying “Now I know how to go on!” is tantamount to saying “Now I understand!”) But in order for the pupil to know how to go on in this way, he has to recognize that writing 1002 is the right thing to do given that he has just written 996, 998 and 1000. Otherwise we cannot speak of his *knowing* how to go on and hence of his *grasping* the concept or *understanding* the meaning of the associated expression. A pupil who had been successfully trained to respond to “+2” by writing sequences of even numbers, but

---

<sup>15</sup> By “understanding linguistic expressions” and “entertaining intentional content” I mean to pick out the kind of understanding and grasp of content which is characteristic of human beings as opposed to animals or artificially intelligent systems. So I mean something stronger than the sense of “understanding” in which a dog can understand commands like “Sit!” or grasp that the cat has run up the tree. This difference is familiar, but difficult to articulate in general terms. Brandom (1994, 88-89) claims that understanding in the stronger sense has a first-personal aspect—the difference between a parrot’s saying “red” to a red thing and a person’s calling the thing “red” is that the significance of “red” is available to the person but not to the parrot—but then explains the difference further in terms of the recognition of inferential connections among the various items of intentional content. I think it plausible that the difference turns on the presence or absence of a first-personal dimension, but I do not think that Brandom’s view succeeds in capturing that dimension (for this, see Ginsborg 2018). For the purposes of this discussion, I want to rely on the intuitive and pre-theoretical idea of the kind of understanding paradigmatically possessed by linguistically competent humans.

who could not recognize, in writing 1002 after 1000, that the number he was writing was correct given what he had written before, would not have grasped the concept of adding two or understood the expression “+2.”

A similar point holds for the other two groups of examples. The spontaneous sorting examples are different from Wittgenstein’s examples of going on in that children do not have to be taught, say, to sort one doll with the other dolls rather than the boats; this is precisely why psychologists label the sorting “spontaneous.” But we can think of what the children do as a kind of going on, like that of continuing a pattern: the child recognizes, after having separated out three of the dolls, that the right way to go on is to add the fourth doll rather than one of the boats. And, as with Wittgenstein’s examples of going on, a child does not have to grasp the concepts *doll* or *boat* in order to recognize that she should put one doll with the other dolls.<sup>16</sup> On the contrary, the capacity to sort the dolls together, and to recognize that this doll “goes with” the other dolls is part of what goes into the child’s acquiring the concept. Without that capacity she would never be in a position to learn the words “doll” and “boat,” nor, a fortiori, to understand sentences like “Let’s put the doll in the boat.” The behavior of the children in the false labelling examples is also a kind of going on: the 16-month-old who objects to the word “shoe” for the cat, and points to her own shoe when she hears the word, is showing that she knows the right way to go on with the word “shoe,” that is, the right way to use it given how it has been used in the past. This kind of example does not presuppose that the child understands the word “shoe” over and above being able to respond to its use as a label for a kind of thing: she does not, for instance, have to understand the experimenter as claiming that the cat is a shoe. Rather, her being able to recognize that the word “shoe” goes with her shoe rather than the cat, is of a piece with her being able to recognize the correctness of

---

<sup>16</sup> On a less demanding understanding of grasping a concept, one on which discriminative behavior is evidence of concept-possession (as in studies of animal cognition), grasping the concept will be a necessary condition of the capacity to sort, with or without the recognition of rightness in one’s sorting. But, as noted in the previous footnote, I have a stronger conception in mind.

sorting one shoe with other shoes: it is part of what puts her in a position to grasp the concept of a shoe and to understand utterances like “this is a shoe.”<sup>17</sup>

### III. Why primitive norms are distinct from other kinds of norms

In the previous section I tried to make clear, by appeal to examples, what I mean by “primitive norms.” and I argued both that these norms are in some sense genuine—in a sense that allows us to speak of their “recognition”—and that they can be recognized by children who have not yet mastered a language. I also argued that recognizing these norms is a prior condition of understanding language and grasping concepts.<sup>18</sup> I now want to say more about what kind of norms these are, and in particular to argue that they cannot be assimilated to any of the more familiar kinds of norms that I mentioned earlier. I think it should be clear that they are not moral norms or norms governing rational belief-formation.<sup>19</sup> Nor are they aesthetic norms, as standardly conceived. Even though it might in some sense “look right” for the doll to be put with the other dolls rather than the boats, or for 1002 rather than 1004 to follow 1000 in the sequence of numerals, there is a difference between recognizing this kind of appropriateness or correctness in an arrangement of objects or signs, and finding it beautiful. But it might seem that they can be assimilated either to norms of rational action, or to the norms governing specific cultural practices, paradigmatically the rules of a game.

I will begin by considering the second of these possibilities. Regarding Wittgenstein’s case of continuing the 0,2,4,6,8... series, it is tempting to suppose that our recognition of the correctness of 1002

---

<sup>17</sup> The point here does not apply only to children; an adult’s understanding of the word “shoe” also depends on a capacity to recognize the correctness of one use of “shoe” in the light of other uses, or (relatedly) of different shoes as “belonging” together. For more discussion, see Ginsborg (forthcoming a).

