

# EPISTEMOLOGY IS NOT OPTIONAL

Assessing Epistemic Institutionalism as a Response to the  
Argument from Analogy

DOMINIC AIDAN BOWLES  
UNIVERSITY COLLEGE LONDON  
MPhil Stud in Philosophical Studies

I, Dominic Aidan Bowles, confirm that the work presented in my thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

# Abstract

The aim of this thesis is to assess Chris Cowie's epistemic institutionalism as a strategy for rescuing moral error theory from the argument from analogy (sometimes called the epistemic companions in guilt argument). I focus on what Cowie calls the internalism-based moral error theory. This is the view that moral judgements are mistaken because they concern reasons that are not suitably related to our desires. The proponent of the argument from analogy argues that epistemic judgements also concern such reasons and yet epistemology is not subject to an error theory, thus, the internalism-based error theorist's argument proves too much.

Epistemic institutionalism is the view that epistemic judgements are normative in the same way as judgements within sports and games, etiquette, fashion and the law. Cowie argues that these judgements are safe from the error theorist's argument and therefore so are epistemic judgements.

I present several challenges to institutionalism: I argue that even if we adopt institutionalism our genuine reasons to believe are very closely related to desires. I provide epistemic analogues to Plato's Ring of Gyges and argue that institutionalism is not equipped to explain our intuitions in cases like these. I also argue that it is a problem for Cowie that epistemology cannot be opted in and out of in the way that sports and games can.

I conclude that epistemic institutionalism is unacceptable. Epistemology is best understood as concerning genuine reasons unrelated to our desires. Therefore, the argument from analogy succeeds and internalism-based moral error theory leads to an unacceptable epistemic error theory. Thus, this version of moral error theory must be rejected.

# Impact Statement

There are two motivations for engaging with epistemic institutionalism as a response to the argument from analogy.

Firstly, the argument from analogy, if successful, would be a valuable tool in responding to moral error theory. Moral error theory has very outlandish consequences. According to error theory it is not the case that it's morally wrong to kill innocent children and it's not the case that it is morally good to effortlessly save someone from drowning. Error theory denies all kinds of seemingly obvious moral claims – claims that are taken for granted in almost all first order ethics. If true, it has major consequences for everyone.

Moral error theory is, therefore, a quite radical form of scepticism. One might think it is comparable with external world scepticism in the centrality of the claims it denies. However, unlike external world scepticism, many philosophers have argued in favour of moral error theory. As much as external world scepticism is debated in philosophy classes, very few people seriously advocate for it. There is a sense in which we are united in the fight against the external world sceptic; we engage with them as a spectre to be banished – a problem to be solved – we must avoid the unacceptable conclusion that we do not have any knowledge of the external world at all.

If the version of the argument from analogy I am considering succeeds, it will show that a particularly influential form of moral error theory leads

to an even more radical scepticism – an epistemic scepticism that no one can seriously endorse. This would be a significant result. It would begin to unite us in the fight against the moral error theorist in the way we are united against the external world sceptic. It would help us to see moral error theory for the unacceptable conclusion it is.

The second motivation for engaging with epistemic institutionalism is based on what institutionalism itself would mean for epistemology. A characteristic feature of sports and games is that they can be opted out of. We are permitted to stop playing them when it doesn't suit us. Many have begun to worry that we live in an age in which the truth matters less and less. Populist political leaders and their supporters seem in many ways to have opted out of responsiveness to truth and evidence when forming some of their beliefs. It may seem that such people are genuinely at fault in doing so. However, epistemic institutionalism, in allying epistemology with sports and games, seems to give up the resources to explain this.

This provides motivation to resist epistemic institutionalism or at least to better understand its apparent tension with our criticisms of the post-truth world. I will argue that this tension is a serious problem for institutionalism and that epistemology is better understood on the model of realist morality than sports and games, etiquette, fashion and the law.

# Contents

1	Introduction.....	8
1.1	Introducing moral error theory .....	8
1.2	Ways of responding to moral error theory.....	12
1.3	Arguments by analogy .....	13
1.4	Introducing the characters .....	17
1.5	Moral error theory.....	17
1.6	The argument from analogy.....	25
1.7	Cowie’s critique of the argument from analogy .....	29
1.8	Summary and plan for the rest of the thesis.....	34
2	Merely Institutional Normativity.....	37
2.1	The genuinely normative/merely institutional distinction.....	39
2.2	The genuine normativity of mere institutions.....	48
2.3	Reasons to comply with individual institutional judgements.....	51
2.4	Reasons to opt into an institution.....	54
2.5	Summary .....	58
3	The Epistemic Ring of Gyges.....	60
3.1	Plato and the Ring of Gyges .....	61
3.2	Barriers to a satisfactory epistemic Ring of Gyges case.....	63
3.3	Epistemic Ring of Gyges case 1: Trivial Falsehoods .....	66
3.4	Epistemic Ring of Gyges case 2: man-made climate change.....	70

3.5	Necessary desires.....	75
3.6	Summary .....	83
4	Opting out of epistemology .....	85
4.1	The Epistemic Sceptic .....	85
4.2	Trivial Falsehoods and No Ordinary Explanations .....	94
4.3	Epistemology as constitutive of belief.....	95
4.4	Joyce and forestalling ‘so what?’ responses .....	97
4.5	Beliefs, shmeliiefs .....	105
4.6	Summary .....	110
5	Conclusion.....	112
6	Bibliography.....	117

# **1 Introduction**

I will assess Chris Cowie's epistemic institutionalism as a strategy for rescuing moral error theory from the argument from analogy (also called the epistemic companions in guilt argument). Epistemic institutionalism is the view that epistemic judgements are normative in the same way as judgements within sports and games, etiquette, fashion and the law (Cowie, 2019). Cowie argues that these judgements are not subject to an error theory and thus neither are epistemic judgements. I will argue that the normativity of epistemic judgements is not optional in the way that the normativity of judgements within sports and games, etiquette, fashion and the law are. It is characteristic of those kinds of judgement, I will argue, that we do not have to follow them when it doesn't suit us: we don't have to play rugby if we don't want to. I will argue that epistemology is not like this. Epistemology is not optional. I will argue that, as a consequence, epistemic institutionalism does not succeed in rescuing moral error theory from the argument from analogy.

## **1.1 Introducing moral error theory**

Moral error theory is a sceptical thesis about moral judgements. Moral error theorists agree with moral realists that when we make moral judgements, we are trying to describe the world, our judgements are capable of being true or false and our moral judgements express beliefs. In this they contrast with non-cognitivist positions like expressivism. Unlike moral realism, error theory is a sceptical thesis; it makes the claim that moral judgements are systematically false. Error theorists diagnose different problems with moral judgements, but they all agree that they



try to describe a moral reality that doesn't exist and thus turn out to be false.

Moral error theory was popularised by JL Mackie in his book *Ethics: Inventing Right and Wrong* (Mackie, 1990). His basic argument is that moral judgements commit us to the existence of objective values and prescriptions which do not exist, therefore, moral judgements are systematically false. His most famous argument is the argument from queerness: objective values and prescriptions would be so queer that we should not believe in them. One way he makes this more precise in terms of the kinds of reasons objective prescriptions would commit us to: moral judgements tell us to do things that we don't want to do and thus rely on reasons for action disconnected from our desires. There are no such reasons, he says, so moral judgements are false. A more detailed version of this argument given by Joyce will be central to my investigation.

Moral error theory has some quite outlandish consequences. According to error theory it is not the case that it's morally wrong to kill innocent children and it's not the case that it is morally good to effortlessly save someone from drowning. Error theory denies all kinds of seemingly obvious moral claims – claims that are taken for granted in almost all first order ethics.

It is, therefore, a quite radical form of scepticism. One might think it is comparable with external world scepticism<sup>1</sup> in the centrality of the

---

<sup>1</sup> The denial that we have any knowledge of the external world, motivated by arguments like Descartes's evil demon or brain in a vat cases.

claims it denies. However, unlike external world scepticism, moral error theory is a position that many people are happy to adopt in philosophical contexts. After Mackie, many other people have argued that moral judgements are false (Joyce, 2001) (Olson, 2014). While, as much as external world scepticism is debated in philosophy classes, very few people are willing to seriously advocate for it. There is a sense in which we are united in the fight against the external world sceptic; we engage with them as a spectre to be banished – a problem to be solved – we must avoid the unacceptable conclusion that we do not have any knowledge of the external world at all.

Why do I say that people ‘adopt moral error theory in philosophical contexts’ rather than simply that people are genuine moral error theorists? Many moral error theorists are both willing and able to engage in thoughtful moral discussion and are motivated by moral considerations in their everyday lives. Perhaps this is simply because they happen to care about moral ends (the wellbeing of others, honesty etc.). However, a more plausible interpretation is that they have a hard time committing to their radical sceptical thesis. Someone who genuinely thinks that they have no reason not to be decent when they don’t want to would behave more like a psychopath than a thoughtful philosopher. A committed error theorist of the sort I am considering would see their actions as irrational when they do not service their desires. We regularly find ourselves in situations where we want something that moral considerations do not allow. Suppose you are on the bus and you would really like a seat but an old woman gets on who needs the seat more than

you do. You give up the seat for her. How would you justify this? It would be a very strange justification of this to say you want her to sit down more than you want to sit down yourself. Which of these things you want more is irrelevant to what you ought to do in this situation. A committed error theorist, however, should take this point seriously; whether they ought to give up their seat *depends* on how much they want the old woman to sit down. It is hard to imagine Joyce and Mackie justifying their actions, even to themselves, in this way.

This parallels the difficulty of committing to external world scepticism outside of philosophy. No one really gives brain-in-a-vat scenarios, say, equal credence with ordinary claims about the external world. When people plan their trips to the shops, they consider things like what the traffic will be like on a Saturday, not whether an evil demon will have replaced the shops with an elaborate replica. To genuinely think that these things merit equal consideration would be very strange indeed.

The disparity between the treatment of moral error theory and external world scepticism, therefore, remains largely within the context of philosophy. I am interested in whether even this disparity is misguided. Should we treat moral error theory as a position that is possibly true or as a philosophical argument with an unacceptable conclusion? Should the moral error theorist be considered alongside the external world sceptic – a spectre to be banished rather than a genuine participant in the debate?

