

## **Mental Agency and Metaethics**

Matt Evans (NYU)

and

Nishi Shah (Amherst College)

1. Normal human adults have the capacity to reflectively step back from our inclinations to believe certain propositions and ask ourselves whether to believe those propositions. When this reflection concludes in a belief, we have exercised our agency over that belief, even if what we end up believing is what we already were inclined to believe. We have not been mere spectators of the formation of our beliefs; we have determined for ourselves what to believe.

Of course, not every kind of reflection on our beliefs involves exercising such agency. When I reflectively conclude that it would be desirable to believe some proposition that I am inclined to believe, my reflection does not itself issue in my believing that proposition. But when I reflectively conclude that there is sufficient evidence that some proposition is true, my reflection normally does issue in my believing that proposition.

Why is only the latter kind of reflection—reflection on evidence—capable of issuing in belief and therefore the only kind of reflection capable of constituting our agency over our beliefs? One of us has argued that the answer to this question lies in the fact that reflection that is aimed at concluding in belief—reflection aimed at exercising one’s agency over one’s beliefs, is structured by acceptance of the following norm:

Believing that  $p$  is correct iff  $p$ <sup>1</sup>

The reason that such reflection is governed by this norm is that the norm is contained, so to speak, in the very concept of belief. Anyone who fully possesses the concept of belief

---

<sup>1</sup> See Shah (2003).

must apply this norm, at least implicitly, when exercising the concept. Aiming to arrive at a belief by means of reflection, and thus aiming to be the agent of one's beliefs, requires that one be guided by the norm of truth. Of course, aiming to arrive at a belief doesn't usually involve explicitly asking oneself *what to believe*. Nonetheless, one cannot aim to form a belief unless one conceives of one's activity as belief-formation. That one conceives of one's activity as determining what to believe, and thus as governed by the norm of truth, explains the structure of one's reflection. That there is evidence that p is pertinent to determining whether, by believing that p, one would be complying with the norm; that it would be desirable to believe that p, however, is not by itself pertinent to determining whether, by believing that p, one would be complying with the norm. Given the way that the norm of truth structures doxastic reflection, the only way to exercise reflective agency over our beliefs is by deliberating about the evidence.

Similarly, when we step back from our intentions and our non-appetitive desires and attempt to determine what to intend or desire, we find that only reflection on considerations relevant to the desirability of the objects of the intentions or desires—the acts intended or things desired— can move us to revise or retain them. As the toxin case makes vivid, reflection on the desirability of an intention, as opposed to the desirability of the act intended, cannot by itself issue in an intention. Similarly, being offered money to desire a plate of mud makes it desirable to desire a plate of mud, but because it does not make having the plate of mud itself desirable, reflecting on this fact does not by itself issue in the desire for it.

As with belief, we think that the explanation for these facts about our capacities to exercise reflective agency over our intentions and desires is a normative one. Reflection aimed at determining what to intend is structured by an agent's acceptance of the following norm:

Intending to phi is correct iff phi-ing is desirable.<sup>2</sup>

---

<sup>2</sup> See Shah (2008).

Reflection aimed at determining what to desire, for those desires that are capable of reflective governance, is structured by an agent's acceptance of the following norm:

Desiring that p is correct iff p is desirable.

The reason that effective reflection is governed by these norms is that these norms are contained in the very concepts of intention and desire that are used to frame the questions that reflection seeks to answer. Asking what to intend or what to desire requires exercising the concepts of intention and desire, and thus requires applying these norms. That one conceives of one's activity as determining what to intend or desire, and thus as being governed by the corresponding norm, explains the structure of one's reflection. That phi-ing is desirable is pertinent to determining whether, by intending or desiring to phi, one would be complying with the norm for intention or desire; that it would be desirable to intend or desire to phi, however, is not by itself pertinent to determining whether, by intending or desiring to phi, one would be complying with the norm. Given the way that the norm of action structures practical reflection, the only way to exercise reflective agency over our intentions or desires is by deliberating about the desirability of the acts intended or objects desired.

Beliefs, desires, and intentions, are the central constituents of human agency. To the extent that we have reflective control over these states we are self-governing creatures. But if our reflective agency is as we described, then self-governance involves normative guidance. Anyone who is capable of exercising reflective agency over these mental states must conceive of them in normative terms, accepting that certain norms govern these states and other norms do not. Reflective control over these states requires that one successfully regulate them in conformity with these norms. Surprisingly, we shall argue, this feature of our reflective agency is inconsistent with any of the prominent forms of normative anti-realism: anyone who accepts the description of reflective agency that we have given must reject normative anti-realism, whether it takes the form of an error-theory, constructivism, or expressivism. After arguing for this conclusion, we conclude

with a discussion of where this leaves us: can we conclude that if our description of the agency we exert over our minds is correct, then realism must be true?

