

Nick Zangwill, Durham

Non-Cognitivism and Consistency

I argue for cognitivism about some normative judgements. I begin with the issues of realism and cognitivism as they manifest themselves in moral philosophy. I then proceed to issues of realism and cognitivism about normative judgements more generally. I describe the norm of consistency in normative judgement, and I argue that this norm means that we must be cognitivist about some normative judgements.

§ 1. The Realism/Cognitivism Debate

There is a debate about the metaphysics of morality. Are moral good, evil, moral obligation, justice, virtue and the like, real or genuine properties of the world? Or are there no such genuine properties? Are these properties mind-independent in some sense? This is supposed to be a metaphysical issue about the world, not an issue about moral thought or language.

However, it is not easy to pursue this debate. It is not clear what the ground-rules are. Furthermore, some philosophers have been sceptical about whether there is a genuine issue here. Metametaphysical scepticism is often just dogmatically asserted as if it is obvious once stated.¹ But there is an interesting form of such scepticism. Moral projectivists, expressivists or prescriptivists, who think that moral judgements are, or express, propositional attitudes that are not beliefs, sometimes say that they are quite happy to allow that we can and should talk in terms of moral truths, facts, states of affairs and knowledge (see for example Blackburn, 1984, 1993, 1998). They are happy to *appropriate* such notions and give them meaning from within their own theories. If so, what terms are there left in which to frame the metaphysical debate? If there are none, then that debate has evaporated (Blackburn 1980, Zangwill 1992a). This is an argument, not just anti-theoretical assertion. I am not endorsing this kind of sceptical

¹ Wittgensteinians and post-Kantians are rather prone to the vice of dogmatic anti-theory.

argument. But it does supply some reason to shift our attention away from the world and its properties to our moral judgements and moral language. This will be my main interest in what follows.

What then is the issue about moral judgements? We make moral judgements that we sometimes express in language by saying things like “Queen Isabella of Spain was evil”. What we are interested in is the nature and explanation of these judgements – of these mental acts. Scepticism about the existence of a genuine issue here would be perverse if not unintelligible. For moral judgements are real phenomena in the world and there is obviously something to be said about them.

It is common to make a distinction between beliefs and desires, even if many contemporary philosophers use the words “belief” and “desire” with a specific meaning that does not coincide with some aspects of the ordinary English uses. Philosophers find a more general distinction between ‘cognitive’ and ‘non-cognitive’ states, of which beliefs and desires are central cases. In an earlier age this general distinction might have been between ‘reason’ and ‘passion’. There are two broad kinds of mental states here. Since Elizabeth Anscombe it has been usual to explain the difference in terms of the different ‘directions of fit’ of the two sorts of state. Anscombe’s example is of a husband who goes shopping with a shopping list. He uses the list to guide his actions. A detective follows him, making a list of the husband’s purchases. The detective’s list is a record of what happens in the world. But the husband’s list is a guide to what he is to do to the world (Anscombe 1959). The general point here is usually said to be that the point of beliefs is to fit to the world, but the point of intentions or desires is to get the world to fit them. There is a lot more to say about the idea of direction of fit, but I shall not pursue this here (see Zangwill 1998).

Perhaps a mental state can be of both kinds, and perhaps there are other kinds of mental state, but neither possibility casts doubt on the intelligibility of the standard debate over moral judgements, for that debate is not entirely characterized in terms of beliefs and desire. What is more important is the *content* of the mental states that are, or are expressed in, moral judgements. We can distinguish between ‘cognitivism’ and ‘non-cognitivism’ about moral judgements as follows. Cognitivists say that moral judgements are beliefs that represent moral facts (things possessing moral properties), and when true they are true in virtue of the moral facts they represent. Non-cognitivists think moral judgements are, or express, emotions, desires or pleasures with non-moral content, which we have in reaction to our cognition of non-moral facts. To put it

in terms of direction of fit, moral cognitivists say that moral judgements must fit the moral world, whereas non-cognitivists say that the point of moral judgements is not to understand the world but to change it – that is, not to understand the moral world but to change the ordinary world, or at least to react to it with feeling.

Since cognitivism says that moral judgements represent moral facts, it leaves open the possibility of an error-theory according to which our judgements aim to represent moral facts but sadly there are no (positive) moral facts (Mackie 1977). This view is analogous to atheism. Our judgements are cognitive but all our (positive) judgements are false. By contrast the non-cognitivist denies that moral judgements aspire to represent moral facts. They are a felt or motivational reaction to our cognition of non-moral facts.

That is an outline of issue over moral judgements. Now this range of issue crops up in an exactly parallel form for norms other than moral norms. There are norms of epistemology, law, aesthetics, etiquette, practical reason, sport, religion, and so on. In each of these areas it seems that we can raise issues about the metaphysics of the normative properties in question; and even if we cannot, in each of these areas we definitely can raise issues about the nature of the normative judgements.