<sup>18</sup> Or at least that it is a condition of understanding the kinds of words and concepts that figure in the examples, those that are learned through activities of sorting and going on, and that apply, say, to middle-sized objects and to simple mathematical operations. But I am assuming that we need to understand such words and concepts in order to acquire other kinds of concepts, in particular those which can be learned only through verbal explanations. Thanks to an anonymous reviewer for prompting this clarification.

<sup>19</sup> Perhaps this is too fast. What if we understand a child who writes 1002 after 1000 as expressing the belief that  $1000+2=1002$ , or who sorts the doll with the other dolls, as expressing the belief that it is a doll rather than a boat? Then her recognition that this is what she ought to do could be the recognition, albeit implicit, that she has reason for the belief. However, in the examples on which we are focussing, the child does not yet grasp the meaning of the expression “+2” or the concept *doll*, at least not to an extent that would allow her to entertain propositions like  $1000+2=1002$  or *this is a doll*; a fortiori, she is not yet in a position to recognize reasons for believing them.

after 1000 amounts to the recognition that it accords with the rule *add two to the previous number*, and that, even if it has not been specified that this is in fact the rule, we tacitly assume that it is because we have a social or cultural practice of counting by twos. We could in principle replace or supplement it with a different practice, that of counting by twos up to 1000 and then counting by fours, in which case the rule in force would be *add two up to 1000 and then add four* (or, for short, *schmadd-two*), and it would be correct to write 1004 rather than 1002. This kind of practice would not come naturally to us, but under certain conditions it might make sense to adopt it. Similarly, we might suppose that a child who puts the doll with the other dolls, or points to her shoe when she hears the word “shoe,” is doing the correct thing only relative to a culturally determined rule that says to sort the dolls together, or to sort the shoes together. There could well be a context in which another rule for sorting might be called for. As I suggested in the previous section, a four-year-old might well devise a game whose rules call for putting the doll with the boats. In the context of a group of children playing this game, it would be incorrect to put the doll with the other dolls.

However, if I am right in arguing that the recognition of primitive norms, as manifested in the examples I have given, is required for grasping concepts and understanding linguistic expressions, then we cannot assimilate the norms to rules of social practices like games. For in that case the recognition of the norms would require being able to recognize which particular rule is in force—whether, say, it is the *add-two* or the *schmadd-two* rule—and to determine whether one’s behavior accords with that rule prior to being able to recognize one’s behavior as correct. We would need to be able to reason along the following lines: “The rule says to add two to the previous number, the previous number is 1000 and 1000 plus two is 1002, therefore I ought to write 1002” or “the rule says the dolls ought to go together; this is a doll and these are dolls; therefore this should go with these.”<sup>20</sup> But of course we already need concepts like “add two” and “doll” to engage in this kind of reasoning, and to grasp those concepts we need to be able to recognize, in situations like those described in the examples, that writing 1002 is correct or that the

---

<sup>20</sup> Note that it is not required, for recognizing rules of social practices, that we actually reason this way when we follow the rules; my point here is just that we must be able to do so.

doll should go with the other dolls.<sup>21</sup> Our ability to conceive of a range of different possible rules for continuing a series of numerals or sorting objects depends on our being able to take for granted the more fundamental level of normativity whose recognition allows us to formulate those rules. The point here is primarily philosophical, but it is reinforced by considerations from developmental psychology. It is highly implausible to suppose that the children in the spontaneous sorting and false labelling studies are capable, even in an inarticulate way, of grasping social rules and reasoning from the existence of these rules to the correctness or incorrectness of behavior. And yet there is every reason to think that they do recognize their own, and others', behavior as correct or incorrect.

Similar considerations tell against identifying the primitive norms with norms of rational action. We might think that the child's recognition that she should write 1002 after 1000 or put the doll with the other dolls amounts to the recognition, even inchoate, of a reason for that behavior. Perhaps the child developing the series recognizes that writing 1002 will please her teacher, and that this is a consideration which favors writing 1002. And even though the children in the sorting and false labelling experiments seem unconcerned about whether the adults nearby will approve of what they do—they don't hesitate, for example, to correct the experimenter who mismatches objects or applies false labels—we might still think that their behavior is motivated, at a more general level, by a desire to conform to the behavior of their social group. The child may recognize that, in putting the doll with the other dolls, she is engaging in behavior which is the same as the behavior of her parents when they put her toys away, or sort the forks and spoons in the cutlery drawer. Given her motivation to conform, that recognition would amount to the recognition of a reason for the behavior. However, in both kinds of case, the recognition of a reason requires a level of conceptual sophistication which, if the argument of the preceding section is correct, depends on, and so cannot be invoked to explain, the child's recognition of primitive norms. The child developing the series may well recognize that her writing 1002 will please her teacher, and she may also recognize that this gives her a conclusive reason to write 1002, or in other words that she "ought" to write

---

<sup>21</sup> This would not be the case if the reasoning were taking place at the subpersonal level; but then it would account at most for our behavior, and not for the recognition of that behavior as correct.