A crucial point in what follows is this. According to the conception of agency just described, one is committed to the relevant norms just by applying the concepts of belief, desire, and intention. Insofar as the metanormative theorist understands the people who he is interpreting as agents, he must apply the concepts of belief, desire, and intention to them. To do this, it is not enough though that he attribute the commitment to the relevant norms to them; he must also endorse those norms himself. This is because the norms are built into the very concepts of those mental states. A commitment to these norms thus is involved in any unconditional application of those concepts. Insofar as he conceives of those he is interpreting as possessing those mental states, which he must if he is to conceive of them as agents at all, he is committed to the applicability of these norms to their mental states. We will use the label ‘Normativity of Attitudes’ to refer to the thesis that each type of attitude ascription involved in attributions of agency involves a commitment to a corresponding normative judgment.

2. Making normative judgments is part of the stream of human life; it is something we do. The metanormative theorist has the primary task of explaining this practice of making normative judgments in terms of its place in the life of a certain kind of creature, human beings, and thereby situating normative judgments within the order of nature.<sup>3</sup>

Potential explanations of this practice can be usefully organized by their relation to John Mackie’s famous argument for a moral error theory. Mackie argued that the objective purport of moral discourse requires that obligations be part of the mind-independent fabric of the universe—that they be metaphysically independent from anyone’s actual or potential attitudes, but that the universe, at least as it has been disclosed to us by the natural sciences, contains no such “queer” entities,<sup>4</sup> Nor have the

---

<sup>3</sup> Of course we do not mean to deny that there are other tasks that one might assign to a metanormative theorist. Our intention is to describe a recognizable philosophical project that many 20<sup>th</sup> century philosophers who we now label ‘metaethicists’ have engaged in.

<sup>4</sup> See Mackie (1977).

natural sciences disclosed any perceptual or other capacities that would allow us to detect the presence of such mind-independent ‘morons’<sup>5</sup>, even if they did exist. Mackie thus claimed that moral judgments are not what they sometimes appear to be: warranted, true statements of objective fact.

Generalized to cover all normative judgments, Mackie’s argument has the following structure:

1. If normative judgments are factual judgments, then correct normative judgments accurately represent attitude-independent normative facts.
2. Normative judgments are factual judgments.
3. There are no attitude-independent normative facts.
4. So there are no correct normative judgments. (1-3)

Assuming that this argument is valid, the metanormative theorist has four different strategies to choose from: constructivism, expressivism, realism, and an error-theory. The constructivist denies premise one. He believes that normative facts are metaphysically grounded in facts about the attitudes of some class of agents, and thus denies that in order for a normative judgment to be correct it must represent a mind-independent normative fact. The expressivist denies the second premise. He believes that the function of normative utterances is to give voice to non-representational practical attitudes, not to depict a realm of normative facts. Therefore he denies that the normative judgments expressed by normative utterances are factual judgments. The realist denies the third premise. He believes there is a realm of mind-independent normative facts that make at least some of our normative judgments true. Lastly, the error-theorist such as Mackie accepts the entire argument. He thinks that, like other defective practices such as witchcraft, practices that presuppose the truth of normative judgments are defective because normative judgments are about a realm of facts that does not exist. To make a normative judgment is to mistake a subjective feeling for objective reality.

---

<sup>5</sup> See Dworkin (1996).

3. Suppose for the sake of simplicity that the order of nature is a system of *facts* standing in a transitive and asymmetric relation of *ground*. This relation would be the metaphysical dependence relation we want to express when we say things like “what makes it the case that p is that q” or “it is in virtue of the fact that q that p” or “the fact that p is prior in the order of explanation to the fact that q.”<sup>6</sup> Some examples might be: “what makes it the case that this is the right thing to do is that it will have the best consequences” or “this experience is painful in virtue of the fact that it stimulates aversive behavior” or “this disjunction’s being true is posterior in the order of explanation to one of its disjunct’s being true.”

It is important to point out that this relation, as we understand it, is not fundamentally *modal*. One might hold, for example, that Plato belongs to the singleton set {Plato} in virtue of the fact that he exists even though it is metaphysically necessary that Plato belongs to the singleton set {Plato} if and only if he exists. Or one might hold *both* that what makes the act of serving the gods pious is that it participates in the form of piety *and* that, necessarily, the act of serving the gods is pious if and only if it participates in the form of piety.<sup>7</sup>

The metanormative theorist’s primary task is to show how normative judgments and normative facts, if there are any, stand in the order of nature. There are some constraints on how this task can be performed, however. One constraint has to do with the category of *mind-dependent* facts. Let us call a fact F *fundamental* if there is no other fact G such that F is grounded in G. In that case it seems correct to suppose that if there are any

---

<sup>6</sup> On the relation of metaphysical ground, see especially Fine (1995, 2001), Schaffer (2009), Rosen (forthcoming), and — for a useful overview — Correia (2008).