In this paper, I shall argue that we cannot be non-cognitivist about *all* normative judgements. *Some* normative judgements must be cognitive. I argue from the requirement of consistency in our normative judgements. In sections 2 to 4 I describe the consistency requirement on moral judgements. In sections 5 to 8 I argue from consistency to cognitivism. It could be that we can argue that if some normative judgements are cognitive, then they all are. But I shall leave that to one side.

§ 2. The Consistency Requirement

Normative judgements have certain *aspirations*. Both cognitivism and non-cognitivism about normative judgements need to say something about these aspirations.

One such aspiration is the aspiration to *correctness*. We aim to make the *right* or *correct* judgement. But we can get it wrong. We are fallible.²

² This aspiration to correctness is closely connected with the idea that the correctness of a normative judgement is 'mind-independent' (See Zangwill 1994).

There is a contrast here between moral thought and our thought about what is funny and what is nice. It is plausible – or at least many find it plausible – that there is no right and wrong for judgements of funniness and niceness.³ I shall not say much about the aspiration to correctness.

The other aspiration, which I want to focus on, is the aspiration to *consistency* between our normative judgements. This is the idea – roughly – that we should make similar normative judgements about similar cases. The aspiration to consistency is *not* just the constraint that we should not think both that X ought to be done and also that X ought not to be done. This is mere *logical* consistency, and it is probably implied by the aspiration to correctness. The consistency requirement is not just a matter of how we ought to judge, but of how we ought to judge given our other normative judgements that are made on a similar basis. The idea of merely logical consistency involves nothing like that.⁴

The consistency requirement is often deployed with effect in moral argument. For example, a person may criticize another person for accepting bribes but think that a similar bribe that he accepted was alright. But if he does not think that who accepts a bribe makes a difference, someone may use that with effect to get him to alter one or other of his judgements. Another example: Peter Singer has persuaded many people to change their view of animal rights by arguing that we think that causing unnecessary pain to children is wrong just because it is pain, but we don't think causing unnecessary pain to animals is wrong even though it is no less pain (Singer 1975). So he argues that we should alter our judgement about causing pain to animals. This argument has been very effective. Appeals to consistency are effective in moral argument, which shows that we are aware that there is a norm that we should respect in moral thinking. We may achieve consistency by adjusting either judgement in the above examples. But consistency is a norm bearing on our judgements and it is separate from the one that enjoins us to make the correct judgement.

The normative judgements in question may be what are called “*pro tanto*” judgements – ones that can be outweighed. For example, if we judge that an act is *pro tanto* morally wrong just because it is theft, then

³ Kant distinguishes ‘judgements of taste’ from ‘judgements of the agreeable’ by appeal to the aspiration to correctness (Kant 1928).

⁴ This requirement of consistency is connected with the idea that we judge that X is N (some normative property) *because* it is G (some grounding natural property). X's being G *makes* it N; it is N *in virtue of* being G (Zangwill 2006, 2008).

other acts of theft must also be *pro tanto* wrong. Perhaps there are some occasions on which all things considered we ought to steal. Perhaps one's children are starving. But that does not mean that we ought not to steal. The initial judgement is that the act is *pro tanto* wrong because it is theft, but another judgement is that a different act is overall permitted because it is theft *and* feeding one's starving children. There is no inconsistency. In the second case the *pro tanto* wrongness of theft does not win against other moral norms so as to be the *overall* obligation determined by the conjunctive situation. There might be a third case in which one feeds one's children without theft; that would be better than feeding them with theft. That shows that there is a *pro tanto* wrongness of theft in the second case that is outweighed.⁵

We noted that moral judgements contrast with judgements of funniness and niceness in respect of their aspiration to correctness. Many philosophers think not only that there is no such thing as correctness in judgements of niceness or funniness but also there is no norm of consistency, or at least none that is as strong as that which binds moral judgements. That is, not only is it not the case that there are judgements that we ought to make but it is also not the case that there are judgements that we ought to make given other judgements about a similar situation. Jokes wear thin with repetition – one finds them less funny – and it seems alright for one's judgements about the funniness of jokes to vary over time. And it is alright for one's judgements about the niceness of the fifth doughnut one eats to differ from one's judgement about the first. It is true that there is more strain in allowing directly *comparative* judgements that flout consistency, such as that a joke is not as funny as it used to be, or that the fifth doughnut was not as nice as the first. Then, it seems that we will look for a natural difference. But there doesn't seem to be the same kind of *robustness* in the demand that we bring different judgements into line with each other as there is in our moral judgements. However, if someone thinks that there are robust norms of correctness and consistency in our niceness and funniness thought, then I would not mind much. So long as there is such robustness in our moral judgements.

⁵ 'Particularists' sometimes say that what makes something *pro tanto* right or wrong on one occasion may not do so on another, given different circumstances (Dancy 2004). But they surely cannot mean to exclude all arguments, such as the bribery and animal-pain arguments, that deploy the norm of consistency with dialectical effect. Even particularists think that if there is a moral difference between two cases, then there must be some other difference that explains it.