## 1.2 Ways of responding to moral error theory

There are different ways of dealing with sceptical arguments against a target discourse (external world knowledge ascriptions, moral judgements or whatever). We can either let the sceptic set the terms of the debate; allow them to suspend judgement about the target discourse and present an argument which we then criticise. Alternatively, we can do as Moore proposes for external world scepticism and refuse to take the target discourse off the table (Moore, 1993). We can argue that the sceptical conclusion is so implausible that, whatever their argument is, it must have gone wrong somewhere. For Moore, this involves arguing from the obvious fact he has hands to the conclusion that the sceptic is mistaken. This might look question begging; from the sceptic's perspective, whether Moore knows he has hands is part of what is up for debate – it cannot simply be used as a premise against them.

It is understandable for the sceptic to view Moore as begging the question but perhaps it is not illegitimate for him to do so. The Moorean shift essentially takes the external world off the table – our knowledge of its existence is more certain than any premises the sceptic might use in their argument against it and thus the sceptic must have gone wrong.

There is an analogue to the Moorean shift as a response to moral error theory: we can argue from the obviousness of the fact that it is morally wrong to kill innocent children to the conclusion that moral error theory is mistaken. One could (quite convincingly) argue that the claim that it is wrong to kill innocent children is much more plausible than any premise

that the error theorist could come up with and thereby produce a legitimate argument against moral error theory.

This strategy will not be my focus here. Even if this is a legitimate response, it will not be very convincing to someone inclined towards error theory. These people are willing to question the reliability of the intuition that it is morally wrong to kill innocent children and thus question the force of this Moorean argument. It is, therefore, dialectically helpful to have more in our arsenal than a Moorean shift.

### **1.3 Arguments by analogy**

An alternative way to argue is by analogy (also called companions in guilt arguments). To argue against a sceptic by analogy one must first choose a candidate discourse with important similarities to the sceptic's target. One must argue that the sceptic's argument will apply equally well to the analogous discourse and thus lead to a broader scepticism. One must then argue that the analogous discourse is clearly unproblematic, therefore, the sceptic's argument must have gone wrong somewhere – it over-generates.

Arguing by analogy, like the Moorean shift, is indirect; it does not target a particular premise of the sceptic's argument but points to its unacceptable consequences and concludes from that that the sceptic has gone wrong. It has benefits over the Moorean shift however in that it does not beg the question against the sceptic; we can leave our knowledge of the external world, moral judgement or whatever on the table and have the argument still go through.

This is not to say, however, that arguments by analogy do not beg questions against those attracted by even broader forms of scepticism. For instance, some of the arguments I will put forward may be seen as question begging by epistemic sceptics (proponents of what we will later call the epistemic error theory). In chapter 3, for example, I will appeal to intuitions that are distinctively epistemic in character which an epistemic sceptic may argue are off the table while we attempt to ground our epistemic judgements. The value of arguments by analogy is not that they begin from a point that *all* sceptics can agree with, it is that they allow us to target those who want to hold a restricted scepticism without appealing to claims that *they* deny. They allow us to argue against the moral error theorist, for example, without appealing to moral intuitions. I will be focussed on an argument against moral error theory by analogy with epistemology. This argument has been prominently made by Cuneo, Rowland and Das amongst others (Cuneo, 2007) (Rowland, 2013) (Das, 2016). The structure of the argument is as follows:

P1. If the argument for moral error theory is sound, then an analogous argument for epistemic error theory is sound too.

P2. Epistemic error theory is not true.

Therefore,

C. The argument for moral error theory is not sound.

(Cowie, 2019, p. 31)

The argument from analogy with epistemology (henceforth, the argument from analogy) is strictly speaking a *kind* of argument. There are many different arguments for moral error theory that identify different features of moral judgements as problematic. For each of these there will be an argument from analogy: epistemic judgements must be argued to share the particular feature of moral judgements the error theorist engaged with finds problematic. To allow me to engage with the arguments in sufficient depth, I will focus on one argument for moral error theory – one based on the wrong kind of reasons – and the corresponding argument from analogy.

I will engage with moral error theory and the argument from analogy through Chris Cowie's 2019 book *Morality and Epistemic Judgement: The Argument from Analogy*. Cowie argues that the argument from analogy is not successful because there are relevant differences between morality and epistemology that shield epistemic judgements from the error theorist. His core claim is that epistemic judgements are normative only in the way that judgements in sports and games, fashion, etiquette and the law are normative. These kinds of judgement are not affected by the error theorist's argument and thus neither are epistemic judgements.

My goal is to reject Cowie's argument. I will argue that epistemic judgements cannot be treated on the model of sports and games, fashion, etiquette, and the law, and in fact have much more in common with morality. As I have said, my criticisms will be centred around the way that one cannot opt out of epistemology. I will conclude that the version of the argument from analogy I am considering is successful.

This will not refute moral error theory (considered as the claim that moral judgements are systematically false) because there may be other arguments for it without unacceptable consequences. This is a limitation of the argument from analogy: it is targeted at *arguments* for error theory rather than its central claim (moral judgements are systematically false). Thus, while the argument from analogy provides us a model that can be adapted to any argument for moral error theory, each of these arguments will need to be made separately.

However, the argument for error theory I will consider has a claim to be one of the strongest, and certainly one of the most influential. It has roots in Mackie's argument from queerness and focusses on a worry that lots of people have had about morality: why should we follow moral rules if we don't want to? If *this* argument were found to over-generate then it would be a significant blow to moral error theory.

By addressing this argument for error theory and the corresponding argument from analogy I hope to lend plausibility to the broader claim that moral error theory is not a special form of scepticism. The problems that error theorists have with moral judgements spread much more widely than morality and thus taking them seriously requires that you adopt a much more radical scepticism. I suspect that error theorists will not be comfortable with this conclusion. Radical scepticism is too high a price to pay for dispensing with morality.

There is not space here to make a full argument for the general claim. Consider this a case study for the argument that moral error theory



commits one to radical scepticism – there is nothing uniquely problematic about moral judgement.

#### **1.4 Introducing the characters**

In the remainder of the introduction, I will introduce the characters and their motivations. I will introduce the variety of error theorist who concerns us and the details of their argument. I will set out the argument from analogy as it applies to them. Much of this exposition will follow Cowie's presentation and I will note anywhere where we depart from how he sees the problem. I will then set out Cowie's main response to the argument from analogy – epistemic institutionalism.

#### **1.5 Moral error theory**

In his book Cowie identifies two different varieties of moral error theory which he calls internalism-based moral error theory and irreducibility-based moral error theory. They differ over what they take to be problematic about moral judgements. Internalism-based error theorists take moral judgements to concern categorical, normative reasons for action. They argue that there can be no such reasons because they would be incompatible with reasons-internalism – a view about how our reasons for action must relate to our psychologies. Irreducibility-based error theorists take moral judgements to be concerned with irreducibly normative properties and relations. They then argue that there are no such properties or relations.

Cowie addresses both kinds of error theory and the corresponding arguments from analogy. He considers his main argument to be helpful

in both cases. However, I will focus exclusively on internalism-based moral error theory. The two arguments focus on different features of moral judgements and objections to one will not necessarily be objections to the other. Therefore, it is sensible to handle them separately. Internalism-based error theory is closer to what Mackie appears to have had in mind with his argument from queerness and that's why I choose to focus on that version of the argument.

It is worth making a brief note on terminology: I will discuss normative judgments, norms and reasons following Cowie's usage: norms are the contents of our normative judgements—moral norms are the contents of moral judgements, epistemic norms are the contents of our epistemic judgements and so on. Reasons are facts which speak for or against taking an action or forming a belief. I will be concerned with justifying reasons, i.e. facts we could appeal to in giving justifications of our actions or beliefs.

### **1.5.1 Internalism-based moral error theory**

Cowie's presentation of internalism-based moral error theory is based on the argument in Richard Joyce's *Myth of Morality* (Joyce, 2001). The structure is as follows:

Commitment Premise: Moral judgements concern categorical, non-institutional reasons for action.

Existential Premise: There are no categorical, non-institutional reasons for action.

Therefore,

Moral judgements are mistaken.

The conclusion follows due to a background assumption that moral judgements are not supposed to be about something that doesn't exist (more on this in a moment).

Let us take each premise in turn.

### **1.5.2 Commitment Premise**

Many of our normative judgements are sensitive to the desires of the agent we are judging about. To use Cowie's example, suppose you and your friend Tyson are in a café and he is deciding whether to choose tea or coffee (Cowie, p. 11). You tell him that he ought to choose tea so he can get a good night's sleep. This is sensitive to your judgement of his desires. If you thought that he did not want a good night's sleep and instead wanted to stay up all night playing cards, then you may judge differently: he should choose coffee.

Moral judgements on the other hand, do not seem to work like this. Again, using Cowie's example, your judgement that Tyson should not keep slaves is not sensitive to your judgement of his desires. He should not keep slaves even if this is the thing he wants most in the world. Moral judgements seem to concern what you ought to do, independently of your desires. This is what Cowie means when he describes the reasons moral judgements are concerned with as categorical – they speak for or against actions independently of the desires of the agents involved.

There is still the 'non-institutional' part of the commitment premise to address. Cowie defines institutions as "sets of prescriptions concerning

which courses of action are forbidden, permissible, or recommended” (Cowie, p. 11). Sports and games are the paradigm cases of institutions with their clearly defined rule sets. Cowie also considers etiquette, fashion and the law among institutions. He considers the norms of institutions to be categorical. In the case of rugby, he says “when your teammate on the rugby pitch tells you that shouldn’t pass the ball forwards, he doesn’t mean that you shouldn’t do so unless you want to” (Cowie, p. 12). The thought is that like morality, sports and games also involve normative judgements that apply to you regardless of what you want.

There is, however, a difference between morality and the normativity of institutions. Moral reasons are, in Cowie’s terminology, not merely institutional but *genuinely normative*. Different writers have given this distinction different names. Joyce has called genuine normativity the ‘authority’ or ‘practical oomph’ of moral reasons (Joyce, 2001). The distinction has also been called the substantive/formal distinction (McPherson, 2011) and the reason-implying / rule implying distinction (Parfit, 2011). The exact nature of the distinction is hard to specify but the important point is that there is a difference in kind between moral judgements and judgements of institutions that roughly corresponds to whether you have to take them seriously. Purportedly, moral judgements always have to be taken seriously. According to realists, they always come with genuinely normative reasons that merit inclusion in our all-things-considered evaluation of action. Whereas institutional

judgements (or *merely* institutional judgements) can be ignored in all-things-considered evaluations. I will say more on this in section 2.1.