<sup>7</sup> It is also worth noting that the relation of metaphysical ground, as we understand it, is not equivalent to the relation of *conceptual* ground. For clearly it is not necessary to establish that certain concepts stand in a relation of ground to each other in order to establish that *the properties these concepts refer to* stand in a relation of ground to each other. (Here we are assuming that non-empty concepts refer to properties, and that the property of being F is grounded in the property of being G if and only if, for any x, if x is F, then x is G in virtue of the fact that x is G. For an interesting argument against the second assumption, see Rosen (forthcoming).) One might deny that the concept of being in pain is grounded in the concept of being in a certain neural state, but accept that the *property* of being in pain is grounded in the *property* of being in a certain neural state. Or one might deny that the concept of being good is grounded in the concept of being pleasant, but accept that the *property* of being good is grounded in the *property* of being pleasant.

*fundamental* normative facts, then all of them are mind-independent. For a fact is mind-dependent — or more specifically, attitude-dependent — just in case it is grounded in an attitudinal fact (i.e., a fact about someone's attitudes). But if it is grounded in an attitudinal fact, then it is grounded in *some* fact, and so is not fundamental. Therefore, if a normative fact is fundamental, then it is mind-independent.

The challenge for the realist is to establish that the facts that make some of our normative judgments correct are attitude-independent normative facts.<sup>8</sup> The nonreductivist realist claims that these normative facts are primitive, whereas the reductive realist claims that they are grounded in some non-attitudinal non-normative facts. For example, assume that Tom judges that lying is wrong and that his judgment is correct. The nonreductive realist claims that what makes Tom's judgment correct is the fact that murder is wrong, which is either metaphysically fundamental or grounded in other normative facts which are themselves metaphysically fundamental. The reductive realist also claims that the fact that lying is wrong makes Tom's judgment correct, but unlike the nonreductive realist, he claims that this fact is itself grounded in some non-normative non-attitudinal facts.

The challenge for the anti-realist is to establish that at the most fundamental metaphysical level, assuming there is one, there are no normative facts. Either there are no normative facts at all, as the expressivists and error-theorists claim, or all such facts are grounded in attitudinal facts and are thus not fundamental, as constructivists claim. Both the expressivist and error-theorist deny that there are normative facts that make Tom's judgment that lying is wrong correct. The error-theorist claims that this means that Tom's judgment is false, whereas the expressivist claims that Tom's judgment is, at the most fundamental level, non-propositional. At that level, the very idea of a normative proposition or fact makes no sense, and it is a mistake to ask, at that level, whether Tom's judgment is true or false. Unlike the error-theorist and expressivist, the constructivist can

---

<sup>8</sup> Some may think that the realist needs to establish a stronger claim, which is that the normative facts that ground correct normative judgments are themselves fundamental. For our purposes, we can leave this question open. Whether to count non-constructivist reductive accounts of normative facts as forms of realism is not our concern in this paper.

accept that Tom's judgment that lying is wrong is correct. If he does so, he will claim that the normative fact that makes Tom's judgment correct is grounded in some attitudinal facts.

In order to place normative judgments and normative facts, if there are any, in the order of nature, the anti-realist must give an account of normative judgments and facts (if there are any) according to which, they are identical to, or are grounded in, non-normative facts. Now it might be thought that the anti-realist can satisfy this constraint only if he can avoid making (or committing himself to) any normative judgments in giving his account.<sup>9</sup> But it is obvious that, if the Normativity of Attitudes is correct, one cannot give an account of normative judgments without making a normative judgment. Here is why. Nobody denies that normative judgments are attitudes; what people disagree about is *what kind of attitudes they are*. Are they cognitive states, such as beliefs about normative facts? Or are they non-cognitive states such as a desires or intentions? An account of normative judgment thus must be an account of some particular kind of attitude. And if the Normativity of Attitudes is correct, then claiming that a judgment is an attitude commits one to whatever normative judgments are required to make that attitude ascription. So even to claim that normative judgments are attitudes entails committing oneself to some normative judgment or other. Anyone attempting to give an account of the attitudes that constitute normative judgments thus commits himself to at least one normative judgment. So if the Normativity of Attitudes is true, then one cannot give an account of normative judgment without committing oneself to at least one normative judgment.