1002 in the sense of “ought” associated with the rationality of action. But if she does recognize that her writing 1002 will please her teacher, this could only be because she knows that her teacher wants her to write the correct number, and that 1002 is correct. So to recognize what she ought to do in the “reasons” sense, she must already recognize what she ought to do in the primitive sense. In the case of the child sorting the dolls, the recognition that she is conforming to the behavior of her social group requires her to grasp that her sorting the dolls in the experimental situation is the same as what her parents do at home with her toys or with the spoons and forks. That presupposes a grasp of the concept *same*, which in turn requires the capacity to recognize different situations as “belonging together” and thus the recognition of primitive normativity. In effect, she has to be able to sort together the situations of sorting dolls and boats in the psychologist’s lab, sorting toys at home, and sorting spoons and forks at home, and she has to be able to reason that her putting the doll with the other dolls is correct because that behavior “belongs with” her parents’ behavior of putting a fork with the other forks. So, here again, the recognition of a reason for her behavior depends on, rather than explaining, the recognition that one thing should go with another in the primitive sense. This is not to deny that children engaged in these activities are in fact highly motivated by a desire to conform to others’ behavior, and that this is a large part of the explanation why they do them. The point, rather, is that the fact of their behavior conforming to that of others cannot be recognized by them as a reason for the behavior unless they already grasp concepts like *conform* or *same*; which implies that if they do recognize that they have reason to behave in a particular way on some occasion, this recognition cannot be identified with the recognition of normativity required for grasp of those and other concepts.<sup>22</sup>

---

<sup>22</sup> In arguing that the recognition of primitive norms is not the recognition of reasons, I have focused on candidate reasons that are social in nature: that the behavior would please others, or conform to the practice of one’s group. This is in part because these are the reasons that I have most often seen suggested as explaining the relevant kind of behavior. However, as an anonymous reviewer helpfully points out, these are only examples, and it is not clear whether my argument generalizes to reasons of other kinds. Moreover, it might be thought that any recognition of behavior as favored amounts to the recognition of a reason, so that, in saying that the child recognizes her behavior (in putting the doll with the other dolls) as appropriate, she is *eo ipso* recognizing a reason for it. I am inclined to respond here that someone counts as recognizing a reason for her behavior only if she recognizes some specific respect in which the behavior is favored, or some specific fact that favors it; a mere recognition of the behavior as

#### IV. In what sense are primitive norms real?

I now want to return to the question of normative realism raised in the first section, and to ask how it applies to the norms I have identified. I have claimed that a normal child in the situation Wittgenstein describes—one in which she has received the usual kind of training in writing numbers and is developing the add-two series in the context of learning simple arithmetical expressions—recognizes that she ought to write 1002 after 1000. Similarly I have claimed that a child engaged in spontaneous sorting of dolls and boats—who arranges them without trying to conform to instructions she has either given herself or received from others—recognizes, of the doll that she has just picked up, that she ought to put it with the other dolls. Does this mean that it is an objective fact about the world that she ought to write 1002, or put the doll with the other dolls, in the same way that it is an objective fact that grass is green or that  $2+2=4$ ? My answer is yes and no: it is an objective fact, but not in a way that allows us to assimilate it to the paradigm objective facts cited in the first section. In saying that it is an objective fact, I mean to say that the *prima facie* reality of the relevant norms cannot be explained away by any of the standard anti-realist strategies. As I am about to argue, we cannot explain the child's supposed consciousness of the norms as non-cognitive; we cannot suppose that it is cognitive but erroneous; and we cannot say that the norm applies in virtue of the children's, or other people's, beliefs, attitudes or inclinations regarding how to go on with the series or how to arrange the dolls and boats. This means that the fact that the pupil ought to write 1002 satisfies the criteria for facthood in a relatively robust sense, one which goes beyond the mere entitlement to preface the statement "the pupil ought to write 1002" with "It is a fact that the pupil ought to write 1002." But I also want to claim that facts about

---

avored would not amount to the recognition of a reason. But any recognition of such a respect or fact—even as minimal as *they are all dolls* or *this is the same as those*—would, as with the examples I have given, require the use of concepts whose grasp presupposes the recognition of primitive norms. I discuss this further in section V of Ginsborg forthcoming b.

primitive norms are fundamentally different from the empirical and mathematical facts that are often invoked, in debates about normative realism, as standards for being an objective fact. So although my position is realist in the sense that I reject the standard anti-realist approaches to understanding the norms, it is not realist in the sense of the dialectic introduced at the beginning of this paper.

Before pursuing that last point, I will explain why I think that the primitive norms are not amenable to the standard anti-realist treatments. Let us begin with the non-cognitivist approach. One natural form such an approach might take would be roughly emotivist: the child's seeming recognition of the correctness of writing 1002 or putting the doll where she does would, in reality, amount to something more like an ungrounded preference for writing 1002 or putting the doll there.<sup>23</sup> On this kind of view, when the child presented with the mismatched objects says "No belongs this way," the right way to understand her is as expressing her dislike for the arrangement, or, depending on the circumstances, her desire to switch the objects around. For our part, when we say that the child is sorting or going on correctly or incorrectly, we are either saying that she conforms or fails to conform to a rule we have in mind—so that the normativity is not primitive—or expressing our own preference or liking for one arrangement or pattern over another. In apparent favor of this approach, it seems plausible that small children do prefer to arrange things in such a way that objects of the same kind are placed together. If they did not, then it is hard to explain why they would in fact arrange them in that way. And even though adults might find these kinds of arrangements boring from an aesthetic point of view, there are still reasons why we might prefer that children — our own especially — arrange things and continue patterns in the developmentally normal way.