§ 3. Refining Consistency

The consistency requirement is a weak constraint; but it is not toothless, given only very weak assumptions. It is controversial whether they are substantive moral assumptions. One weak restriction is that spatial and temporal properties are not morally relevant in themselves. So spatial or temporal differences, by themselves, cannot make a moral difference. Given that assumption, the idea of consistency is not trivial for the reason that no two spatially distinct things are ever exactly the same, or because things are never exactly the same at different times.

I shall use the word “natural” in what follows just to mean those properties that are non-normative (or at least properties that contrast with the normative properties that we are considering). This will simplify this discussion.

We can imagine someone asserting that every single natural feature of a thing is morally relevant. So the precise spatial and temporal location of a thing can make a moral difference in itself. Let us call this ‘spatial and temporal extremism’. Someone might even say that being *actual* is morally significant. So two things that are indistinguishable in every natural respect, including spatial and temporal respects, except that they are in different possible worlds, might be morally different. Call this ‘world extremism’. It might be said that these extreme views do *not* abandon the norm of consistency. Instead, it just turns out that it is easy to conform to the norm – extremely easy in fact. Perhaps we should distinguish these toothless forms of consistency from non-toothless (or toothful) forms of consistency – those that have bite.

In my view it is a conceptual requirement on moral thinking that a non-toothless form of consistency requirement holds, and so it is a conceptual requirement that mere spatial, temporal or world location cannot make a moral difference.⁶

⁶ There is a difficulty over the biblical commandments to keep the Shabbat day holy and not work on it. This seems to privilege pure time positioning as making a moral difference. One might think that it cannot be *just* because it is the seventh day that it is special. The biblical reasons for the commandment appeal to two previous events – the creation of the world and the era of slavery in Egypt. But then there is a general problem about how past events by themselves are able to cast a normative straightjacket over the present. The enduring appeal of retributivism about punishment shows that we allow that past situations have moral power over the present (Sher 1987). Even those who shy away from retributivist views about punishment tend to think that we have a duty to make up for past injustices to groups of people.

An argument for the conceptual status of a non-toothless consistency requirement is this. Norms can function to guide behaviour – in a broad sense in which behaviour might include intentions, beliefs, desires and feelings. But they could not do this unless one case could be a *precedent* for another. It must be possible that the natural basis of a judgement can recur so that we should make a similar judgement. (See Hart 1961)

Once we move away from spatial, temporal or world extremism, the requirement of consistency will have *some* teeth. However, teeth can be more or less sharp. What we might call a ‘minimal’ position would be that the morally relevant determining natural properties are complete ‘total’ conjunctive properties, which include absolutely every natural property of a thing – intrinsic and extrinsic – apart from mere spatio-temporal location. But this minimalism is as impractical as the extreme view. For practicality, we must operate with a notion according to which the natural properties are properties that can be repeatably instantiated. So they must fall short of including *all* relational properties. On a non-minimal view, the natural properties are more specific local properties. But once we move away from the minimal view, it becomes a substantive matter how extensive the morally relevant determining natural properties need to be. On the minimal view, the natural similarity between two cases, which binds us to make a similar moral judgement, is *exact* similarity (like that between two ties of the same design from the same factory); but on a non-minimal view, it is just similarity in more restricted respects. Not every relational property counts; and not every intrinsic property counts. For example, if we think an act wrong because it is a stealing, then we might think it irrelevant that it took place in

We think that newspapers sometimes ought to publish apologies. And we think that a medal may be appropriate when someone has saved a life. These issues have echoes in Wittgensteinian’s concerns about linguistic meaning: how can my past use of an expression cast a normative straightjacket over my current use (Wittgenstein 1953)? The general problematic idea is that events at one time make appropriate events at another. In the case of Shabbat, even if we owe the past something, why is it that certain specific days are special and not others? Suppose that everyone were somehow frozen in time for 24 hours (Shoemakers 1984) one thousand years ago so that everyone has been keeping Shabbat on the wrong day. Would that matter religiously? Maimonides’ account of the value of the Shabbat might, I suspect, at least in part, be in terms of the present and future *benefits* of keeping it (Maimonides 1958, part III). If so, actually getting the day right would not matter much. But there may be other respects in which the right day matters. I gather that there is no consensus about this issue among Jewish religious authorities.

Sydney, or on Tuesday, or that the thief was wearing a blue hat, or that the stealing was fast or slow or clever or stupid. Or one might think that the racial group of the thief is not relevant. Steps beyond the minimal view are substantive steps. But that there are some practically applicable requirements of consistency is a conceptual truth. If it were not so, no one would be moved by the bribery and animal-pain arguments: when people see the inconsistency in the combination of the judgements that they hold, they feel pressure to modify one or the other of them.

The requirement of consistency is an essential part of our moral thought. It is partly constitutive of it. That we have this aspiration in our moral thought is theoretically neutral between cognitivism and non-cognitivism. We can derive it from reflection on moral thought – from conceptual analysis. In asserting that the consistency constraint has conceptual status, I seem to be placing myself in controversial territory, since I seem to be committing myself to some kind of fundamental distinction between the conceptual presuppositions of moral thought and ordinary substantive moral commitments, however general and central. But it is not clear that there is too dramatic a cleavage. We can allow that the conceptual and the substantive overlap. In that case, we can say that making any moral judgement at all commits us to certain general and central controversial substantive moral assumptions. The important thing is that so long as we want to make moral judgements, we must recognize the consistency requirement. It is a conceptual requirement, whether or not it is also a substantive moral commitment.