I take the commitment premise to fall out of a version of moral realism. As I will be engaging in arguments over moral error theory, I will assume a roughly realist structure to moral reality (or purported moral reality). That is, I will consider the error theorist in dialogue with the moral realist. In this debate both sides accept that morality purports to be authoritative, universal, inescapable, objective and importantly for our purposes, categorical and non-institutional. People have considered views on which morality is not like this: Foot, for instance, argues that morality is not categorical (Foot, 1972). Fictionalists and non-cognitivists also challenge the common ground between the realist and the error theorist. However, all of these views are out of step with standard moral realism and thus I will set them aside as beyond the scope of my investigation. All of the participants in our debate, therefore, will take the commitment premise for granted. While this may be unsatisfying if you hold these views, bracketing these kinds of concerns will be necessary to cover the rest of the topic in satisfactory depth.

### **1.5.3 Existential Premise**

The internalism-based error theorist's argument for the existential premise is based on reasons-internalism which is a view about how our genuinely normative reasons for action relate to our psychologies. Roughly put, reasons-internalists think that reasons for action must bear an appropriate relation to the desires of the agent to whom they apply or to their desires suitably idealised.

Reasons-internalism is most closely associated with Bernard Williams (Williams, 1981) but many people have given arguments for it. In section 4.4 I will discuss a particular motivation from Joyce (Joyce, 2001). A common argument is that reasons for action are supposed to be the sorts of things that agents could, in principle, be moved by. The only things that meet this criterion, an internalist might say, are facts that are related to an agent's desires. As I have said, we see traces of this kind of argument in Mackie's original presentation of error theory.

Internalism does not automatically lead to a moral error theory, even if we accept the commitment premise. The commitment premise says that moral judgements concern categorical, non-institutional reasons for action, that is, reasons that we have independently of our particular desires. The only way to square this with reasons-internalism is to posit that there is some constraint on our desires that means we always end up with moral reasons that bear an internalist-appropriate relation to our desires.

The idea is that we all either have, or under suitable circumstances (full information, perfect rationality etc) *would* have, desires to comply with moral judgements. In this case, it would not matter what your desires were within these limits, you would end up with moral reasons because you had moral desires.

As Cowie explains, error theorists are suspicious of these necessary desires claims. It seems possible to imagine someone without moral desires: Joyce gives an example from Gilbert Harman of a Mafioso who

lacks the necessary desires to give him reason to cease his life of crime (Harman, 1977). Such an agent is aware of the impacts of his crimes and yet still wants to go on with them. Thus, it is implausible that even idealising this agent would produce the desires necessary. This suggests, Joyce says, the burden is on the proponent of necessary desires to argue that such desires exist. Whether this burden shifting argument is successful, it is worth noting that not only internalism-based error theorists but also many of their opponents deny the plausibility of necessary moral desires.

If there are no such desires, then reasons-internalism would rule out the reasons that, according to the commitment premise, are required by moral judgements and as such, establish the moral error theory. I will not engage with the plausibility of reasons-internalism directly, nor the plausibility of necessary moral desires. However, in chapter 2 I will consider what reasons-internalism means for our genuine reasons to believe and in section 3.5 whether necessary epistemic desires are more plausible than necessary moral desires.

#### **1.5.4 Varieties of genuinely normative reason**

It is worth pausing for a moment on what the stock of genuine reasons looks like for different participants in our debate. Internalists think that the entire stock of an agent's genuinely normative reasons are related to their desires or their desires suitably idealised. Moral realists think that we have genuinely normative reasons to conform to moral judgements, regardless of what we want. Error theorists think that these two

thoughts are incompatible – the reasons that moral judgements concern cannot exist in light of internalism.

I will occasionally refer to these reasons as ‘moral reasons’. By this I mean nothing more than the external reasons that error theorists and many moral realists think must exist in order for moral judgements to be true. I do not mean to say that there is anything special about moral reasons that distinguish them from other kinds of genuinely normative reasons. I wish to remain silent on this. All that matters about moral reasons for my purposes is that they are purported to be genuinely normative – they belong in our all-things-considered assessments of action – and they are external – they do not depend on the particular desires of the agents for whom they are reasons.

I will also refer to ‘prudential reasons’. I use this term because Cowie uses it in his argument for institutionalism. He talks about prudential reasons as though they are genuinely normative. He does not, however, suggest that prudential reasons are problematic for error theorists. This is an interesting point of his book. Prudential reasons concern our self-interest – we have prudential reason to do something when it is good for us. This is not necessarily palatable for the reasons-internalist. Again, following Cowie, I have framed the reasons-internalist as particularly focussed on desires.<sup>2</sup> It is not clear however that my desires always line up with what is good for me. For instance, I regularly do not want to go

---

<sup>2</sup> This too may be a contentious point. We might also have stated reasons-internalism in terms of relating genuine reasons to an agent’s interests or ends rather than their desires. This would not radically change what I have to say in the rest of the thesis though thus I follow Cowie’s usage.



for a run even though it may be good for me to do so. Therefore, it is not immediate that any given reasons-internalist can tolerate prudential reasons as genuinely normative. I will set this question aside though. Prudential reasons are certainly much easier for the reasons-internalist to explain than moral reasons are. Thus, while strictly speaking the error theorist may need an argument in order to appeal to them, I will grant them this point. I will consider prudential reasons to be part of the common ground between the moral error theorist and the moral realist. The error theorist can, therefore, appeal to them when looking for genuine reasons to follow epistemic norms.

### **1.6 The argument from analogy**

Now we have focussed in on a particular argument for error theory we can fill in the details of the basic argument from analogy:

Parity: If the argument for internalism-based moral error theory is sound, then an analogous argument for epistemic error theory is sound too.

Epistemic Existence: Epistemic error theory is not true.

Therefore,

C. The argument for internalism-based moral error theory is not sound.

(Cowie, p. 33)

This argument will be the focus of the remainder of the thesis. Cowie calls it the internalism-based argument from analogy. I will mostly refer to it

simply as the argument from analogy. Again, I will consider each premise in turn.

### **1.6.1 Parity Premise**

This is the analogy part of the argument. To make the case for it we must show that epistemic judgements share the features of morality that internalism-based error theorists object to, that is, the reliance on categorical, genuinely normative reasons. Cowie also emphasises that it must be argued that epistemology's reliance on categorical genuinely normative reasons is equally problematic to the reliance of moral judgements on categorical genuinely normative reasons.

Firstly, the reliance on categorical, genuinely normative reasons. Epistemic judgements and their corresponding reasons concern the relation of one's beliefs to truth and evidence. Whether you have epistemic reason to believe something depends only on the evidence you have, not your particular desires. In this sense, epistemic reasons are categorical – they apply to you independently of your desires. This is widely accepted and even Cowie in his critique of the argument from analogy accepts this point; this is just how epistemology is done. Perhaps the harder point to argue is that epistemic reasons are genuinely normative.

Proponents of the argument from analogy argue that there is a contrast between epistemic norms and the norms of institutions. Das for example writes:

Like moral reasons of certain kinds, a wide range of epistemic reasons stand in sharp contrast to the norms that govern chess and tennis...

(Das, 2017, p. 64)

They think that epistemic judgements are not just categorical in the way that the rules of rugby are but in a more substantial sense, the way morality is. Whether epistemology is genuinely normative will be the subject of much of the thesis so we shall leave it for now.

Let us also briefly consider whether categorical genuinely normative reasons are as problematic in the epistemic case as they are in the moral case. Cowie thinks it is necessary to make this point because epistemology is in the business of evaluating beliefs whereas morality is in the business of evaluating action. The arguments Cowie considers for internalism make reference to how we explain people's *action* and how reasons must be able to move people to *act*. Our belief forming practices do not work in quite the same way as our ordinary action (for example we have much less voluntary control over our beliefs) and thus at least *prima facie* the same arguments against categorical genuinely normative reasons for belief may not apply.

Many writers have thought that the same concerns *do* apply. They propose that there is no relevant difference between reasons for belief and reasons for action such that we may admit external reasons in one domain and not the other. Cowie gives an argument against this parity in chapter 4 of his book, but this will not be the focus of the thesis. I will take

for granted that internalists object to all categorical, genuinely normative reasons and discuss whether epistemology makes such a commitment.

### **1.6.2 Epistemic Existence Premise**

The parity premise aims to show that the error theorist's argument extends to epistemology – if we accept the argument for moral error theory, we should accept an epistemic error theory too. We might respond to this in two ways: reject the argument for moral error theory or accept an epistemic error theory. The epistemic existence premise serves to force us into the first of these options; we cannot accept an epistemic error theory.

There are various quite convincing arguments that epistemic error theory is untenable. Many of them argue that it would be self-undermining – put very briefly, it has as a consequence that we would have no reason to believe it. The details of these arguments are not directly relevant for my purposes, therefore, I mostly omit them. I will, however, mention two of them in sections 2.2 and 4.3.1.

They are not directly relevant because we can give the argument from analogy force without explicitly arguing against the epistemic error theory. We can present the argument as a diagnosis of an inconsistency in the moral error theorist's position. Consider the following quote from Cuneo:

J. L. Mackie was a moral nihilist . . . But he expressed no sympathy for epistemic nihilism . . . I doubt that he would've had much

sympathy with it. For Mackie believed in the power of argument and reason.

(Cuneo, 2007, p. 122)

We can say, as a matter of fact, people who are attracted to moral error theory in philosophy are not attracted to epistemic error theory. The argument from analogy is a valuable response to such people; if successful, it shows that they cannot keep their scepticism localised.

### **1.7 Cowie's critique of the argument from analogy**

Cowie critiques the argument from analogy in a number of ways. His main argument is focussed on denying that epistemic judgements concern categorical, non-institutional reasons and thus denying the Parity Premise.

At the core of his strategy is what he calls the basic rationale:

The basic rationale: Epistemic judgements, unlike moral judgements, are normative only in the sense that judgements within etiquette, fashion, sports and games, and the law are normative.

(Cowie, p. 52)

I will discuss how exactly this is supposed to avoid the parity premise in section 2.1. In the terminology introduced in section 1.5.2, it involves arguing that epistemology is merely institutionally normative. That is, it does not make demands on agents that cannot be ignored in all things considered evaluations. He calls his view epistemic institutionalism.

Cowie has two arguments for his position: the argument from ordinary explanation and the argument from trivial truths. Both involve finding features of epistemic discourse that it has in common with merely institutional judgement and differs from morality. I will eventually provide arguments against institutionalism that directly parallel these two arguments. For a summary see section 4.2.

### **1.7.1 The argument from ordinary explanation**

The core claim of Cowie's first argument is as follows:

Ordinary explanation (epistemic): The ordinary explanation of why, if at all, one ought to follow epistemic norms is the same as the ordinary explanation of why, if at all, one ought to follow the norms of etiquette, fashion, sports and games, and the law. It is that there is moral or prudential reason to do so.