But does this show that anti-realists— constructivists, expressivists, and error-theorists— cannot succeed in giving an account of normative judgment? No. For recall that the basic anti-realist stance is simply that there are no attitude-independent normative facts. Thus the only basic constraint that the anti-realist must satisfy, in giving account of normative judgment, is to avoid making any judgments *that would, if correct, entail that*

---

<sup>9</sup> Cite Drier's objection to Blackburn's expressivism as an example of an argument that seems to rely on such an assumption.



*there are some attitude-independent normative facts.* Call this **the coherence constraint**. So the real issue is whether or not an anti-realist can meet the coherence constraint if the Normativity of Attitudes is correct.

Notice that the Normativity of Attitudes does not entail that there are any normative facts or even that making a normative judgment is a cognitive act that commits one to the existence of normative facts. Interpreting the normative commitments involved in attitude ascriptions is a job for the metanormative theorist. Realists and error-theorists will interpret these commitments as judgments about mind-independent normative facts; the realist claiming that some of these judgments are true and the error-theorist denying that any of them are true. Constructivists will interpret the normative commitments involved in attitude ascriptions as judgments about mind-dependent normative facts, and expressivists will interpret them as non-cognitive attitudes that are neither true nor false.

Although no particular metanormative claims are contained in the Normativity of Attitudes, we will argue that only the realist can both accept the Normativity of Attitudes and meet the coherence constraint. No other metaethical position can provide a stable interpretation of the normative commitments involved in attitude ascriptions.

4. Start with the error-theorist. According to him, there are no correct normative judgments because any normative judgment, if correct, accurately represents some attitude-independent normative fact, and there are no such facts. What must one do to arrive at a fully general error-theory of normativity? It is not sufficient that one assert that propositions attributing normative properties are all false. To see this, think about something that we are all error-theorists about, witch-discourse. We think that people who believed (or continue to believe) that certain people are witches were mistaken. To arrive at this claim, we must have been able to attribute these beliefs to certain people. Unless we were able to identify practices in which some people held such beliefs, there would be no sense to be made of our claim that anyone has ever been *mistaken* about the existence of witches. At a minimum, to arrive at an error-theory about normativity one therefore must do the following two things:

- 1) Attribute normative judgments.
- 2) Demonstrate that these judgments are systematically false.

The question is whether these two conditions can be jointly fulfilled if the Normativity of Attitudes is correct. Let us focus on a particular normative judgment:

- 3) Lying is wrong.

In what does such a judgment consist? According to the error-theorist, the judgment that lying is wrong expresses a belief about an attitude-independent normative fact.

What then is required to ascribe such a belief? Let us focus on a particular belief ascription of this type:

- 4) Tom believes that lying is wrong.

What is involved in making this judgment? According to the Normativity of Attitudes, ascriptions of belief require making normative judgments: to ascribe a belief one must judge, implicitly at least, that the mental state so classified is correct if and only if it is true. Judging that Tom believes that lying is wrong thus commits one to the normative judgment that:

- 5) Tom's belief that lying is wrong is correct if and only if lying is wrong.

How does the error-theorist interpret this commitment to a norm of correctness? Since it takes the form of a judgment, he must claim that the commitment to a norm of correctness is a belief whose content is that the belief one has ascribed has the normative property of being correct if and only if its content is true. Thus, ascribing to Tom the belief that lying is wrong would require believing the following normative proposition:

- 6) Tom's belief that lying is wrong is correct if and only if lying is wrong.

If one were an error-theorist about normative judgments—believing amongst other things that there is no normative property of wrongness and thus no true propositions attributing this property to anything—one therefore would be committed to believing the following two inconsistent claims:

- 7) There are no true normative propositions.
- 8) Tom's belief that lying is wrong is incorrect.<sup>10</sup>

The latter claim is itself a normative proposition that must be false if the former claim is true. And this case obviously generalizes to attributions of all normative beliefs. Attributing normative beliefs to others commits one to believing normative propositions oneself.

The error-theorist cannot demonstrate that all of our normative judgments are false without attributing some normative judgments to us. If the Normativity of Attitudes is true, then, by the error-theorist's own lights, his attributions of normative judgments commits him to the existence of attitude-independent normative facts about the correctness of the beliefs that constitute those judgments. Thus, the error-theorist cannot attempt to demonstrate that all normative judgments are false without making a normative judgment that, if correct, implies the existence of attitude-independent normative facts. So if the error-theorist is correct, he is incorrect: he cannot satisfy the coherence constraint.