§ 4. Explaining and Justifying Consistency

Since the requirement of consistency is an essential part of normative thought, views that reject it or that cannot account for it have a problem. Such views must embrace an error theory. Either the requirement of consistency is a theoretically neutral part of our thought, which both cognitivists and non-cognitivists must respect, or else it is a central substantive commitment that both must respect.

There are two in principle separable projects here: first that of explaining the fact that we are committed to adhering to the norm of consistency; and second, that of justifying our commitment to that norm (Zangwill 1993).

Cognitivists and non-cognitivists offer different explanations *and* justifications of consistency.

Cognitivists explain and justify our felt obligation to be consistent in our normative judgements in terms of a metaphysical fact about the normative properties that we represent in our judgements: that it is their nature to have a modal structure such that an instantiation of a normative property has necessary consequences for the instantiation of normative properties of other things, at other times, in other worlds. (This is sometimes called ‘supervenience’.) For the cognitivist, these necessitation relations explain and justify the consistency requirement. It is *because* the properties are modally structured in this way that we ought to be consistent in our judgements. Why? Because, if we fail to heed the requirement of consistency, our judgements fail to track the necessitation relations among the normative properties; and that means that we are likely to make judgements that couldn’t all be true.

The non-cognitivist explanation and justification goes in the opposite direction. Our talk in terms of a metaphysical dependence relation between normative and natural properties is a consequence of a requirement of consistency in our thought – *which has some other explanation and justification*. Simon Blackburn and Alan Gibbard have made various interesting suggestions concerning about the source of this requirement in moral thinking given their projectivist or expressivist views. None of these suggestions draws on metaphysical facts about moral properties. (See Blackburn 1985, Gibbard 1990, 2003.)

Blackburn has been particularly resourceful. For a non-cognitivist, moralizing is, as Blackburn says, a ‘practice with a purpose’ (Blackburn 1985, p. 19). If we allowed fickle and inconsistent judgement, our practice of moralizing would not serve its purpose (see also Hume 1742, part 3, book I, chapter 3). One such purpose is the co-ordination our behaviour. We would be unable to rely on those with fickle attitudes. Joint projects would be impossible if people had many inconsistent judgements. And as a consequence, things would be worse for those engaging in the joint activities. Mutual knowledge of consistency in attitude over time enables us to avoid the prisoner’s dilemmas that naturally arise. This can be supplemented with an evolutionary story that shows how prisoner’s dilemmas arise in natural selection, and species or groups that find ways to avoid them do better. It’s good for a species or a group if its members squawk when they see a predator, even though it is bad for the individual. Such dispositions enabled our forebears to act altruistically when it

was against their individual survival prospects but in favour of the group survival prospects. It enabled the group to avoid prisoner's dilemmas. Our tendency to subject ourselves to moral norms might be thought to be a development of such co-ordination dispositions. Co-ordination is a matter of a certain kind of consistency in action between one person and another. It is not such a large step from there to a requirement of consistency in judgements about those actions between one person and another. And it is not such a large step from there to the requirement that the judgements of one person should be consistent with each other. I don't know if such a story will work. It is an interesting suggestion. I think the program has certain problems, but let us put them to one side.

I want to raise a problem for non-cognitivism about norms, which concedes that some such explanation and justification of the consistency requirement is successful.

§ 5. The Regress Argument

As I mentioned before, norms come in varieties or brands. Norms can be moral, epistemological, practical, political, legal, semantic, etiquette, aesthetic, sporting or religious. What I am interested in is the *structural relation* between different normative properties and judgements.

I assume that we can agree that moral properties and judgements count as normative if any do. Morality is a paradigm of the normative.

My argument takes off from the idea that, whatever the source of the requirement of consistency in our *moral* judgements is, this is another norm. We *ought* to be consistent in our moral judgements about what we ought to do (or feel etc.). (As I use the word "norm", it denotes an obligation or permission, not a *judgement* about an obligation or permission.)

Let us leave open the nature of the second-order norm of consistency. Perhaps it is a second-order *moral* norm, as in Blackburn's early account of moral *modus ponens* (Blackburn 1984, which is discussed by Hale 1986, 1989, Schueler 1988, Zangwill 1992b). Or maybe it is a norm of *rationality*.

Whether norms of consistency are second-order moral norms, rational norms, or some other kind of norm, non-cognitivism about morality is crucially *incomplete*. On the face of it, we need an account of the metaphysics of these norms of consistency; and even if that issue makes no sense, we need to know whether judgements of norms of consistency are

cognitive or non-cognitive. For example, if norms of moral consistency are rational norms, we need to give an account of judgements of rationality.

If our judgements of a certain sort are constitutively governed by a norm, we can always, at least in principle, *make it explicit*; and once we have done that, we can inspect that normative commitment and ask questions. We can make *explicit* the norms that are *implicit*; and we can then reflect on their nature.