---

<sup>23</sup> In focusing on this version of non-cognitivism, I am oversimplifying in two ways. First I am not considering other potential non-cognitivist accounts, e.g. those (like Gibbard's) which identify seeming grasp of norms with the undertaking of plans. Second, as an anonymous reviewer has pointed out, many non-cognitivists would claim that mere preferences are not enough to ground normative judgments: the child's "no belongs this way" has to express something more fine-grained than a simple liking for the arrangement. See e.g. Björnsson and McPherson 2014. However, although the point deserves further discussion, I believe that at least the second reason I offer below for rejecting non-cognitivism generalizes to more sophisticated versions of noncognitivism than that sketched here.

However, the fact that children prefer some ways of sorting things over others does not mean that we can dispense with the idea that they recognize primitive norms. This is for two reasons. First, it is psychologically implausible to suppose that children have a basic preference for same-kind arrangements of objects or repeating patterns in the same way that they have a basic preference for food over hunger, stimulation over boredom, and social contact over isolation. It is much more plausible that, if the child prefers to put the doll with the other dolls, it is because she recognizes that this is where it belongs, or that this is the correct place to put it. Gopnik and Meltzoff (1997) make a similar point when they describe the children's sorting responses, in the spontaneous sorting studies, as "removed from any immediate reward" (p. 172). Instead, they claim, "categorization is its own reward" (ibid.). There is no reason to think that there is anything intrinsically satisfying about writing 1002 after 1000 or putting the doll with the other dolls. Rather, whatever satisfaction the child feels is the satisfaction of having got something right, that is, of having met a normative demand. The second reason we cannot dispense with the recognition of primitive norms [[ primitively normative attitudes]] is that, if the argument I made in section II is correct, this recognition is necessary for understanding language and grasping concepts. If the child does not recognize that she ought to write 1002 after 1000, and instead merely enjoys doing so, then, although we can predict that she will go on as we do with the sign "+2", we cannot say that she *knows* how to go on with the sign, since this requires knowing that 1002 is the correct thing to write. So we cannot say that she understands the sign. Similarly, a child who likes putting one doll with the other dolls but does not recognize that the one doll "belongs with" (should be put with) the others, is not in a position to grasp the concept *doll*. In brief, if we do not accept a broadly cognitivist view of the primitive norms, then we have to give up the idea that we know the meanings of the expressions that we use or grasp the corresponding concepts. This leads, arguably, to an unacceptable skepticism about meaning and understanding; and, in the context of the debate about normative realism, it leaves us unable to explain how we can understand the paradigmatically fact-stating claims in terms of which the debate is framed.

The argument I just described applies also against the error-theoretic approach. What we need, in order to make sense of the possibility of understanding and meaning, is not just that the child *take* 1002 to

be the correct thing to write after 1000 but that she *recognize* the correctness of 1002. Otherwise, again, we cannot say that she *knows* how to go on. All that we can say is that, in going on as she does, she *believes* that she is going on correctly, and this is not enough for her to *understand* the expression “+2” as it is used in connection with the series, or to *grasp* the principle of the series. In a nutshell, if the primitive norms are illusory, then so are linguistic meaning and grasp of concepts. The upshot is a form of skepticism about meaning that is, arguably, self-defeating and, at the very least, intuitively unappealing.<sup>24</sup>

That leaves what I called the subjectivist approach, on which the norms do correspond to facts, but where the facts are in some distinctive way dependent on our beliefs, attitudes, desires or inclinations. On such an approach, we would see the primitive norms as grounded either on our inclinations to go on and to sort things, as we do, or on our normative judgments regarding the correctness of our particular ways of going on and sorting things. At a first, crude, approximation, the correctness of the pupil’s going on with 1002 would be grounded simply on the fact that he does go on with 1002, or that he judges it appropriate or fitting to go on with 1002. On a view like this, Wittgenstein’s aberrant pupil, who goes on with 1004, is also going on correctly, since writing 1004 is what he is inclined to do, or judges appropriate. On a slightly less crude version of the approach, the subjectivist might hold that what determines the norms for any individual human being is not the inclinations or judgments of that particular individual, but rather the inclinations of her community, or of human beings in general. It is the fact that the aberrant pupil diverges from the rest of us in his judgment about how to go on after 1000 which makes his judgment wrong.

An initial reaction we might have to these versions of subjectivism is simply that they are unmotivated. Why should the actual inclinations or judgments, either of an individual, or of human

---

<sup>24</sup> In case this rejection of the error-theoretic approach seems too fast, I should point out that I find the error-theoretic approach to ethics equally implausible, since it requires us either to reject our ordinary conviction that murder is not wrong, or to suppose that metaethical talk creates a special context in which we can say that murder is not wrong without contradicting our ordinary conviction that it is (see Joyce 2021, 3.2); neither of these options strike me as defensible. The standard example of an error theory of ethics is Mackie’s view. But as noted earlier, I agree with Berker (2019) in reading Mackie as a subjectivist rather than an error theorist.

beings in general, determine how the individual ought to judge? On the first version of the view, where the pupil ought to write 1002 because he is inclined to, or judges that he ought, to write 1002, there seems to be no normativity in the picture at all; in order for there to be normativity there has to be at least the possibility of a gap between what a person does, or thinks she ought to do, and what in fact she ought to do. On the second version of the view, there is such a gap: the aberrant pupil is not doing what the rest of us do. But it remains unexplained why what we do should constitute a standard for what the pupil ought to do. There seems to be no content to the supposedly normative claim that the pupil ought to write 1002, over and above the claim that, in writing 1002, he is doing something different from what we do.