---

<sup>10</sup> The inconsistency can be made more transparent if we frame the two claims the error-theorist is committed to in terms of properties. Here are the two claims reformulated:

- 7\*) There are no normative properties.
- 8\*) John's belief that lying is wrong has the normative property of being incorrect.

Those whose metaphysical scruples cause them to recoil at talk of the existence of properties obviously will reject this formulation of the error-theorist's commitments.

5. Is either the constructivist or the expressivist in any better position to satisfy this constraint? Yes. For the constructivist holds that correct normative judgments, including those normative judgments involved in attributing normative judgments to others, accurately represent *attitude-dependent* normative facts, and the expressivist holds that normative judgments, including those normative judgments that are involved in attributions of normative judgments, are *non-factual* judgments. So neither the constructivist nor the expressivist, in identifying the practice of making normative judgments, needs to make any judgments that would, according to his own account, accurately represent attitude-independent normative facts. Unlike the error-theorist, then, the constructivist and the expressivist can satisfy the coherence constraint.

Nonetheless, the Normativity of Attitudes poses an obstacle for both of these views. Consider first the constructivist. Like the expressivist, the constructivist needs to show that all norm-judgmental facts are grounded in non-normative facts. But unlike the expressivist, he thinks that at the most fundamental explanatory level, normative judgments are beliefs, and thus are on the cognitive side of the cognitive/non-cognitive divide. If one is not an error-theorist, the claim that normative judgments are beliefs commits one to the existence of some *normative* facts. For if normative judgments are beliefs, then correct normative judgments are made correct by the normative facts that make their contents true. This means that if there are any correct normative judgments, there are normative facts. Since the constructivist is not an error-theorist, he believes that at least some normative judgments are correct. Therefore he is committed to the existence of some normative facts. But unlike the realist, he claims that normative facts are grounded in attitude-dependent facts.

Unfortunately, the constructivist cannot make good on this claim if the Normativity of Attitudes is correct. In claiming that normative facts are attitude-dependent, the constructivist commits himself to grounding normative facts in attitudinal facts. A popular constructivist strategy, for example, attempts to ground facts about what an agent has reason to do in facts about what that agent would *desire* under certain conditions. To

claim that normative facts are grounded in such desires is not merely to commit oneself to claims of the following type:

S has reason to phi iff S desires to phi under conditions C

It may be that there are states of affairs that our desires co-vary with when we are under certain non-trivially specified conditions. If so, such bi-conditionals allow us to perspicuously represent these connections. That such bi-conditionals express truths about certain states of affairs does not imply that the facts corresponding to those states of affairs are grounded in the relevant attitudinal facts.<sup>11</sup> To say that these facts are grounded in attitudinal facts would be to provide a possible explanation of why these bi-conditionals hold. This stronger claim that the attitudinal facts ground the normative facts entails that the grounding attitudinal facts — in this case the desires of a certain set of agents — must themselves be metaphysically prior to the normative facts.

But if the Normativity of Attitudes is correct, the constructivist cannot coherently maintain that attitudinal facts, including facts about what certain agents desire, are metaphysically prior to normative facts. Remember that according to the Normativity of Attitudes, it is a conceptual truth about desire that:

The desire to phi is correct iff phi-ing is desirable.

If correctness is a property, as the constructivist supposes, then it follows from this conceptual truth that a desire to phi has the *essential* property of being correct iff phi-ing is desirable. In the language of facts, this means that desire facts are composed of

---

<sup>11</sup> It is difficult to know how to understand a no-priority view, such as that endorsed by McDowell (1985) and Wiggins (1987), which claims that neither attitudinal facts nor normative facts are more fundamental than the other, and that neither kind of fact is grounded in any other facts. Their view entails that normative facts are not grounded in attitudinal facts or any other facts, which might make it appear to be a realist view. But a realist claims that some attitudinal facts, those that correspond to correct normative judgments, are grounded in normative facts. This is a claim that the no-priority theorist rejects, but it is difficult to see how to make sense of this. If normative facts are metaphysically fundamental, as this view implies, it is hard to see how to avoid the claim that correct normative judgments, which are attitudes, are grounded in some of those normative facts.

normative facts. In claiming that reason facts are grounded in desire facts, the constructivist thus is committed to claiming that reason facts are grounded in facts that are themselves composed of correctness facts.

To make this clear, let us focus on an example. Assume that Tom's judgment that he has reason to lie is true. The constructivist proposal we described says that his judgment is made true by the fact that he would desire to lie under ideal conditions. If the Normativity of Attitudes is true, this claim can be unpacked as follows:

Tom's judgment that he has reason to lie is made true by the fact that under ideal conditions he would be in a state that is correct iff lying is desirable.