(Cowie, p. 63)

The thought is that when someone asks why they should follow an epistemic norm, it is natural to appeal to non-epistemic considerations. He gives an example from Railton of a student, Gary, in an epistemology class asking the following:

[P]hilosophers seem to agree that there are certain standards of belief. . . They spend all their time disagreeing about what exactly these standards are, but they just seem to assume that we'll want to follow them. Suppose I don't? What can they say to me?

Railton suggests that the teacher might respond as follows:

Well, you understand why you should eat your vegetables? You may not care about these epistemic standards as such, but you do care a lot about other things. And you're more likely to get what you want if you have warranted beliefs.

(Railton, 1997, p. 294)

Cowie thinks that this is a good response. In fact, he claims further that it would be strange *not* to give this response to someone pressing you on why they should follow epistemic norms. The ordinary way to explain your reasons to follow epistemic norms is to appeal to moral or prudential, that is, non-epistemic reasons to do so.

He thinks that this is exactly how we would deal with someone pushing us on why they should follow the rules of a mere institution – we would appeal to reasons from outside that institution. If someone asks why they should follow a rule of chess, for example, we might answer in prudential terms – because it's fun.

Cowie thinks that moral rules are not like this: it is rarely possible and indeed peculiar if someone asks why they shouldn't kill someone and you respond in terms of prudential reasons. For instance, were you to say you shouldn't kill people in order to avoid social exclusion you would seem to miss the point about the wrong of murder. The ordinary explanation for following moral norms is in terms of moral reasons themselves – it is because there are moral reasons to do so.

Cowie thinks this provides evidence for epistemic institutionalism – the reasons to follow the norms are most naturally identified as coming from

sources separate from the institution itself – from our moral and prudential reasons.

The discourse with Gary the epistemic sceptic will form a central part of my critique of Cowie. In section 3.4 I will argue that the response the teacher gives to Gary is not always available when we need it and in section 4.1 I will question the intelligibility of conversing with Gary at all if he is a genuine epistemic sceptic.

### **1.7.2 The argument from trivial truths**

The core claim of Cowie's second argument for epistemic institutionalism is as follows:

Trivial truths (epistemic): Epistemic norms, much as the norms of etiquette, fashion, sports and games, and the law, are the kinds of things that, unless there is a moral or prudential reason to follow them, don't generate genuinely normative reasons.

(Cowie, p. 66)

Cowie claims there are some trivial truths which even though we have strong evidence for, we lack genuinely normative reason to believe. Trivial truths are those which we have no moral or prudential reason to believe. In support of this claim, he gives the following example from Adam Leite:

I'm standing next to a door at a convention centre. I idly notice that all of the many people I've seen come out of the door have been accompanied by dogs. I am certainly not being irrational if I fail to form the belief that the next person to come out of the door



is likely to be accompanied by a dog. More generally, as I go through my day, I gain all sorts of evidence supporting all sorts of beliefs. But I don't form most of them, and it hardly seems plausible that I have any reason to do so, given that they are about matters of complete indifference to me.

(Leite, 2007, p. 460)

Cowie argues that in this case epistemology tells you to believe that the next person will be accompanied by a dog (because it is supported by your evidence) and yet due to the triviality of the proposition we have no genuine reason to form that belief.

He also points to examples from Papineau about the names of the kings of Assyria and the blades of grass on a lawn:

I have already cited the examples of the kings of Assyria and the blades of grass on my lawn. If nothing practical hinges on these matters for me, and I am going to keep my opinions to myself, I say that there is nothing wrong with my having false beliefs on these matters.

(Papineau, 2013, p. 68)

Cowie argues that Papineau too is making the point that truth and evidence in these cases do not generate genuine reasons to believe truly.

He claims that this is analogous to our commitment to other mere institutions; when you have no moral or prudential reasons to follow the

rules of rugby, you have no genuinely normative reasons to do so. He argues that these cases support trivial truths (epistemic).

Cowie contrasts this with morality. He says there is no moral analogue to trivial truth cases, that is, it is never the case that morality tells us to do something so trivial that we have no genuine reason to do it; if morality tells you to do something, you always have reasons to do it, at least, so the commitment premise says.

Since epistemic norms share this feature with merely institutional norms and not with morality, Cowie says this provides further evidence for epistemic institutionalism.

I will argue in section 3.3.1 that trivial truths (epistemic) looks much less plausible in other cases.

## **1.8 Summary and plan for the rest of the thesis**

So now we can see the lay of the land. Internalism-based error theorists object to moral judgements on the grounds that they rely on categorical, genuinely normative reasons. The corresponding argument from analogy says that this should lead them to also accept an epistemic error theory. Since accepting an epistemic error theory is absurd, this means we should reject the internalism-based moral error theory. Cowie denies this implication because epistemic judgements are merely institutionally normative; they are normative only in the way that sports and games, etiquette, fashion and the law are normative.

This is the plan going forward:

In chapter 2 I will lay the groundwork for my case against Cowie. I will explore how exactly sports and games, fashion, etiquette and the law are normative and what explains the fact we sometimes have genuine reasons to conform to their judgements but not always. This will be necessary if we are to understand Cowie's basic rationale. I will emphasise how epistemic institutionalism means that our genuine reasons to believe must be closely related to our desires even though Cowie claims that epistemic reasons are categorical. The consequence of this is that we don't need to follow epistemic norms when they do not line up with our desires – they would be, in this sense, optional.

In chapter 3 I will attempt to provide an epistemic analogue of Plato's Ring of Gyges, that is, a case where we lack moral or prudential reasons to follow epistemic norms but it still seems we should follow them. In other words, we have to follow epistemic norms even though it doesn't suit us. Using the groundwork from chapter 2, I will argue that the model of merely institutional normative judgements cannot explain this.

In chapter 4 I will present a different challenge. I will argue that there are conceptual problems with opting in and out of epistemology in the way we can opt in and out of some mere institutions. I will argue that this provides another reason to think that epistemic institutionalism will not help moral error theorists avoid the argument from analogy. I will also consider the argument that epistemic norms are constitutive of belief and how this interacts with institutionalism.

I will conclude that the model of games and sports, etiquette, fashion and the law is not suited to epistemic judgements. Epistemology has much more in common with morality. Therefore, this version of the argument from analogy succeeds.

## 2 Merely Institutional Normativity

Let us briefly recap the story so far:

I am engaging with the form of the argument from analogy which targets internalism-based moral error theory. Internalism-based error theorists object to moral judgements on the grounds that they concern a problematic kind of reason – reasons that lack an appropriate relation to the desires of the agent to whom they apply. They say that since there are no such reasons, moral judgements are systematically false.

Recall also that I am assuming a background of moral realism about purported moral reality; I am taking for granted that moral judgements are supposed to be in the business of truth, they make demands on individuals that are independent of their desires etc. I am, therefore, taking for granted the Commitment Premise of the error theorist's argument discussed in section 1.5.2. As I said there, in doing this I set aside various metaethical views including non-cognitivism and fictionalism.

The argument from analogy attacks the error theorist's argument on the grounds that it proves too much; epistemic judgements also concern reasons without the relation to desires that the internalist requires, thus if there really are no such reasons, then epistemology is in error too. Since, the argument goes, epistemology is not subject to an error theory, the moral error theorist's argument must have gone wrong somewhere.

Cowie's response to the argument from analogy is to propose that epistemic judgements are normative not in the same way as morality but

only in the same way as sports and games, etiquette, fashion and the law. He calls the second class of judgements 'merely institutional'. According to Cowie, institutions are "sets of prescriptions concerning which courses of action are forbidden, permissible, or recommended." (Cowie, p. 11). On this definition, morality can be considered an institution, as can Rugby Union, Chess, the laws of the United Kingdom, and so on. They all commend some actions and prohibit others. Judgements are described as *merely* institutional when they express *nothing more* than what follows from some list of rules; they lack the purportedly problematic reasons commitments of morality. Cowie claims that merely institutional judgements can be true independently of the desires of the agents to whom they apply and yet are unproblematic for the error theorist. Therefore, if epistemic normativity is merely institutional, it is protected from the error theorist's argument and thus the argument from analogy fails to show that internalism-based moral error theory over-generates.

The aim of this chapter is to consider the exact sense in which mere institutions are normative. This will help us to understand how epistemic institutionalism is supposed to avoid epistemic error theory. I will be particularly interested in how the norms of mere institutions can become genuinely normative, that is, how the norms of mere institutions, despite being only lists of rules and their consequences, sometimes merit consideration in our all-things-considered evaluations. Section 2.1 will be devoted to explaining what Cowie means by merely institutional normativity and how he thinks this protects a set of norms from error

theory. In the remainder of the chapter, I will depart from Cowie's exposition and consider ways in which mere institutions like sports and the law become genuinely normative. I will argue that moral error theorists think their arguments have genuinely normative force and I will emphasise the way in which institutionalism is not immune to problems of tying reasons for belief to our desires, even though Cowie calls epistemology categorical. This will also set the groundwork for chapter 3 where I consider epistemic analogues of Plato's Ring of Gyges.

## **2.1 The genuinely normative/merely institutional distinction.**

Cowie claims that merely institutional judgements are both categorical (they obtain independently of the desires of the agents to whom they apply) and are not subject to an internalism-based error theory. I will consider each of these claims in turn.

### **2.1.1 Institutional judgements are categorical.**

Cowie claims that merely institutional judgements are categorical. By this he means that our judgements "don't vary with the desires of the person to whom they apply" (Cowie, p. 11). For example, he considers normative rugby judgements to be categorical. He says that "when your teammate on the rugby pitch tells you that shouldn't pass the ball forwards, he doesn't mean that you shouldn't do so unless you want to" (p. 12) – the rules of Rugby prohibit everyone from passing the ball forwards no matter what they want.

Since Cowie thinks that normative rugby judgements are categorical and therefore not sensitive to the agent's desires, we might think that normative judgements about Rugby should not change as my perceptions of your desires change. However, if this were Cowie's position the following looks like a problem case:

The Unwilling Participant: I see you playing Rugby and initially make various judgements about what you ought to do, one of which is that you should not pass the ball forwards. However, I then learn that while you are playing rugby, you would much rather not be. Perhaps someone told you that you'd have a lot of fun playing rugby, but it turns out that you find rugby no fun at all. Furthermore, you owe nothing to the other players; they are cruel to you and laugh as you flinch away from tackles. On learning this about you it seems that my judgement of what you ought to do will change. I may now judge that it doesn't matter which direction you pass the ball in; what you ought to do is stop playing rugby and do something else.