A less crude, and *prima facie* more attractive, view would have the correctness of writing 1002 depend not on the pupil's inclinations or judgments regarding what to write after 1000, but on inclinations or judgments about what to do in other situations. Here I am thinking of a model along the lines of Sharon Street's Humean constructivism in metaethics, on which normative facts about an individual's reasons are determined by the judgments she makes about her reasons—judgments that may in part be determined by her inclinations—but where the link between the judgments about reasons and the actual facts about reasons is far less direct than on the picture we started with.<sup>25</sup> On Street's view, our judgments about reasons set standards for one another in a way which allows us to distinguish the fact of my judging that I have reason to do something in any given instance from the fact of my actually having a reason to do it. If I judge that I have conclusive reason to live a healthy life, then I am mistaken in judging that I have reason to keep on smoking and to refrain from exercise, at least assuming that the first judgment is more deeply held and more part of my identity than the other two judgments. Practical reason has a rich formal structure that allows us to make sense of the idea that judgments about reasons can be consistent or inconsistent with one another; this possibility of inconsistency allows us to make room for the possibility that some of our judgments can be in error, as required if there is to be genuine normativity in the picture. The erroneous judgments will be those that are inconsistent with our most deeply held

---

<sup>25</sup> E.g. Street 2008.

judgments, those with which we most fundamentally identify: these latter judgments are the source of the rational norms by which our actions are governed.

However it is a key element of Street's account that we must be able to understand the consistency or inconsistency among our different judgments about reasons in a way which is not itself normative. Deciding that a judgment about one's reasons is consistent with, or "withstands scrutiny from the standpoint of" one's other judgments does not itself require a substantive normative judgment about whether the judgment is rational given one's other judgments (2008, 232). If it did, the constructivist account would be "unhelpfully circular" (2008, 231-232). This feature of constructivism about reasons means that the approach cannot be carried over to the case of the primitive norms. We can see this by considering how we might try to give a constructivist account of primitive norms. At least following Street's formulation of constructivism, we would need to say that the pupil ought to write 1002 not simply in virtue of his judgment that he ought to write 1002, but in virtue of that judgment's withstanding scrutiny from the point of view of his other judgments. What might those other judgments be? A plausible candidate would be other judgments about how to go on in the kinds of series and pattern-continuation exercises described by Wittgenstein. We might suppose that the pupil had previously learned to develop the series of natural numbers, and that, in doing so, he had written the sequence 999, 1000, 1001, 1002, and judged that writing 1002 in that context was correct. Or we might suppose that, in drawing patterns of alternating dots and dashes, he invariably found it correct to draw a dash after a dot and vice versa, no matter how long the pattern was, and, in particular, that he did so even after five hundred iterations of the dot-dash pattern.<sup>26</sup> In the case of the child finding it correct to sort the doll with the other dolls rather than the boats, we might see her judgment as consistent with other judgments she made in other activities of sorting or early language use. She might, for example, know to respond to the expression "show me your doll" by showing her doll, rather than a boat. Or, in spontaneously sorting toy

---

<sup>26</sup> On this attempt at a constructivist approach to primitive normativity, the aberrant pupil might be correct in writing 1004 after 1002 if, for example, he continued the series of natural numbers by writing ...999, 1000, 1002, 1004, or if, after five hundred iterations of the dot-dash pattern he started to draw two dots followed by two dashes (and, in each case, judged that what he was doing was correct).

animals and vehicles, she might take the sheep to belong with the cow and the goat, rather than with the car and the aeroplane. Her judgment about where the doll should go would be correct in virtue of its consistency with these other judgments. The problem, though, is that in the case of the primitive norms, deciding that one judgment is consistent with another itself involves a substantive normative judgment, and indeed a judgment of the very same kind as those whose consistency we are trying to assess. In deciding whether the pupil's judgment that it is correct to write 1002 after 1000 in the context of the series 0,2,4,6,8... is consistent with, or withstands scrutiny from the point of view of, his judgments about what to write in other series- or pattern-continuation cases, we are making another judgment of primitive normativity: whether the pupil is "going on" correctly if, having recognized the correctness of ...999, 1000, 1001, 1002 in developing the series of natural numbers, he now judges that 1000, 1002, 1004 is correct in developing the series of even numbers. Similarly, in taking the child's judgments across the various sorting activities to be consistent with one another, we are in effect taking it to be correct to sort those various activities together: we take it that responding to the label "doll" by showing a doll can be sorted with the activity of putting the dolls together in the same location, or with the activity of sorting together the different animals. So regardless of whether Street's own metaethical constructivism about reasons avoids circularity, the attempt to apply the same model to primitive norms fails to do so.

I have not ruled out that primitive norms could be amenable to a different kind of subjectivist treatment.<sup>27</sup> But given the relative attractiveness of Street's constructivism as a form of subjectivism, I think I have done enough at least to cast doubt on the prospects for a subjectivist account of primitive normativity. The upshot, if my arguments against non-cognitivist and error-theoretic treatments of these norms are accepted, is that the standard anti-realist accounts of normativity fail in the case of primitive normativity. This allows us to say that, in a relatively robust sense of "objective fact"—one which does not apply on anti-realist approaches—it is an objective fact, say, that the pupil in Wittgenstein's example ought to write 1002 after 1000.