That the desire to lie has this correctness condition is partly what makes it a desire. If it didn't have this correctness condition, it wouldn't be a desire, but some other mental state. But this means that desires are themselves normative properties, and desire facts are normative facts. If the constructivist tries to ground the correctness facts constitutive of desire facts in further attitudinal facts, he will face the same problem grounding the correctness facts that constitute these further attitudinal facts. So, if the Normativity of Attitudes is true, there is no way for the constructivist to claim that Tom's judgment that he has reason to lie is grounded in a desire and also claim that Tom's judgment is grounded in non-normative facts.

To summarize, the constructivist claims that all the normative facts are attitude-dependent. Since a fact is attitude-dependent only if it is grounded in some attitudinal fact, the constructivist must hold that all the normative facts are grounded in (some) attitudinal facts. But if the Normativity of Attitudes is correct, then all the attitudinal facts are themselves normative facts. The constructivist thus is in no position to show that all of the normative facts are grounded in non-normative facts. If the attitudinal facts are themselves normative facts, and, as the constructivist claims, all normative facts are grounded in attitudinal facts, then at most what follows is that all normative facts are

grounded in further normative facts. But the constructivists goal of grounding all normative facts in non-normative facts cannot be reached.<sup>12</sup>

6. Can the expressivist, unlike the error-theorist and constructivist, carry out his explanation of normative judgment if the Normativity of Attitudes is correct? Like the constructivist, the expressivist needs to show that all of the facts about normative judgments are grounded in some of the non-normative facts. But unlike the constructivist, the expressivist holds that normative judgments are, at least in the most explanatorily fundamental sense, conative rather than cognitive states.<sup>13</sup> So the expressivist is not required to claim, along with the constructivist, that there are normative facts that make some normative judgments true, and thus not required, as the constructivist, is, to ground these normative facts in non-normative facts. If normative judgments are not beliefs, much less beliefs about normative propositions, then the question of whether correct normative judgments are grounded in normative facts doesn't arise.

Let us then focus on the expressivist's central claim that normative judgments express non-cognitive attitudes. The Normativity of Attitudes, if true, applies to whatever attitudes are expressed by normative judgments, whether they are desires, intentions, or some cousin of these. Deliberating about whether to judge that, for example, one ought to lie, is deliberating about whether to have one of these attitudes, and whatever norms of correctness govern these attitudes will govern one's deliberation about the matter. If judging that one ought to lie expresses an intention to lie, then deliberating about whether to judge that one ought to lie is governed by the norm for intention. Applied to this case,

---

<sup>12</sup> Here is another problem. The claim that all normative facts are grounded in attitudinal facts, combined with the Normativity of Attitudes, entails that all normative facts are grounded in some subset of the normative facts. And this seems incoherent, at least if the set of normative facts is finite. If the set of normative facts is finite, then the subset of normative facts that ground all the normative facts must ground themselves. But grounding is an asymmetric relation, so this impossible. But what if the set of normative facts is infinite? If this is the case, for each normative fact, including each norm-judgmental fact, since those are normative facts, there will be a distinct normative fact that grounds it, consisting in whatever attitudinal fact the constructivist puts forth as the ground of normative truth. But this implies the existence of interminable metaphysical grounds, something that, as we pointed out earlier, no constructivist should want to accept.

<sup>13</sup> This may exclude some recent versions of expressivism, but it does not exclude the most recent views of Gibbard and Blackburn, the two most prominent expressivists. Both Blackburn (1998) and Gibbard (2003) commit themselves to the claim that at the most fundamental explanatory level, normative judgments are non-representational mental states.

the norm is that intending to lie is correct iff lying is desirable. Given that it is a conceptual truth about the attitudes expressed by normative judgments that they have standards of correctness, ascriptions of normative judgments such as that Tom judges that he ought to lie are themselves normative judgments. Given that the expressivist is a non-factualist about normative judgments, he must claim that ascriptions of normative judgments are not factual judgments expressing beliefs about some fact or state of affairs, but instead are themselves expressions of non-cognitive attitudes.

The question is whether the expressivist can consistently deny that ascriptions of normative judgments are factual judgments. One problem he faces is that this denial is in tension with his central claim that normative judgments are intrinsically motivating. It is in part because normative judgments, unlike other judgments, are intrinsically motivating, that the expressivist puts them on the non-cognitive side of the cognitive/non-cognitive divide. But the claim that normative judgments are intrinsically motivating seems to entail that normative judgments have causal powers. After all, how could a judgment motivate if it had no causal powers? And it seems possible to make a factual judgment about anything that has a causal power. Certainly, the expressivist about normativity has no ground to claim that causal judgments in general are nonfactual judgments. So if the expressivist is right that normative judgments are intrinsically motivating, he must accept that ascriptions of normative judgments are ascriptions of things that have causal powers. Absent some independent argument that causal judgments are themselves nonfactual judgments, he cannot justifiably assert therefore that ascriptions of normative judgments are nonfactual judgments.