Let us build on this. There is a second-order norm bearing on moral judgements. But if there is such a norm, we can be aware of it. (This is the idea of possible reflection; see Railton 1984, Velleman 1989, Korsgaard 1995.) We can make judgements corresponding to the norm. We can make second-order judgements of consistency about our first-order judgements. But what about the norms bearing on these second-order judgements? If we judge that a norm of consistency has been flouted in one pair of moral judgements, then we ought to judge the same in another similar pair. If one pair of moral judgements is inconsistent, then any other pair of judgements similarly related is also inconsistent, and we should judge accordingly. Our judgements of consistency between our moral judgements ought to be consistent. And the same applies to our judgements of consistency about our judgements of consistency about our moral judgements. They also ought to be consistent.

I want to show that there is a vicious regress here, but only as long as we keep giving non-cognitive analyses of normative judgements.

We conceive of our normative judgements at any level as being subject to norms of correctness and consistency. That is, I ought to make certain normative judgements rather than others, and I ought to make certain normative judgements given other normative judgements about a natural basis. When I make any normative judgement, I am aware that I *could* make normative judgements at a level above. I need not actually do so, but *there is a permanent possibility of normative reflection*. But we cannot think of these permanently possible higher-level normative judgements as *arbitrary*. We must think of them as right and proper. In this sense, we conceive of the justification of our practice at the N level as being dependent on its being justified at the N+1 level. The norms on one level depend on the norms on the level above. But then some account needs to be given of why the norm at the N+1 level holds.

That account will be either cognitivist or non-cognitivist. If it is cognitivist, then there are simply normative properties of the relevant sort.

Those properties are the source of the norm. (I say more about the realist account in section 8.) But if the account is non-cognitivist, then the only way that norms at the $N+1$ level can support N -level practice is if normative judgements at the $N+2$ level would be justified. This unleashes a regress. For example, if norms of consistency in moral judgements are rational norms, then judgements of rational norms are themselves subject to a norm of consistency – and we need to unearth the source of *that* norm. But for a non-cognitivist, judgements about rational norms are only possible because judgements on the level above *would be* justified. If it is not the case that judgements at the $N+1$ level *would be* justified then judgements at the N -level *are not* really subject to the normative constraints that we assume they are subject to. But those constraints are partly constitutive of N judgements. So our making N judgements would be undercut if $N+1$ level judgements were not justified, and $N+2$ level judgements about $N+1$ level judgements would be undercut if $N+3$ level judgements were not justified. And so on.

The problem iterates. The regress has to come to an end. And it has to come to an end in a cognitive judgement about a brute dependence of a real normative property on a non-normative property. That is, first-order normative practice is constitutively governed by higher-order norms. If there *is* a property corresponding to the higher-order norms then the first-order normative practice is legitimate. Its validity depends on the existence of the higher-order properties. If not, the legitimacy of first-order practice depends on our making certain second-order judgements about first-order practice. But then something must be said about those second-order judgements. And so on.

§ 6. Stopping the Regress?

This is the regress argument for normative realism. I want to address perhaps the most obvious reply to this argument, perhaps the one that is likely to be the most common. We will then be in a position to address some less obvious responses.

The objection is that we can stem the regress by giving a purely causal account of the base level normative judgements. *This is just what we do*, says a Wittgensteinian. It is just a fact about our lives that we play this language-game or engage in that conceptual-practice, where it is part of that game/practice that we think of that practice as normatively regu-

lated. We think of the practice as responsive to norms; but – on this story – we are not *justified* in doing so. Norms only emerge from *within* the language-game/conceptual-practice. They are not something we *impose* on it *from outside*. It is not that we play or practice as we do because we recognize that there are norms that demand it. The response is that at some point we must move from trying to *justify* our language-game or conceptual-practice to being content to *describe* it. We simply practice as we do – we just play a certain language-game which has certain constitutive features (Wittgenstein 1953, 1967). Or perhaps, at some level, we simply have hard-wired dispositions to think as we do (Hume 1742, Goldman 1986). Some normative practices might be grounded on others, but at the base normative level, there is a purely psychological explanation of our practising as we do and of our seeking consistency in our normative judgements.

This response is interesting but I think unsatisfactory. There is an instability in the view that seeking consistency is just something we just do, which can be described and causally explained but not justified. Suppose that there are only two levels of norms. It means that we could adopt a second-order ‘nihilism’, by which I mean that there is no such thing as correctness or incorrectness at the second level. If so, we must accept that it is *not* the case that we *ought* to respect the consistency constraint in first-order normative judgements. We cannot hold that second-order nihilist view while blithely carrying on with first-order practice. For if there is no *reason* to be consistent, then I don’t see how anyone can regard it as a *flaw* to be inconsistent. There would be no reason why we should not abandon the consistency requirement. We should view it as a *pathological tendency*, which makes no real claim on us. But then the integrity of first-order judgement is jeopardized.