---

<sup>27</sup> In particular, I have not considered whether a form of Kantian constructivism might work for primitive normativity. Perhaps the materials for it could be found in Kant's *Critique of Judgment*. But I suspect that, in that case, it would not qualify as subjectivist nor, more generally, anti-realist.



However, the considerations I have raised in arguing for the genuineness of these norms, and specifically that we need to recognize them in order to make sense of the possibility of meaning and understanding,<sup>28</sup> suggest that the corresponding normative facts cannot be assimilated to paradigmatic descriptive facts like “grass is green,” “protons have positive charge,” and “ $2+2=4$ .” This is not just because they are normative, but because—unlike other normative facts—they are in a certain way more fundamental than the paradigm facts. They are more fundamental in the sense that there can be no grasp of the paradigm facts without a prior grasp of the primitive normative facts. If it is not possible to recognize, as the children in the examples do, that 1002 “comes after” 1000 in the 0,2,4,6,8... sequence or that one green thing “belongs with” other green things in a way that it does not belong with blue things, then it is not possible to grasp the concepts of *adding two* or *green*, nor the propositions that  $2+2=4$  or that grass is green. So the normative fact that 1002 comes after 1000 in the series 0,2,4,6,8... is not on all fours with the fact, say, that  $1000+2=1002$ : it is not an arithmetical fact, or a fact about numbers, but rather a fact which must be recognized if arithmetical facts are to be so much as thinkable by us.<sup>29</sup> Moreover, if

---

<sup>28</sup> The idea that we need to recognize primitive norms, and to regard them as objective, in order to make sense of the possibility of meaning and understanding has some affinity with David Lewis’s view (1983) that we need to invoke an objective distinction between natural and non-natural properties in order to solve the “problem of interpretation” and thus to answer the skepticism about meaning raised by Kripke’s Wittgenstein. The most important difference is that on my view the primitive norms are prior to, and explain, the privileged status of properties like adding and green over quadding and grue, whereas for Lewis, the privileged status of these properties is metaphysically fundamental and, presumably, accounts for the correctness of going on with 1002 or applying “green” to a newly perceived green rather than a newly perceived blue thing. On Lewis’s picture, as I understand it, Wittgenstein’s pupil ought to write 1002 after 1000 because it is correct for him to interpret his previous behavior as adding-two rather than schmadding-two, and it is correct for him to do that because adding is a more natural property than schmadding. So for Lewis, as on my view, there is an objective “ought” governing the pupil’s present behavior; but in contrast to my view, the “ought” is not primitive, but derives from the metaphysical fact that adding is natural. I see two respects in which Lewis’s view is problematic. The first, which is closely related to the “problem of missing value” raised by Dasgupta (2018), is that we could recognize that adding is natural, and still deny that the pupil ought to write 1002: perhaps, rather than interpreting his past usage in terms of the more natural property of adding two, he should interpret it in terms of the more schmatural property of schmadding two, and thus write 1004. The second is that Lewis’s view does not account for the apparent psychological fact that children can know how to go on as they ought, and to sort things correctly, without having the conceptual resources to interpret their previous behavior or that of their teachers. My view avoids these difficulties by taking the oughts, and their recognition, as basic. We can recognize the correctness of going on as we do in the 0,2,4,6,8... series or in our use of the sign “+” without first having to determine what function we were computing in our past use or a referent for “+”, and this is sufficient for avoiding skepticism about meaning à la Kripke’s Wittgenstein. For more on this, see Ginsborg forthcoming a.

<sup>29</sup> Wittgenstein may be making this kind of point when he writes, at *Zettel* §357, “Z 357: We have a colour system as we have a number system.// Do the systems reside in our nature or in the nature of things? How are we to put it?-- Not in the nature of numbers or colours.” His answer leaves open that the number system (of which the correctness

we identify facts as true propositions, and if we suppose further that it is essential to something's being a proposition that it can in principle be grasped, then the sense in which the normative facts are more fundamental than the paradigm descriptive facts can be stated even more strongly: the existence of the descriptive facts itself depends on the existence of the normative facts.<sup>30</sup> To speak of the existence of a fact that grass is green, on this way of thinking, is to refer to a potential content of thought, the proposition that grass is green, and to say that a belief with that content would be true. But for there to be potential contents of thought, it must be possible in principle to have a contentful thought, and, I have been arguing, this is excluded if there are no primitive norms.<sup>31</sup>

I have argued, then, that primitive norms are real in a relatively robust sense: they cannot be accounted for along non-cognitivist or error-theoretic lines, and they are not grounded in our subjective inclinations or normative judgments. But, I have also argued, they are not real in the sense of the dialectic, because normative facts are different in kind from the paradigmatic descriptive facts. This, to repeat, is not just because they are normative. We could agree with the moral realist that, say the fact of an action's being wrong is on a par with the fact of its being carried out on a Tuesday or its being a case of gratuitously inflicting pain on someone else, and still deny that the fact of its being correct to go on with 1002 after 1000 is on a par with the arithmetical fact that  $1000+2=1002$ . The difference has to do

---

of writing 1002 after 1000 in the 0,2,4,6,8... series is presumably a part) is "in the nature of things" very generally construed, even though the norms defining the number system (e.g. that 1002 should come after 1000) are not part of arithmetic. Interestingly, Bernard Williams, citing this passage at the conclusion of his influential "Wittgenstein and Idealism" (1973, 95), misquotes the last sentence to read "Not in the nature of things," suggesting that Wittgenstein may be more of an idealist than he actually is. (The misquotation is corrected in later editions.)