The other, more obvious problem with denying that there are any norm-judgmental facts (or that ascriptions of normative judgments are factual judgments) is that the expressivist ends up denying the very data that he is supposed to explain—he ceases to be a metanormative theorist. For the task of the metanormative theorist is to show how the norm-judgmental facts stand in relation to the other facts that there are. The distinctive task of the expressivist is to show that the norm-judgmental facts either are, or are grounded in, some (non-normative) conative facts. Obviously, this task cannot be



accomplished if *there are no* (non-normative) conative facts. So, if the expressivist denies that there are any norm-judgmental facts, then he cannot give an account of normative judgment, much less one that shows all of them to be grounded in non-normative facts.

7. We conclude that if the Normativity of Attitudes is correct, then none of the prominent anti-realist views —error-theory, constructivism, or expressivism—can succeed. This leaves realism. Clearly, realism is consistent with the Normativity of Attitudes. The realist claims that there are attitude-independent normative facts that make some of our normative judgments correct. If attitudes are normatively constituted, then some of the attitude-independent normative facts are attitudinal facts, and our correct attitudinal judgments represent attitude-independent normative facts. Furthermore, if the Normativity of Attitudes is correct, there would appear to be no coherent alternative to realism. So, one might conclude, realism must be true. But this would be a mistake. To see why, let us examine what does and does not follow from our argument against the error-theorist.

A Pyrrhonian-style skeptic seeks to destroy our beliefs about our normative practice. To do this, he needn't commit himself to any claims about our practice. He merely has to point inconsistencies in what we believe about our practice. Unlike such a skeptic, whose project is merely destructive, the error-theorist seeks to explain our normative practice, albeit in a way that may show the practice to be defective. He thus commits himself to the truth of some claims about our practice as well as the claims that he invokes to explain our practice. It is these commitments that derive from attempting to explain our practice—a project the skeptic does not share—that get the error-theorist in trouble. If the Normativity of Attitudes is correct, the claims involved in characterizing the very data that the error-theorist seeks to explain commit him to claims that are inconsistent with his explanation.

The error-theorist claims that all normative judgments attributing normative properties are false because normative judgments are attributions of attitude-independent normative properties to actions, objects, or states of affairs, but no such properties are

ever instantiated. The realist, on the other hand, claims that there exist instantiated attitude-independent normative properties that make some of our normative judgments true. If the Normativity of Attitudes is correct, in attributing the normative judgments he takes to be false, the error-theorist must make a normative judgment himself, and this normative judgment commits him, by his own lights, to the existence of attitude-independent normative facts. The error-theorist thus cannot consistently claim that there are no attitude-independent normative facts and that all normative judgments are therefore false.

If this argument is correct, does it entail the realist's claim that some of our normative judgments are actually true? No. All the argument shows is that one cannot believe one thing without believing another: I cannot believe that John believes that lying is wrong without myself holding a normative belief, namely that John's belief is correct if and only if lying is wrong. Nothing follows as to the truth or falsity of either belief, or even whether it is possible to have these beliefs.

An error-theorist must attribute normative judgments and demonstrate that they are all false. The argument we described aims to show that if the Normativity of Attitudes is correct fulfilling the first condition makes it impossible to satisfy the second condition; but, as we just pointed out, this tells us nothing about whether the realist's central claim is true or false. Are there further conditions that an error-theorist must satisfy, and is there a way of moving from the claim that an error-theorist must satisfy these further conditions to the truth of at least one normative proposition? Presumably, in asserting that all normative judgments are false, the error-theorist is aiming to assert the truth. The error-theorist, it might be argued, thus takes himself to believe that all normative judgments are false, which means that he must take himself to have a mental attitude towards the proposition that all normative judgments are false that itself has the normative property of being correct if and only if all normative judgments are false. This, again, is merely a 'pragmatic' contradiction— in order to sincerely assert that all normative judgments are false, one must take one normative proposition to be true; it does not follow that any normative judgments are true.