This looks like a normative judgement about Rugby changing on finding out about your desires. This would seem to contradict Cowie's claim that normative judgements in rugby are categorical. Furthermore, there is no parallel case for moral judgements. When I judge that you morally ought not keep slaves, there really is nothing I could learn about your desires that would lead me to change my judgement.



However, contrary to appearances, The Unwilling Participant does not undermine Cowie's claim about categorical rugby judgements. This is because there are two different notions of normativity at work here. Cowie's claims about Rugby judgements not changing with perceptions of desires make sense when they are understood as concerned only with what the rules of Rugby have to say—the rules of Rugby do not differentiate between players based on their desires and, therefore, restating the Rugby judgement under consideration as “According to the rules of Rugby, you shouldn't pass the ball forwards”, it becomes clear that *this* kind of judgement will not change when I realise you don't want to play; your desires have no impact on what the rules of Rugby tell you to do.<sup>3</sup>

On the other hand, the normative judgements in The Unwilling Participant cannot be understood like this. It is not according to the rules of Rugby that I change my judgement; I still recognise that Rugby does not allow you to pass the ball forwards, I just now recognise there are factors other than the rules of rugby to consider. In this case I am not making judgements simply about what the rules of the game dictate, I am making genuinely normative, or all-things-considered judgements.

As I have said in section 1.5.2, different writers use different terms for this notion. All-things-considered assessments are a kind of final court of

---

<sup>3</sup> This may not be quite true; by their own lights the rules of rugby only apply to rugby players. It might be that your desires have an impact on whether you count as a Rugby player and thus whether the rules of Rugby tell you to do anything at all. If, for example, one of the things it takes to be a Rugby player is wanting to score tries and conversions etc then it might be that your desires do impact what the rules of rugby have to say to you. However, for now it is enough to say that the rules of rugby do not depend on your desires *given that you are a rugby player*.

appeal for assessing action; they are the level at which genuine mistakes are determined. I will follow Cowie in talking about ‘genuinely normative’ reasons and talk about ‘all-things-considered’ judgements. My all-things-considered judgements depend at least in part on the desires of the subject and thus my judgement in the Unwilling Participant can be sensitive to your desires.

We could represent this by using different ‘oughts’ for institutional Rugby judgements and genuinely normative judgements. We could say “You ought<sub>Rugby</sub> not pass the ball forwards” as shorthand for “According to the rules of Rugby, you shouldn’t pass the ball forwards.” and “you ought<sub>all-things-considered</sub> not pass the ball forwards” for the all-things-considered judgement. The former is always true when you are playing rugby and does not depend on your desires whereas the latter is sometimes true and is sensitive to changes in your desires.

Thus, Cowie claims that merely institutional judgements do not change with our perceptions of the desires of the agents to whom they apply. However, since it is a contingent matter whether the norms of a mere institution are outweighed by other considerations (or, as we shall see, have any force over you at all), this does not entail anything about our all-things-considered normative judgements.

This means that Cowie’s epistemic institutionalism can be framed as follows: Epistemic judgements concern what we ought<sub>epistemology</sub> to believe. The epistemic judgement ‘you should believe what your evidence supports’ can be restated as ‘according to the rules of

epistemology, you ought to believe what your evidence supports'. It is a contingent matter for Cowie whether the epistemic ought is outweighed by other considerations (or has any force at all) in our all things considered judgements.

### **2.1.2 Institutional judgements are not genuinely normative.**

Internalism-based error theorists object to moral judgements on the grounds that they concern a problematic kind of reason. They object to reasons which lack the right relationship to the desires of the agents concerned. In section 2.1.1 we saw that institutional normative *judgements* do not have a relation to the desires of the agents to whom they apply. This is only problematic for internalists if these judgements come with objectionable reasons. We also saw that the norms of Rugby can come apart from our all-things-considered normative judgements. We saw that we may lack all-things-considered reason to comply with Rugby norms. Cowie wants to claim further that there is a sense in which we may have no reasons at all to comply with merely institutional judgements.

This relies on the distinction between having *a* reason and having all-things-considered reason. Recall from section 1.4, I am understanding a reason to be a consideration that speaks for or against an action or belief. Having all-things-considered reason takes into account all your individual reasons, weighing them up against each other and giving an overall recommendation. One can have a reason to do something but lack all-things-considered reason to do it; your individual reason may be outweighed. For example, the fact you enjoy the zoo may be a reason to

go to the zoo, however, you enjoy the cinema more and thus you do not have all-things-considered reason to go to the zoo. We have seen that you can lack all-things-considered reason to follow merely institutional judgements. The claim now is that you can lack any reasons at all; there can be nothing at all that speaks for complying with the judgement.

Cowie's position is slightly more complex than this; when talking about epistemic judgements he agrees with Cuneo that "if an agent S is aware that some set of considerations implies that a proposition is true or likely to be true . . . then S has an epistemic reason to believe p" (Cuneo, 2007, p. 206) and says that epistemic judgements (which he thinks are merely institutional) always come with *epistemic* reasons. Thus, we have a notion of epistemic reason which is disconnected from an agent's desires that always accompany epistemic judgements. This is not to give in to the argument from analogy though. This is because, according to Cowie, epistemic reasons are merely institutional, rather than genuinely normative. That is, he distinguishes between genuinely normative reasons and merely institutional reasons as well as between genuinely normative judgements and merely institutional judgements. Merely institutional reasons just fall out of institutions. Rugby reasons, for example, are considerations which, according to the rules of rugby, speak for or against an action. Whether these are genuinely normative reasons – considerations that *genuinely* speak for or against your action – is an open question. This is clear in the following passage:

[S]uppose that I judge that there is excellent evidence for Tyson to believe some proposition, p. According to the epistemic

institutionalist view I am thereby straightforwardly committed to the existence of an epistemic reason for Tyson to believe p. But this is a merely institutional reason. There isn't necessarily any genuinely normative reason for Tyson to believe p.

(Cowie, p. 60)

Here it is clear that Cowie thinks it is possible for there to be *no* genuinely normative reasons to comply with a merely institutional judgement, whether that is in epistemology, in rugby or something else.

The clearest case where you have no genuine reasons to comply with rugby judgements is when you are not playing rugby. There may be other cases though:

Bystander: I see you playing Rugby and initially make various judgements about what you ought to do, one of which is that you should not pass the ball forwards. However, I then learn that while you are playing rugby, you have just seen bystander collapse on the sidelines, and you are best placed to help them. Furthermore, to do so will require breaking some rules of rugby. At this point my normative judgements change. While previously I judged that you ought to comply with the rules of rugby, now I have learned about your desires I judge that you ought to help the bystander, violating the rules if necessary.

One possible explanation of this case is that when you see the bystander collapse, the rules of rugby cease to be relevant to our assessment of your actions. We might say that in these cases any genuine reasons you had to

follow the rules of rugby are trumped by the moral emergency on the sidelines. It is not that you have genuine reasons to follow the rules which are outweighed, it is that you cease to have any genuine reasons at all to follow the rules.

This is in contrast with morality; consider the moral judgement “You should not murder”. If this is true then according to moral realists it follows that you have genuinely normative reasons not to murder (indeed, strong, or even conclusive reasons not to). This is an entailment of the Commitment Premise of the moral error theory (which, recall from section 1.5.2, I am taking for granted).

Thus, we may distinguish moral judgements, which are purported by realists to always generate genuine reasons, from mere institutions like Rugby, which only sometimes come with genuine reasons.

Cowie says that internalists are not concerned with merely institutional reasons that are disconnected from desires, only with genuine reasons. They are not troubled by the fact that we sometimes hold people to standards separate from their desires in the law, sports and so on. They are concerned with considerations that are supposed to count in all-things-considered justifications which bear no relation to the subject’s desires.

This is how mere institutions are shielded from error theory; they concern only merely institutional, categorical reasons and not genuinely normative reasons. Since merely institutional reasons are not problematic for internalists, neither are merely institutional judgements.

Moral judgements are still subject to their argument because they rely on genuinely normative, categorical reasons.

### **2.1.3 Summary**

If it is becoming hard to keep track of the different kinds of judgements and reasons we are working with, let us summarise the picture so far.

Firstly, there are moral judgements. These are the error theorist's intended target. Moral judgements purport to be such that we always have genuine reasons to comply with them, regardless of our desires. According to moral realists, whenever we are assessing someone's action all-things-considered, moral reasons must weigh in the balance. Internalism-based error theorist's object to moral judgements because they think that anything that is included in an all-things-considered assessment of an agent's action must have an appropriate relation to that agent's desires. Moral reasons, error theorists and many realists agree, lack this feature.

Secondly, there are merely institutional judgements. These are unproblematic for the error theorist. While the judgements themselves are independent of the desires of the agents to whom they apply, they do not generate reasons that must be weighed in the balance in all-things-considered assessments of people's actions. Therefore, internalism-based error theorists need not object to merely institutional judgements because they merely express rules or conventions that do not necessarily have any force over an agent. Merely institutional judgements come with merely institutional reasons. These are not considerations that always

merit inclusion in all-things-considered assessments of actions. These only speak for or against an action according to the institution in question. Whether they have genuine force over an agent is an open question.

We can begin to consider where epistemic norms fit within this picture. The proponent of the argument from analogy says that epistemic judgements are genuinely normative and thus vulnerable to error theoretic arguments. Cowie, on the other hand, argues that epistemic judgements are merely institutional—of a kind with judgements in sports and games, etiquette, fashion and the law. Cowie and his opponents are agreed that epistemic judgements are independent of the desires of the agents to whom they apply. They also both agree that epistemic judgements always generate epistemic reasons. However, the status of these reasons is at issue between them. The proponent of the argument from analogy claims they are genuinely normative reasons of a kind with moral reasons. Cowie thinks they are merely institutional—they only speak for or against beliefs according to the institution of epistemology. Reasons of the first kind are problematic for error theorists whereas reasons of the second kind are not. This, then, is how epistemic institutionalism shields epistemology from an internalism based error theory – it claims that epistemic judgements do not concern genuinely normative reasons.

## **2.2 The genuine normativity of mere institutions.**

In section 2.1 I explained how Cowie distinguishes genuinely normative judgements from merely institutional judgements. In the remainder of



this chapter, I will depart from Cowie slightly and consider the interaction between these notions. Mere institutions can become genuinely normative for us. When I am playing chess, I consider the merely institutional reasons of chess to be genuinely normative: I respond to them, accept criticism on the basis of them etc. This is more than simply my mistaking mere conventions for genuinely important considerations. It is possible for my genuine reasons to align with what an institution tells me to do, giving me genuine reasons to comply with its directives. My aim is to explore the ways this can happen and what this means for epistemology.