<sup>30</sup> Here I am taking for granted a non-minimalist sense of "fact"; but I believe that we need to do so in order to for the debate about normative realism so much as to make sense. See Boghossian 1990 ("The Status of Content").

<sup>31</sup> This might seem like an unacceptably idealist view. Here are two considerations which might help allay that concern. First, the existence of the descriptive facts does not depend on the existence of human beings, but only on its being the case that, if there were human beings they would be subject to the primitive norms. It would be a fact that protons had negative charge and that  $2+2=4$  even if no human beings had never evolved. We need to appeal to the notion of human beings (or possibly of thinking beings more generally, although I will not pursue that) in order to make sense of the notion of a fact, but that does not make the facts themselves depend on the actual human beings. Second, if there were no primitive norms—if, roughly speaking, any way of sorting things or of going on in a series was as good as any other—there would still be protons, blades of grass, numbers, and so on, and there would also still be properties like being green or having positive charge. Any given blade of grass would still be green, as well as having any number of grue-like properties. What would be missing would be the (non-minimalistically-construed) fact of the grassblade's being green, which I take to be something over and above the grassblade and the set of green things, something that is (to borrow a term from John McDowell) "thought-shaped." (See e.g. McDowell 2018).

with the order of priority for which I have been arguing between the normative facts and the descriptive facts. Even if we reject the strong metaphysical priority suggested at the end of the previous paragraph, on which the existence of the arithmetical facts or colour facts depends on the existence of the normative facts about how we should go on or what should be sorted with what, I think it is still enough to note that our grasp of the descriptive facts depends on our grasp of the normative facts in order to justify the claim that these facts are not on a par. The point can be reinforced by noting that there can be no question, as there might be in the moral case, or in other more familiar cases of putative normative facts, of grounding facts about primitive norms on the descriptive facts. A moral realist—perhaps of a naturalist kind—might hold that an action is wrong in virtue of the fact that it gratuitously inflicts pain on others. But the primitive correctness of writing 1002 after 1000 or putting one green thing with the other green things does not depend on the fact that  $1000+2=1002$  or that the thing is green. It is only if we assume that the correct way to develop the series is to add two rather than to schmadd two, or that it is the green, rather than the grue, things that should be sorted together, that the correctness of writing 1002 after 1000 or putting one green thing with the other green things follows from these descriptive facts. And in that case the normativity is not primitive, but rather amounts to accordance with a general rule (*add two to the previous number, sort the green things together*), grasp of which presupposes recognition of the primitive norms governing what to write after 1000 or where to put this particular green thing.

The upshot of these considerations, as regards the status of the primitive norms, is that they are, in an intuitive sense real—and that in particular their existence cannot be explained away through non-cognitivist, error-theoretic, or subjectivist approaches—yet that they are not real in the sense of the dialectic introduced at the start of this paper. My position regarding these norms is thus, in a sense realist, but different from the kind of position typically taken by realists about the more frequently-discussed kinds of norms, paradigmatically moral norms or the norms that apply to us in virtue of there being reasons. This difference, however, does not bring me any closer to an anti-realist view of these norms. On the contrary, I see in the line of thought I have presented a potential argument against normative anti-

realism: not just anti-realism about the particular kind of norm on which I have been focussing, but anti-realism of other kinds. I will conclude by sketching that argument.

I claimed at the start of this paper that, as debates about realism are traditionally conducted, both the realist and the anti-realist take for granted that there are straightforwardly descriptive properties and facts which serve as a standard for what counts as real. The question of normative realism is often posed as the question whether properties like wrongness, or being a reason for an action, or aesthetic value, are out there in the world in the same way that properties like being the sum of 2 and 2, or being positively charged, are out there in the world. Bracketing the possibility of anti-realism about numbers or theoretical physics, realists and anti-realists typically take the existence of such properties, and the corresponding facts, for granted; and they both typically take for granted that these properties and facts are knowable by us. The problem the anti-realist sees with the putative normative properties and facts is that there is supposed to be something puzzling about them in comparison with their descriptive counterparts, and, in particular, that it is hard to explain how we can come to know them, as we can come to know the descriptive facts and properties.