Consider the rules of a game. There are rules of chess, defined as norms of correct (permissible) moves for that game, and moves that win or lose. But those rules are arbitrary. (Well, they are not *arbitrary for chess*; perhaps chess *must* have the rules it does. But there can be other games with different rules. For example in “cylindrical chess”, one can go off one side of the board and come back on the other, as if the board were a cylinder.) However, if we reflect on the rules, we can see that they admit of a justification, which is that playing that game (chess rather than ‘schmess’) is interesting or pleasurable. The rules of the game only bind us given that they make playing the game worthwhile. Stupid games with stupid rules make no claim on us. If we play a game, we typically

think that it is worth doing so – unless we are just playing to please another, or we are under duress. Suppose someone (but not me) invents a stupid pointless game according to which one must walk in a zigzag path down the street. No one else knows of this game. Is it then the case that *I*, who have never heard of the game, ought to walk in that fashion? Well in *a* sense, I ought to, according to the game (compare Foot 1972); but if that game is entirely pointless then I have no reason to play that game rather than the non-zigzag game according to which one must walk in a non-zigzag path down streets. Unless there is a second-order reason to play the game, and one game rather than another, the rules or ‘norms’ of the game are not binding. Similarly, abiding by a consistency constraint is one of the rules of the ‘game’ of making moral judgements. But there must be a reason to play a game with such a constitutive rule. Consistency is a constitutive rule of moralizing just as there are constitutive rules of chess or the zigzag game. But when we make moral judgements, play chess, or the zigzag game, and thereby assume the burden of adhering to the constitutive rules, we do so only because, at least in our reflective moments, we think that the activity has a point. In the chess case, that might be pleasure. In the case of making moral judgements, it might be that our judgements depict a kind of properties. Or it might be that our first-order practice is subject to second-order assessment. But then we can ask: what drives and legitimizes that second-order assessment?

Consider etiquette. Some etiquette has a point. But some is notoriously arbitrary. If we were to arrive at the second-order realization that some particular requirement of etiquette lacks any rationale, that inevitably *infects* first-order practice. We can’t go on as before, feeling bound by the norm. Even if we find that for psychological reasons we cannot abandon the practice, we will come to see it as a pathological compulsion, but not as something that we ought to respect. Similarly, if we come to have the second-order belief that human beings have a near universal inclination towards theological beliefs only because those beliefs provide social cohesion or solace in the face of death, we can hardly return and engage in our first-order theological thought entirely unaffected. Similarly, if there is only a causal account of our commitment to consistency in normative judgements, that undercuts our justification for obeying the consistency constraint in our normative judgements, and embracing the purely causal account would undermine our commitment to making normative judgements.

Suppose the emotivist or projectivist says that seeking consistency is necessary if moralizing is to fulfil its role of enabling human beings to co-ordinate their behaviour, which enables things to go well. Seeking consistency has good effects. It might be argued that this would show that we *ought* to seek consistency. It is better if we do. Seeking consistency has benefits. But this seems not to be enough to support *my* adherence to the consistency constraint if I become aware of it in reflection. Am I to think “The reason I should bring my judgements into line with each other is that if many people do this, society will flourish”? That shows that it is good if *people in general* seek consistency. But it is not a reason *for me* to be consistent. So even though this view may not be nihilistic on the social level, it is individualistically nihilistic about the consistency requirement. Such a norm of consistency does not bind me. I might even be glad that *others* have such tendency. I could view it as a beneficial pathological tendency. The realist answers the question “Why ought I to be consistent?” by pointing to the fact that it is what the properties dictate. The non-cognitivist answers that it is a good idea if we all have a tendency to respect consistency, that is, if we tend to *think* that there is such a norm. But why should this be convincing to *me*? Why should that bring *me* to see that certain combinations of my judgements or attitudes are such that *I* should approve or disapprove of them?

Second-order nihilism is inconsistent with what is necessary for first-order practice. Second-order nihilism means that there is no reason not to cease making the first-order normative judgements and that there is no reason to seek consistency among those judgements. The lack of justification at a higher level normative level *trickles down* to all lower levels of normative judgement.

So a purely causal explanation of our language-game or conceptual-practice is not sufficient. In Kant’s terms, we need a *de juris* ‘deduction’ of the norm of consistency, not (or not just) a *de facto* explanation.

§ 7. Two Other Possible Replies

It might be suggested in reply that a projectivist could hold a *coherentist* view, so that there is a mutually supporting, non-linear holistic web of norms, each of which derives support from the others member of the set. But this is illusory. For the gravitational metaphor of support is given sense by the ‘level’ or ‘order’ of justificatory support, and that is a matter

of whether a norm applies to other normative judgements. A first-order norm does not. A second-order norm is one that bears on first-order normative judgements. A third-order norm is one that bears on normative judgements about normative judgements that are not about normative judgements. And so on. The only source of normative support for a normative judgement is from the level above. The norms at any level can only be supported by a norm at the level above. A fifth-level norm cannot be supported by a first-level norm; it must be supported by a sixth-level norm. So a coherentist model of normative support will not work. Normative support is hierarchical, not holistic.