This is important because if Cowie is really to save epistemology from the grip of the error theorist's argument it will not be enough to say that epistemology is a mere institution—little more than a game for us to work out the rules of. He must further explain why epistemology is sometimes the right game to play.

Recall from section 1.6.2, one way of posing the argument from analogy is as diagnosing an inconsistency in the error theorist's position. Here is such an argument in more detail:

You say that you find moral judgements unacceptable because moral reasons are not appropriately related to an agent's desires and yet you present arguments and expect them to have force. The force they have is that of epistemic reasons and epistemic reasons also lack an appropriate relation to the agent's desires. Therefore, your position is incoherent—you must either

withdraw your complaint about moral reasons or accept that your arguments do not have the force you think they have.

I claim that error theorists think that their arguments have genuinely normative force. Error theorists do not think that the takeaway from their argument is simply “according to the rules of epistemology, you should believe moral judgements are systematically false” – you ought<sub>epistemology</sub> believe in moral error theory. They think it is “You should genuinely believe that moral judgements are systematically false” – you ought<sub>all-things-considered</sub> believe in moral error theory. It is not just that epistemology says that you should, it is that you are doing epistemology, the rules of epistemology apply to you and therefore you *genuinely* ought to believe that moral judgements are systematically false (or at least you have some genuine reasons to think so). If error theorists don’t think that we have genuine reasons to believe in error theory, then we might reasonably ask why they believe it and present arguments for it.<sup>4</sup>

This is not devastating to the institutionalist project. Cowie can accommodate genuine reasons overlapping with merely institutional reasons as we shall see. However, if error theorists think we have genuine reasons to believe in moral error theory, then they need a story about when the rules of epistemology apply to us, that is, when epistemology becomes genuinely normative. I will ultimately conclude that the routes into normativity for mere institutions are not up to the task. First though we must explore the options.

---

<sup>4</sup> This argument is closely related to the self-undermining arguments given against epistemic error theory I alluded to in section 1.6.2.

It is also important to emphasise that even if an error theorist adopts institutionalism as a way to avoid epistemic error theory, they are still forced to tie our *genuine* reasons to believe very closely to desires. Institutionalism allows them to say that we always have *epistemic* reasons to believe in accordance with our evidence, say, but not that we are always *all-things-considered* justified in following our evidence. Cowie is careful to distinguish epistemic institutionalism from a view according to which “epistemic reasons aren’t even categorical; they are constitutively dependent on someone’s possession of a relevant desire or goal.” (Cowie, p. 60). He says that this kind of view is not plausible. In the remainder of this chapter, I hope to show that Cowie’s view shares a lot of the problems of views like this.

Institutions can become genuinely normative in different ways. I will divide the space into two broad categories. The first is characterised by costs and benefits of following individual rules of the institution – it will look for things of genuine value that are at stake for particular rule violations. The central case of this category will be the law. The central case of the second category will be games and sports. I claim that these become genuinely normative through some kind of commitment the agent has to the institution itself.

### **2.3 Reasons to comply with individual institutional judgements.**

As I said in section 2.1.2, merely institutional reasons are considerations that speak for or against actions only according to a set of rules. We always have these reasons but someone who violates them or ignores

them entirely need not be making any genuine mistake. Sometimes, however, we get genuine reasons to comply with the directives of an institution—reasons which we *cannot* ignore without a genuine mistake. For example, Cowie considers legal judgements to be merely institutionally normative, however, we often have genuine reasons to follow the law. There are severe penalties imposed for breaking the law which give you prudential reasons to comply.

We can acquire genuine reasons to comply with the directives of almost any institution in a similar way. If someone incentivises you to follow the directives of an institution through punishments or rewards, then some or all of its merely institutional reasons can become genuinely normative for you. You need not care at all about the institution of the law as such to get these reasons as long as you care about avoiding the punishments or seeking the rewards.

This does not only occur via explicit incentives. I might get genuine reasons to comply with an institution in all sorts of ways. It may be that I wish to avoid social exclusion which is inflicted on people who violate the rules of a club, say. Or I might find I have moral reasons not to commit dangerous tackles that are prohibited by the rules of Rugby. What these methods have in common is that I have genuine reasons to comply even if I don't care about the institution itself.

The first thing to note about this route into normativity is that it is not all or nothing; depending on the specifics of the incentives, your genuinely normative reasons may not exactly line up with your institutional

reasons. If some laws are enforced but not all of them, for example, then individuals can lack prudential reasons to comply with the unenforced laws. If all genuine reasons to follow the law were prudential then this would leave agents with no genuine reasons at all to follow those laws.

A second feature of this route is that even when incentives are in place, institutional requirements will not always correspond to all-things-considered requirements; while the law requires me not to steal, it is not automatic, even when I am given strong incentives to follow the law, that I am all-things-considered required not to steal. If my children will starve if I do not steal a loaf of bread for them then even though I have genuine reasons not to steal coming from my desire to avoid prison, the reasons I have to protect my children can outweigh them at the all-things-considered level.

Error theorists can reconcile mere institutions becoming genuinely normative in this way provided the genuine reasons I have to comply with institutional judgements have the appropriate relation to my desires.

This contrasts with the normativity of morality. While we have genuine reasons to follow moral norms, according to moral realists, they do not only come from sources outside of morality. While we sometimes have prudential reasons to comply with moral judgements, we also purportedly have external moral reasons to comply with them even when it is not in our best interests.

If epistemology becomes genuinely normative in the way described in this section, then an error theorist who adopts institutionalism may be forced to accept that epistemology is not all or nothing: we will lack genuine reasons to conform to some epistemic norms when we lack desire-based reasons to do so. It may also turn out that even when we *do* have genuine reasons to follow epistemic norms they can be outweighed by other non-epistemic considerations. It is possible there are cases where this happens (for an example, see the discussion in section 3.2.1) but it is important to stress that error theorists must accept that this happens *whenever* we lack the right desires. In sections 3.3 and 3.4 I will present cases where this generates unintuitive results.

The takeaway from this discussion should be that if we get genuine reasons to follow epistemic norms in the way that punishments give us reasons to follow the law then our genuine reasons will be highly dependent on whether we care about or can avoid the consequences of our individual epistemic norm violations. If nothing you care about will be lost by believing contrary to your evidence, for instance, there will be no genuine reasons to do so.

We start to see that institutionalism is not so different from a theory in which epistemology is not categorical; our *genuine* reasons to believe go away when our desires do not align with the truth.

## **2.4 Reasons to opt into an institution.**

Not all institutions are best understood as becoming genuinely normative through this contingent alignment of incentives. Sometimes

the institution itself is necessary in explaining the source of our genuine reasons. Suppose we are playing a game of chess. It is not an important game, nothing substantial turns on who wins and neither of us will receive a major penalty if we break the rules. In this case it is not clear what genuine cost there is to breaking the rules. Nevertheless, we might ask whether we can be said to have genuine reasons to follow the rules of chess.

I think that we can. If we have reasons to play chess and we choose to do so, we obtain genuinely normative reasons to follow the rules. Chess is a rule governed activity – part of what it is to play chess is to follow the rules. Conforming to the rules of chess therefore is a necessary means to playing chess. Thus, when you have genuine reasons to play chess you have genuine reasons to follow the rules. What is given up when you violate a chess rule is that in the limit you don't count as playing chess anymore. Thus, you find yourself in violation of your genuine reasons.

A natural way for internalists to make sense of this is to say that to play chess is to adopt a certain goal. A goal like 'checkmate your opponent by making moves within the rules while avoiding being checkmated yourself'. When you have reasons to adopt this goal and you adopt it, institutional chess reasons become genuinely normative; you have a goal which you cannot satisfy without following the rules. We might express this by saying that chess has a constitutive goal—a goal which, once adopted, commits you to its norms. You get reasons to follow the rules by opting in.

We can tell a similar story for other sports and games too. This route is distinguished from the one above because reasons arise out of a commitment to the institution – adopting its constitutive goal or having reasons to play its game – rather than by some other means. In these cases, if you don't have reasons to care about the institution as such you do not get genuine reasons to follow the rules.

With the possibility of opting into an institution though comes the opportunity to opt out. When you lose your reasons to play chess, you lose your reasons to follow the rules. When it stops being fun you can stop playing.

Error theorists can accept the normativity of sports and games on this picture because whether an institution's reasons become genuinely normative is dependent on the agent opting into an institution. The agent must adopt a certain goal. In other words, the norms kick in only when they bare an appropriate relation to the agent's desires.

This also contrasts with the normativity of morality. There is no need to opt into the moral space; you are always subject to moral criticism regardless of the goals you adopt. Certainly, the opportunity to opt out of morality when it doesn't suit your desires would be unacceptable to the moral realist. Morality is arguably *most* important in cases where it tells you to do things you don't want to do.

This gives us an idea of what a game-like epistemology would be like. If epistemology is like a game that offers us some general benefit for playing (for instance, it allows us to efficiently achieve our goals), then



we should be able to stop playing when that benefit is outweighed by something else. In the way that in Unwilling Participant (introduced in section 2.1.1) we judge that you ought to stop playing rugby when you'd be better served by doing something else, the institutionalist who proposes epistemology becomes normative in this way is open to cases where we ought<sub>all-things-considered</sub> to stop playing epistemology and play something else instead. Again, we see that a merely institutional epistemology would tie our genuine reasons to conform to epistemic norms very closely to our desires.

There is an interesting parallel here with the distinction between act and rule utilitarianism. The route into normativity in section 2.3 corresponds to act utilitarianism. It suggests we ought to do exactly as our underlying moral and prudential reasons dictate, if they overlap with the rules of the institution, then so be it. The route in this section corresponds to rule utilitarianism – it suggests we do as the rules of the institution tell us because this will lead to satisfaction of our genuine underlying reasons that couldn't be had otherwise. In section 3.5.6 I will return to this distinction.

The salient point is that when our underlying reasons (which for the error theorist, means desires) aren't best served by conforming to an institution then we ought to opt out of it. We ought to stop playing rugby, say, and do something else instead. In the case of epistemology this means that when the winds of our desires blow unfavourably, we should opt out of following epistemic norms. In the next chapter we will see that this is an unintuitive result. In section 3.5 I will consider whether our

desires are such that we always have genuine epistemic reasons but it is important to note that Cowie's strategy, like the desire based epistemic reasons account he dismissed as implausible, leaves our genuine reasons to believe to live and die by our contingent desires. As we will see in the next chapter, this is bad for institutionalism.