If we could show not merely that the error-theorist must take himself to believe that all normative judgments are false, but that he must actually believe that all normative judgments are false, would this establish that there is at least one true normative proposition? No, what would be established is that it is impossible for the error-theorist's central claim to be true and believed— in order to believe that all normative propositions are false, including the normative propositions that must be true if one is to count as believing anything, one must satisfy one of the conditions that one is committed to claiming cannot be satisfied, one must believe something; from this claim it does not follow that any normative judgments are true—it does not follow that anyone ever does satisfy this condition, that anyone ever does believe anything. Establishing the conditions that need to be met in order to arrive at an error-theory at best tells us that if anyone believes an error-theory then at least one normative proposition is true; it would not tell us whether the antecedent of this conditional is true—whether anyone ever has believed an error-theory— and therefore would not tell us whether there are any true normative propositions.

If the Normativity of Attitudes is correct, then as long as normative judgments are attitudes, there is at least one true attitude-independent normative proposition. If anyone has ever made a normative judgment, there will be an attitude-independent normative truth about the correctness of the attitude constituting his normative judgment. In fact, if anyone ever has had an attitude about anything, it follows that there is at least one attitude-independent normative truth. While this certainly ups the ante for denying realism, it does not entail that realism is true. What follows is that the truth of realism is tied to the existence of attitudes. The reason that no form of anti-realism can succeed is that all of them assume the existence of some attitude or other. This very assumption commits them to realism. But it doesn't follow from this that there are any attitudes, and thus it doesn't follow that realism is true.

Furthermore, the Normativity of Attitudes seems to rule out giving a non-question begging argument for a fully general normative realism. The realist, like the error-

theorist, thinks normative judgments express beliefs. Whereas the error-theorist seeks to debunk all of our normative beliefs as false, the realist seeks to vindicate some of them as true. The realist, though, is not merely trying to demonstrate that a particular subset of normative beliefs must be true if some other subset of normative beliefs is true, because this would leave open the possibility that none of them are true. If he is to provide a non-question-begging argument for his view, the realist thus needs to occupy a normatively detached point of view that does not already commit him to the truth of any normative judgments. In order to carry out either task, normative judgments first need to be attributed in a way that leaves room for such a detached assessment. The problem is that if the Normativity of Attitudes is correct, then it is impossible to reach such a normatively detached perspective from all of our normative judgments. To attribute judgments at all is already to commit oneself to the truth of at least one normative judgment.

8. At the end of his essay “How To Be an Ethical Anti-Realist”, Simon Blackburn writes:

“But I hope I have said enough to show that nature and our theory of nature surround our ethical commitments in a way that gives us a *place* from which to theorize about them.” (174)

If thinking of ourselves as agents— creatures with the attitudes of belief, desire, and intention—involves making normative judgments, there is no such space for us to occupy. The project of giving a fully general explanation of our normative judgments, whether realist or anti-realist, presupposes that we can achieve a normatively detached point of view from which to describe all of our normative judgments. But describing a normative judgment is describing a belief, desire, or intention, and thus requires a commitment to normative judgments. Describing our normative judgments thus cannot be a normatively detached activity. If engaging in metaethical theorizing requires one to take up a normatively detached perspective, then metaethics very well may rest on a mistake.

## References

Blackburn, Simon (1993). 'How to Be an Ethical Anti-Realist', *Essays in Quasi-Realism*. New York: Oxford University Press.

—(1998). *Ruling Passions*. Clarendon Press.

Correia, F. (2008). 'Ontological Dependence', *Philosophy Compass* 3: 1013-32.

Dworkin, Ronald. (1996). 'Objectivity and Truth: You'd Better Believe it', *Philosophy and Public Affairs* 25: 87-139.

Fine, Kit. (1995). 'Ontological Dependence', *Proceedings of the Aristotelian Society* 95: 269-90.

—(2001). 'The Question of Realism', *Philosophers' Imprint* 1: 1-30.

Gibbard, Allan. (2003) *Thinking How to Live*. Harvard University Press.

Mackie, J.L. (1977). *Ethics: Inventing Right and Wrong*. Penguin.

McDowell, John. (1985). 'Values and Secondary Qualities', in Honderich, Ted, ed. *Morality and Objectivity*. London: Routledge and Kegan Paul.

Rosen, Gideon. (forthcoming). 'Metaphysical Dependence: Grounding and Reduction'.

Schaffer, Jonathan. (2009). 'On What Grounds What', in D. J. Chalmers, D. Manley, and R. Wasserman (eds), *Metametaphysics*. Oxford: Oxford University Press: 347-83.

Shah, Nishi. (2003). 'How Truth Governs Belief', *The Philosophical Review* 112: 447-82.

—(2008). 'How Action Governs Intention', *Philosophers' Imprint* 8: 1-19.

Wiggins, David. (1987). 'A Sensible Subjectivism?' *Needs, Values, Truth*. Oxford: Blackwell.