Another possible reply would be to say that some explanation and justification applies at every level at once. That is, there might be some norm – which the non-cognitivist can explain and justify – that enjoins consistency among normative judgements of whatever level. If so, then it might be argued that although there is a non-cognitivist regress, it is not a vicious one. This reply also fails. Perhaps there could be some such norm, which had such an infinite reach and which applied to itself. Nevertheless, this sideways-on variable-level norm would stand in need of justification. There would be work to do in justifying *this* norm. So the regress argument would be back on track. For a normative non-cognitivist, norms are never self-justifying. Their justification derives from elsewhere. So even if there are neutral-level norms of consistency, it does not help to evade the regress argument. Such a normative principle would be universally quantified in form: it would say that for all objects of evaluation (whether judgements or not) we ought to be consistent in our judgements about them. Perhaps. But there are still questions about the legitimacy of that neutral-level universally quantified normative principle.

§ 8. Consistency and Realism

A realist can supply that justification because the realist appeals to normative facts. Someone who engages in a sort of normative judgement is aware that they thereby take on the burden on respecting a norm of consistency in those judgements. Why? Not from some pathological tendency, but because they understand that they ought. And the realist has an account of why they ought. For a realist, normative thinkers know that when they are inconsistent in their judgements they offend a norm

dictated by the properties and their modal structure, and that means that their pattern of judgements comes to possess a negative second-order normative property. The nature of normative properties dictates their modal structure.

It is the normative properties with their essential modal structure that justifies us in conforming to that modal structure when we make our normative judgements about those normative properties. The necessities that bind moral property instantiations of different things, times and worlds, are that in virtue of which the consistency norm for those normative judgements holds. That modal structure is that on which the second-order normative properties depend. That modal structure is the truth-maker of our judgements of consistency. It is where justification comes to an end.

In brief: so long as a non-cognitivist wants to avoid convicting ordinary normative thought of an error, it must respect the consistency requirement. But then *that* norm of consistency needs analysis. If the analysis of that norm is non-cognitivist, we will be forced up another level. At some point we must stop: some normative judgements must be cognitivist. The consistency requirement, at some level, stems from the fact that our judgements are about genuine properties, and it is in the nature of these properties to be modally structured (something that flows from the dependence of normative on non-normative properties). Respect for consistency is licensed by necessities binding normative property instantiations. You cannot be a non-cognitivist all the way up. This is not an argument for *moral* realism. The argument is that you cannot generalize non-cognitivism. At least some normative judgements must be cognitive. The justification of norms must come to an end in cognitive judgements about the instantiation of normative properties. We ought to be consistent in *those* normative judgements simply because the normative properties have the modal structure they do.

The justificatory directions are different for cognitivists and non-cognitivists. For a non-cognitivist, justification must flow down from the upper levels. But for cognitivists and realists, the source of justification for upper-level judgements derives from the base-level normative properties and their modal nature. The justification flows up from the base level for a cognitivist and realist, and downward for a non-cognitivist.

A qualification: the regress argument does not show that we cannot *abandon* a certain sort of normative judgements – that if that kind of normative properties is real, the normative judgements will have some

automatic hold on us (Foot 1972). People sometimes fail to heed certain normative demands, and this is no less true if the norms are realistically construed. The point is just that in so far as we make normative judgements about normative properties, we know that we should heed the modal structure that is essential to them.

A worry: are realists stuck with their own regress? I can be reflective. I know that moral judgements about moral properties are vulnerable to second-order assessment. I can form the thought that I ought to be consistent. But, on the realist theory, this judgement invokes another normative property. And the same will be true of the judgement that I ought to be consistent about my judgements of consistency. Once we are normative realists, it seems that we must embrace an infinite series of nested normative properties! If we ascend to any arbitrary level of reflection, a normative property binds us at that level. There is no escape. There is a normative property of the 115th level that attaches to the 114th level of reflection. For the realist, there is an infinite tower of norms. Any normative property has an infinite number of accomplices, like reflections in a mirror of reflections in a mirror.

However, I think that this is less worrying than it might appear. For these higher-order normative properties are *hypothetical* normative properties. For any normative property, *if someone were* to assert its instantiation, then the way that judgement *would* fit with other possible judgements about similar normative properties *would* mean that those judgements *would* possess higher-order normative properties. So *if we were* to reflect at, say, the 227th level, those judgements *would* possess certain yet higher-order normative properties. The realist can be happy that all these hypothetical normative properties exist, although their antecedents are never all instantiated, except perhaps by God if He is infinitely reflective. Is this unacceptably strange? Well there is nothing queer about the base level normative properties since they are comfortably realized in ordinary natural properties. But if that is right, second-order normative properties are realized in metaphysically unproblematic first-order normative properties. And so on. So while higher-order normative properties may be complex, for a realist they are not queer. (Such an idea seems generally unproblematic: a determinate property such as being 6 foot tall, determines the infinite disjunctive property of being between 5 and 7 foot tall. There is no problem with the idea that more basic level properties can determine all sorts of complex properties.)