## **2.5 Summary**

In summary, Cowie protects epistemology from the error theorist's argument by arguing that it is merely institutional. It is an open question whether we have genuine reasons to conform to epistemic judgements and therefore they do not concern reasons that the internalist rejects.

I argued in section 2.2 that error theorists think their arguments have genuinely normative force rather than merely institutional force – we are genuinely supposed to accept their conclusions. Therefore, I explored the ways we get genuine reasons to conform to merely institutional judgements. The key takeaway from this discussion is this: when all your genuine reasons come from desires, if you don't want to play a game anymore or if you do not care about the costs of violating an institution's rules, those rules cease to be genuinely normative for you. This is the fate that awaits epistemology on the institutionalist picture combined with internalism-based moral error theory.

On realist assumptions, it is now even clearer that morality cannot be understood as a mere institution. According to the realist, morality remains normative even when you don't want to play. It also remains normative even when there are no costs to rule violations that you care

about. This second point can be seen in Plato's Ring of Gyges story: given a ring that removes all the ordinary personal costs of wrongdoing, no one will ever know that it was you that did what you did, you will feel no guilt and so on, the realist says that it is still wrong to violate the directives of morality. As we will see in chapter 3, mere institutions are not similarly robust to removing consequences of our actions.

The way that you can ignore mere institutions when they do not line up with your desires is central to my critique of Cowie. Our genuine reasons to comply with institutions are not particularly robust – we only have them when they line up with our desires. While Cowie distinguishes his view from one on which epistemology is not categorical, when we consider genuinely normative reasons, his view turns out very similar to a view like this. In the next chapter I will argue that epistemic judgements cannot be ignored in this way; we appear to have reasons to comply with them even when we have no other reasons to do so.

### 3 The Epistemic Ring of Gyges

Cowie's aim was to show that error theorists could accept the truth of epistemic judgements while denying the truth of moral judgements. In section 2.2 I argued that the real challenge is to show that error theorists can explain the genuine normative *force* of epistemic judgements while denying the normative force of moral judgements. We have seen that Cowie's strategy has the consequence that for error theorists any genuine reasons we have to comply with epistemic judgements must come from sources outside of epistemology. For error theorists this must come from reasons which are appropriately related to an agent's desires. We have seen two ways that this could come about. One is that there may be non-epistemic, desire-related reasons that speak for compliance with particular epistemic judgements (as with the normativity of the law). The other is that there may be non-epistemic, desire-related reasons to opt into the whole epistemic institution (as with the normativity of games and sports). Either way, any genuine normativity of epistemic reasons will have to be tied very closely to the desires of the agents at hand.

The goal of this chapter is to provide cases where this model is inadequate, in other words, cases where we have no moral or prudential reasons to comply with epistemic judgements and yet we still seem to have genuine reasons. I will frame my investigation around Plato's discussion of the Ring of Gyges in book 2 of the Republic (Plato, 1974, p. 360).

### 3.1 Plato and the Ring of Gyges

The Ring of Gyges is a fictional ring that grants the wearer invisibility. In the story, Gyges uses the ring to amass power and kill the king. The story is introduced by Glaucon to argue that when the social and legal costs of injustice are removed (we will not be punished, nor will we be shunned by society) we have no reasons to behave justly. He claims that anyone given such a ring would use it in this sort of way. Plato resists the view that he expresses through Glaucon. I will not spend much time on exactly what Plato wants us to take from the story; it is not clear cut that he would explain it in terms of external reasons to be moral (he argues that we will not be *fulfilled* if we do not behave justly which could be interpreted as an appeal to internal reasons). What I am interested in is what a modern-day moral realist can take from the story.

The purpose of the Ring of Gyges, I take it, is to remove prudential reasons to be moral and test our intuitions about what we ought to do. In modern times with DNA and fingerprint technology we might have to beef up the ring's powers a little to make the point: we may need a ring that will guarantee you will leave no trace of your crimes, there is no chance of being found out, and furthermore, we may need a ring that will free you from guilt and other psychological distress that results from your immoral actions. But once the ring has been strengthened the same question that interested Plato remains; once all the prudential, self-interested reasons to be moral are removed, do we still have reasons to be moral? Should you behave morally even when you possess the Ring of Gyges?

The moral realist answers both questions ‘yes’. I think they have the intuitions on their side; it does not seem that it becomes acceptable to kill, rape and steal just because you personally will be freed from consequences. The moral realist can use this as an argument that there are external reasons to be moral that do not go away when there’s nothing at stake for you. This is perhaps quite a good argument against internalism-based moral error theory; the intuitions that it’s still wrong to murder people when you know you will get away with it are quite strong. Perhaps stronger than the complex intuitions about what kinds of reasons exist that the error theorist relies on. However, as I said in section 1.2, this kind of challenge to the error theorist is off the table; the error theorist is aware that this is the kind of bullet they have to bite. I am interested in finding epistemic bullets that they have to bite as well. Their scepticism cannot remain localised.

The Ring of Gyges provides a template for this; we need to find a case where it looks like we have reasons to follow epistemic norms even when we have no moral or prudential reasons to do so (either because there just aren’t any of these reasons at play or because they have been artificially removed by a magic ring). This would be an epistemic Ring of Gyges case.

We might call a normative judgement that we have reasons to conform to even when reasons from other sources are removed, *robust* against the Ring of Gyges. Realists think that moral judgements are robust in this way. I will argue that merely institutional judgements are not. If we can provide epistemic Ring of Gyges cases, this will provide evidence that

epistemic judgements are robust against the Ring of Gyges. This will be a mark against epistemic institutionalism.

### **3.2 Barriers to a satisfactory epistemic Ring of Gyges case.**

While I believe that the normativity of epistemic judgements is importantly similar to the normativity of moral judgements, there are also differences that present challenges to us here.

#### **3.2.1 Epistemic requirements are not all-things-considered requirements.**

The first barrier is that it is less clear that epistemic requirements are genuine requirements than it is that moral requirements are genuine requirements. What I mean by this is that when you are morally required to do something, many moral realists think that you are all-things-considered required to do it; if you are morally required to refrain from killing in cold blood, then you are all-things-considered required to refrain from killing in cold blood. Whereas it is less clear that if you are epistemically required to believe what your evidence supports then you are all-things-considered required to believe what your evidence supports. Consider the following counter-case, adapted from Cowie: Tyson is a boxer preparing for a fight. His evidence supports a low credence that he will win and so we might think he is epistemically required to have that low credence. However, he will have a better chance of winning if he believes wholeheartedly he will win. As such, it appears rational for him to believe that he will win despite his epistemic reasons. One way to interpret this is as a case where prudential reasons outweigh epistemic reasons. There may also be cases where moral

reasons outweigh epistemic reasons; sometimes we have moral reasons to trust people even when they have shown themselves in the past to be unreliable. You ought<sub>epistemic</sub> to distrust their testimony but you ought<sub>all-things-considered</sub> to take them at their word.

This creates a challenge to giving an epistemic Ring of Gyges case. The lesson that moral realists take from the original Ring of Gyges is that you ought<sub>all-things-considered</sub> not kill someone in cold blood regardless of the personal benefits it brings you. However, in light of cases like Tyson's it may not be the case, even under normal circumstances, that you ought<sub>all-things-considered</sub> not make epistemic errors regardless of the prudential or moral benefits. This means that there may be fewer cases where our intuitions about what we ought<sub>all-things-considered</sub> to do align with our epistemic judgements when we do not have moral and prudential reasons to comply; moral and prudential reasons against compliance may outweigh epistemic considerations.

### **3.2.2 Moral interference**

A second barrier to a satisfactory epistemic Ring of Gyges case is the complications caused by moral reasons. In the original Ring of Gyges case, we removed prudential reasons for moral compliance using the ring and thus were able to find a normative intuition that was best explained, at least according to moral realists, by genuine moral reasons. Genuine epistemic reasons, if there are any, were not relevant in that case because epistemology is concerned with beliefs and not ordinary action. However, moral reasons, if there are any, can interfere with what we have reason to believe. I have already mentioned moral reasons to



trust epistemically shaky testimony. However, we may also have moral reasons *in favour* of epistemic compliance. It is particularly important, we might think, to believe that cutting the brakes of someone's car will cause a crash because it is morally important that you don't cause someone to crash their car.

The challenge this creates for providing an epistemic Ring of Gyges case is that if we remove prudential benefits for epistemic compliance then there will still be cases where our judgements that we should comply with epistemic norms can be explained via moral reasons. What this means in practice is that when we find an intuition that we ought to comply with epistemic norms, we may find we are simply pulling on moral intuitions for believing well. This is not the strategy I want to pursue. Error theorists insist that our moral intuitions are in error, therefore, they will reject appeals to moral intuitions from the outset. Recall the appeal of the argument from analogy was that it promised to meet the error theorist on their own terms. Relying on moral intuitions in favour of epistemic compliance will not do this. This means we must be careful when choosing our examples.

### **3.2.3 Doxastic involuntarism**

The third barrier is the interference of psychological limitations. We generally have quite a large degree of control over our actions. While Glaucon is wrong to claim that everyone *would* act immorally given a ring of Gyges, it is plausible that everyone is *capable* of acting immorally under those circumstances. On the other hand, we have much less voluntary control over our beliefs. We cannot just decide to believe

something because it suits us; even if I offer you £1 million if you believe that the sky is falling it is very hard and perhaps impossible for you to do so while in possession of exclusively evidence to the contrary.

This is a barrier to providing a satisfactory epistemic Ring of Gyges case because certain epistemic practices may look strange just because they are not the sorts of thing anyone can do. This provides an alternative explanation for our intuitive revulsion to certain belief practices. It is not that you *shouln't* believe in these ways, it is just that you physically *couldn't* believe in these ways.

I will say more about how to avoid this barrier in a moment.

### **3.3 Epistemic Ring of Gyges case 1: Trivial Falsehoods**

I have set out some of the challenges we must overcome to provide an epistemic Ring of Gyges case. I have two proposals for possible cases which I will explore in the following sections. I will set out the cases and assess how satisfactory they are in light of the challenges in section 3.2.