But if I am right both about the existence of the primitive norms, and that we need to recognize them in order to grasp the descriptive facts and properties, then someone who is an anti-realist about some normative domain cannot help accepting the reality, in an intuitive sense, of the primitive norms. That is, the anti-realist about, say, ethical normativity, or the normativity associated with reasons for action, cannot also be an anti-realist about primitive norms. For that would prevent her from taking for granted the paradigm descriptive facts, and our knowledge of them, as standards with respect to which the putative normative facts in her domain of interest fall short. Her denial of the reality of primitive norms would undermine her reliance on our grasp of descriptive facts as the contrast case or foil against which her denial of reality to ethical norms, or to the normativity of reasons, makes sense. This argument does show that anti-realism, regarding any given domain, is untenable. But it does show that anti-realists must allow the reality, in an intuitive sense, of at least some norms, specifically those primitive norms whose

recognition is required for the possibility of meaning and concepts.<sup>32</sup> This in turn means that the justification for anti-realism in a given domain cannot rest simply on the putatively normative character of that domain. If the anti-realist sees a problem with the reality of reasons, moral norms, or aesthetic values, it must be a problem specific to morality, reasons or aesthetics. The anti-realist cannot appeal to the sheer normativity associated with reasons, morality, or aesthetics as a ground for questioning the reality of the corresponding properties and facts.

---

<sup>32</sup> An anonymous reviewer has suggested that the structure of this argument is something like that of an objection Gibbard considers, in light of his view that meaning is normative, against his own anti-realist account of normative claims as expressing plans (Gibbard 2012, 227ff.) The objection arises when we consider that expressivism is itself a thesis about meaning, so that Gibbard's claims about the meaning of normative discourse are themselves normative, leading to at least the appearance of a circularity: Gibbard seems to depend on the possibility of making normative claims (about meaning) in order to argue for an anti-realist conception of normativity generally. Gibbard's response is to accept the circularity and to deny that it is vicious. The normative claims he makes in support of his expressivist view are themselves to be understood as expressions of plans for the use of sentences, and even though that claim in turn is a claim about meaning, it can in turn be understood as the expression of a plan. The reviewer asks whether a similar, circularity-embracing, strategy could be used to deny the reality of primitive norms. This is an interesting question, but I have doubts both about whether Gibbard's strategy is successful in its own terms, and, if it is, whether it could be applied in connection with my view, given the differences between how I see the anti-realist's reliance on primitive norms and how Gibbard sees the dependence of his own anti-realism on norms of meaning. Unfortunately space does not allow further discussion here.

## REFERENCES

- Berker, Selim (2019), 'Mackie was not an error theorist', *Philosophical Perspectives*, 33 (1).
- Björnsson, Gunnar and McPherson, Tristram (2014), 'Moral Attitudes for Non-Cognitivists', *Mind* 123 (489),
- Blackburn, Simon (1993). *Essays in Quasi-Realism* (New York: Oxford University Press).
- Boghossian, Paul A (1990), 'The status of content', *The Philosophical Review*, 99 (2), 157-84.
- (2008), 'Epistemic rules', *The Journal of Philosophy*, 105 (9), 472-500.
- Brandom, Robert (1994), *Making it Explicit* (Cambridge, Mass.: Harvard University Press).
- Dasgupta, Shamik (2018), 'Realism and the Absence of Value', *Philosophical Review*, 127 (3), 279-322.
- Gibbard, Allan (2012), *Meaning and Normativity* (Oxford: Oxford University Press).
- Ginsborg, Hannah (2018), 'Normativity and Concepts.' In Daniel Star (ed.), *The Oxford Handbook of Reasons and Normativity* (Oxford: Oxford University Press).
- (2021), 'Conceptualism and the Notion of a Concept', in Christoph Demmerling and Dirk Schröder (eds.), *Concepts in Thought, Action and Emotion* (New York and Oxford: Routledge), 42-59.
- (2022), 'Going on as one ought: Kripke and Wittgenstein on the normativity of meaning', *Mind & Language* 37, 876–892.
- (Forthcoming a), Putting Wittgenstein Back Into Kripkenstein: Meaning Skepticism and Knowing How to Go On. In Claudine Verheggen, (ed.), *Kripke's Wittgenstein on Rules and Private Language at 40* (Cambridge: Cambridge University Press).
- (Forthcoming b), 'Spontaneity without rationality: a Kantian approach to self-consciousness and perceptual content', in Johannes Roessler Andrea Giananti, and Gianfranco Soldati (eds.), *Perceptual Knowledge and Self-Awareness* (Oxford: Oxford University Press).
- Gopnik, Alison and Meltzoff, Andrew N (1997), *Words, Thoughts, and Theories* (Cambridge, MA: MIT Press).
- Joyce, Richard (2021), 'Moral Anti-Realism ', in Edward N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy* (Winter 2021 Edition).
- Lewis, David (1983), 'New work for a theory of universals', *Australasian journal of philosophy*, 61 (4), 343-77.
- Mackie, J.L. (1977), *Ethics: Inventing Right and Wrong* (Harmondsworth: Penguin).

McDowell, John (2018), 'Travis on Frege, Kant, and the Given', in Rasmus Thybo Jensen Johan Gersel and Søren Overgaard Morten S. Thaning (eds.), *In the Light of Experience* (Oxford: Oxford University Press), 23-35.

Scanlon, T. M. (2014), *Being Realistic About Reasons* (Oxford: Oxford University Press).

Street, Sharon (2006). 'A Darwinian dilemma for realist theories of value', *Philosophical Studies*, 109-66.

--- (2008), 'Constructivism about reasons', in Russ Shafer-Landau (ed.), *Oxford Studies in Metaethics, Volume 3* (Oxford: Oxford University Press), 207-45.

Williams, Bernard (1973), 'Wittgenstein and idealism', *Royal Institute of Philosophy Supplements*, 7, 76-95.