Why is this a virtuous regress, while the non-cognitivist regress is a vi-

cious one? The reason lies the direction of normative dependence. For non-cognitivists, norms at level N are legitimate *only if* there is some norm at level N+1. But for a realist, norms at level N *determine* a norm at level N+1. This is a significant difference. For the non-cognitivist, N level norms depend on N+1 level norms; but for the realist, N+1 Level norms depend on N level norms. Hence only the non-cognitivist regress is vicious, since, from the realist point of view, the N level norm is there, fixed by the natural facts, and so there is nothing to prevent higher-level norms unfolding. While for a non-cognitivist, there is no legitimate level onto which to build; the legitimacy of any level depends on how it is with the level above.

Coda

Does this argument establish normative realism? It establishes *cognitivism* about some normative judgements or other. But this leaves open the possibility of an error theory according to which there are no normative facts of *any* sort. My view is that there are no decisive reasons – metaphysical or epistemological – to deny the existence of genuine normative facts. The arguments against normative realism are not conclusive. But it is also true that we need a reason to believe in the facts that normative cognitivism commits us to. (Similarly, we might show that theological judgements are cognitive and that the arguments against God's existence are no good; but showing that God exists would be another matter.) I have tried to establish normative cognitivism here – no more.

For a realist, we realize that the properties that we are talking about mean that if we engage in a normative conceptual-practice or playing a normative language-game, then our judgements are *vulnerable* to second-order normative assessment. That is, we know that our combinations of judgements may come to possess positive or negative second-order normative properties that dictate that we should modify our first-order judgements. We *fear* the second-order normative properties to which our first-order practice necessarily makes us vulnerable. The ticket of entry to the normative language-game is potential guilt! We practice *in fear* of the norm.⁷

⁷ This paper was given as talk at the University of British Columbia, Northern Arizona University, the University of Illinois at Chicago, The Hebrew University of Jerusalem, Nottingham University, and the British Society of Ethical Theory. I am grateful for questions asked by members of the audiences at these occasions, and also for a discussion with Jimmy Lenman.

References

- Anscombe, Elizabeth 1959: *Intention*, Oxford.
- Blackburn, Simon 1984: *Spreading the Word*, Oxford.
- 1985: Supervenience Revisited, in: Blackburn 1993.
- 1993: *Essays on Quasi-Realism*, Oxford.
- 1998: *Ruling Passions*, Oxford.
- Dancy, Jonathan 2004: *Ethics Without Principles*, Oxford.
- Foot, Philippa 1972: Morality as a System of Hypothetical Imperatives, in: *Philosophical Review* 81, 305–316.
- Gibbard, Alan 1990: *Wise Choices, Apt Feelings*, Oxford.
- 2003: *Thinking How to Live*, Oxford.
- Goldman, Alvin 1986: *Epistemology and Cognition*, Oxford.
- Hale, Bob 1986: The Compleat Projectivist, in: *Philosophical Quarterly* 36, 65–84.
- 1989: Quasi-realism and Truth, in: John Haldane and Crispin Wright (eds.), *Reality, Truth and Projection*, Oxford.
- Hart, Herbert 1961: *The Concept of Law*, Oxford.
- Hume, David 1742: *Treatise on Human Nature*, Oxford 1948.
- Kant, Immanuel 1928: *Critique of Judgement*, transl. James C. Meredith, Oxford.
- Korsgaard, Christine 1995: *The Sources of Normativity*, Cambridge.
- Mackie, John 1977: *Ethics: Inventing Right and Wrong*, Harmondsworth.
- Maimonides, Moses 1958: *Guide For the Perplexed*, London.
- Railton, Peter 1984: Alienation, Consequentialism and the Demands of Morality, in: *Philosophy and Public Affairs* 13, 134–171.
- Schueler, Fred 1988: Modus Ponens and Moral Realism, in: *Ethics* 98, 492–500.
- Sher, George 1987: *Desert*, Princeton.
- Shoemakers, Sydney 1984: Time Without Change, in: *Identity, Cause and Mind*, Cambridge.
- Singer, Peter 1975: *Animal Liberation*, New York.
- Velleman, David 1989: *Practical Reflection*, Princeton.
- Wittgenstein, Ludwig 1953: *Philosophical Investigations*, Oxford.
- 1969: *On Certainty*, Oxford.
- Zangwill, Nick 1992a: Quietism, in: *Midwest Studies in Philosophy* 17, 160–176.
- 1992a: Moral Modus Ponens, in: *Ratio*, 5, 177–193.
- 1993: Quasi-realism, Justification and Explanation, in: *Synthese* 97, 287–296.
- 1994: Moral Mind-Independence, in: *Australasian Journal of Philosophy* 72, 205–219.
- 1998: Direction of Fit and Normative Functionalism, in: *Philosophical Studies* 91, 173–203.
- 2006: Moral Epistemology and the Because Constraint, in: Jamie Dreier (ed.), *Contemporary Debates in Moral Theory*, Blackwell.
- 2008: Moral Dependence, *Oxford Studies in Metaethics* 3, 109–128.