The first case concerns our reasons not to believe trivial falsehoods. Recall from section 1.7.2 that when making his trivial truths argument for institutionalism, Cowie argues that some truths are so trivial that we have no moral or prudential reason to believe them even when we have strong evidence. This case concerns the converse, that is, our reasons not to believe trivial falsehoods. Recall, also from section 1.7.2, Papineau's case of forming beliefs about the kings of Assyria. It is plausible that nothing morally or prudentially turns on an agent's belief about this. As far as prudence or morality is concerned you can believe whatever you

want about it, no matter your evidence. However, there are some beliefs that still look very strange for you to form. It would be very strange, for instance, for you to look at a list of the kings of Assyria from a reliable source and form the belief that all the kings of Assyria were called 'Nigel', despite what you had read on the page. If someone did this it seems to me that the person judging would have made a genuine error, even though nothing prudentially or morally turns on this belief.

However, perhaps there was something illegitimate about how we set up the case. We all have limited powers of deduction and a finite memory. There is, therefore, a small prudential cost to forming any belief. Thus, perhaps we were wrong to say that there are no prudential reasons not to form this belief. Perhaps this can be used to explain the genuine error – you are genuinely at fault because you are using up limited resources unnecessarily.

This might be the point at which to introduce a magic ring. Given a ring that removes the prudential costs of belief formation – that gives you infinite memory and the ability to judge instantly and costlessly – would you still be acting against your all-things-considered reasons in judging that all the kings of Assyria were called 'Nigel'? I think you would. I think you have reasons not to inflict this kind of deliberate damage to your belief system.

This case is limited in two ways. Firstly, by the strength of the intuitions: it is a much smaller bullet to bite that there is nothing wrong with these weird belief forming practices when nothing turns on them than it would

be for there to be nothing wrong with killing. Secondly, in this case the psychological limitation point kicks in: it seems implausible that anyone *could* form this kind of outlandish belief. Both limitations give us ways to call into question the intuition.

One way to combat worries caused by psychological limitation is to simply strengthen our magic ring. Just as in the moral case I suggested a ring that removes guilt and other psychological limitations on acting immorally may be necessary to provide a prudentially unpolluted intuition, in this case too we may allow that the ring removes psychological barriers to belief formation. Given a ring that allows you to form beliefs at will, even in the face of huge amounts of evidence, should you use its powers to believe that all the kings of Assyria were called 'Nigel'? Again, I think you should not. Epistemic institutionalism combined with moral error theory does not have the resources to explain this.

Of course, denying this intuition is not necessarily a devastating cost for the institutionalist. However, I think it is a cost nonetheless. In this case Cowie has to say that someone who used the powers of such a ring would have made an *epistemic* error but since nothing morally or prudentially turns on the question they would not have made a genuine error. All-things-considered, they have no reason not to form beliefs like this. This seems to me a point against his position.

### 3.3.1 Responding to the argument from trivial truths

This case makes clear why Cowie's argument from trivial truths fails.

Recall from section 1.7.2, the core claim of that argument:

Trivial truths (epistemic): Epistemic norms, much as the norms of etiquette, fashion, sports and games, and the law, are the kinds of things that, unless there is a moral or prudential reason to follow them, don't generate genuinely normative reasons.

(Cowie, p. 66)

Cowie thinks that trivial truth cases provide evidence for this. Due to the nature of his examples, he must be considering an epistemic norm along the lines of 'believe everything your evidence supports'. He argues that in trivial truth cases where we have no moral or prudential reasons, we do not have any genuine reasons to comply with this norm.

However, now consider a different epistemic norm: 'believe *only* what your evidence supports'. This is at least as plausible as a requirement of epistemology as 'believe everything your evidence supports'. But in this case, we find that, when we have no moral or prudential reasons, it still seems like we shouldn't violate this norm. This is contrary to Trivial truths (epistemic).

The kings of Assyria being called 'Nigel' is one counter-case. We can also provide cases that correspond to the other examples in the argument from trivial truths. Take Leite's example of standing outside a convention centre as people accompanied by dogs stream out (see section 1.7.2). Cowie agrees with Leite that despite the evidence, there is no particular

reason you should form the belief that the next person to leave will be accompanied by a dog. I claim that, intuitively, there *is* genuine reason *not* to believe that the next person to leave will be accompanied by a rhinoceros, even if nothing morally or prudentially turns on this belief.

Similarly, in Papineau's blades of grass case, while you may have no genuine reason to form an accurate belief about the number of blades of grass on your lawn, I think you *do* have genuine reason *not* to believe that it consists of just *one* blade of grass after walking across the lawn each morning and feeling the grass underfoot.

For Trivial truths (epistemic) to be true, it must be true for all epistemic norms – these cases suggest it is not true, or at least is very unintuitive, for the norm 'believe *only* what your evidence supports'. We shall see in section 3.5 that merely institutional normativity cannot explain this phenomenon and, therefore, this is evidence against epistemic institutionalism.

### **3.4 Epistemic Ring of Gyges case 2: man-made climate change.**

The case in the section 3.3 relied on intuitions about quite trivial propositions. These are not that important and thus it is not a huge cost if the error theorist must dispense with these intuitions. In this section I will consider a case where the intuitions in favour of epistemic compliance are stronger at the expense of straying into more morally polluted waters.

Wishful thinking, wilful ignorance and self-deceit are often considered not just epistemic vices but genuine vices. However, we regularly have desires that will be served by these practices: the weight of many truths about the world is very mentally taxing. The existence of man-made climate change, the extent of racial, gender and class inequality and the fact that there aren't simple solutions to these problems are just a few examples. One only has to turn on the news to find many more. It can be very troubling to believe these things and most people quite reasonably don't want to experience the mental anguish. In these cases, remaining blind to the truth is an excellent means to this end. This is not merely fictional either. Populist political movements are fuelled by people exhibiting epistemic vices of this sort. This provides the opportunity for another epistemic Ring of Gyges case.

For instance, our evidence supports a belief in man-made climate change. Top scientists have studied the topic in detail and come to that conclusion. Their findings have been heavily publicised by reliable sources, and we have read those reports. We clearly have epistemic reasons to believe in man-made climate change. Do we have genuine reasons? I think that we do. Intuitively, we ought<sup>all-things-considered</sup> believe in man-made climate change.

However, the question is polluted by both prudential and moral reasons. Climate change will have a large impact on our lives and the lives of others. Our beliefs about it will have an impact on our actions; if we believe that the things we do affect the progression of climate change then we are more likely to change those actions. Since it matters both

morally and prudentially what we do with regards to climate change, there are genuine reasons to conform our belief to the evidence in this case.

Despite the interference of these reasons, I do not think that they are devastating to the case. Firstly, let us consider prudential reasons. As I have said, beliefs about climate change can be mentally troubling. The negative impact of these beliefs on us are quite significant. Furthermore, the effects of our actions as individuals have very small effects on how climate change will impact us; it is possible that all the recycling I ever do will have no impact at all on whether my house is destroyed by rising sea levels. Therefore, on the balance of my prudential reasons alone it may well be better for me to deceive myself in this case. My desires are better satisfied by breaking epistemic norms than keeping them.

Moral realists can appeal to moral reasons to explain the intuition in favour of following epistemic norms here. Many moral realists think that we morally ought to take an interest in our impact on the natural world. However, error theorists have no such option available. It seems that they are forced to say that even in the face of huge amounts of evidence, it is possible that we ought<sub>all-things-considered</sub> not believe in man-made climate change.

This is a highly unintuitive result. However, perhaps it is par for the course for an error theorist. Error theorists are happy to accept that people can lack genuine reasons not to throw cigarette butts at children,



so perhaps it is no great cost to add to this that we might have no genuine reasons to believe in man-made climate change.

Nevertheless, I think that it is an extra cost. Error theorists might have hoped that while morality turns out to be a sham, at least no one can deny the force of scientific evidence when it is put in front of them without genuine error. The man-made climate change case suggests they can't. They are forced to call the behaviour of a climate change denier who goes on polluting perfectly rational.

To put this point explicitly in the form of a Ring of Gyges case, if you had the power to deceive yourself, take on a false belief whenever it suited you, would you have reasons not to use this power? The error theorist is forced to say that there are no such reasons. This conflicts with the intuition that there are things it matters that we believe in, including man-made climate change.

#### **3.4.1 Responding to the argument from ordinary explanation**

This showcases a limitation of Cowie's argument from ordinary explanation. Recall from section 1.7.1 Gary the epistemic sceptic questioning why he should care about epistemic norms. His teacher tried to respond to him as follows:

Well, you understand why you should eat your vegetables? You may not care about these epistemic standards as such, but you do care a lot about other things. And you're more likely to get what you want if you have warranted beliefs.

(Railton, 1997, p. 294)

Gary might persist saying,

Suppose that believing in man-made climate change is greatly mentally troubling for me, furthermore, I don't care about the truth as such of my belief about climate change. Why shouldn't I ignore all the evidence that is presented to me and be dogmatic in my belief?

In this case, the teacher's response does not work: the things Gary cares about (including his mental health) are not best achieved by complying with epistemic norms. He may even care about the environment and will be more likely to act to protect it if he believes in climate change, but judge that the impact of his false belief on his mental health is far greater than the impact his true belief would have on the climate. Therefore, as far as Gary is concerned it is far better to violate the norm than follow it. The teacher might try and appeal to moral reasons to push Gary towards epistemic compliance. But if Gary is also a moral error theorist, then this is not an option. It seems like there should be *something* that the teacher can say to Gary here but according to internalists there is not. We can imagine the teacher getting exasperated at this point. What he would really like to say to Gary, challenging him on why he should believe that there is man-made climate change is 'because it's true!' or 'because scientists have thought long and hard on evidence gathered over a very long time and come to that conclusion' or perhaps 'it doesn't much matter that you personally will gain from your false belief – this pales in comparison with the evidence'. However, for Cowie this is to appeal to considerations from within the institution of epistemology; this is not

how the genuine normativity of mere institutions works. According to Cowie, Gary should remain unsatisfied.

What this case shows is that the moral error theorist's scepticism is very expansive, even if they adopt institutionalism. They cannot simply appeal to truth or evidence as such in a justification of a belief; they must appeal to something suitably related to the desires of the agent in question. You cannot be an internalism-based error theorist and think that other people must be moved by the power of truth and argument if they don't want to be. Epistemic institutionalism cannot help the error theorist prevent the spread of their scepticism in this way.

### **3.5 Necessary desires**

In the previous two sections, I argued that there are cases where we appear to have genuine reasons to conform to epistemic norms even when we have no moral or prudential reasons to do so. On a simple reading of epistemic institutionalism there is no way to accommodate this. As I argued in sections 3.3.1 and 3.4.1, it is inconsistent with both Cowie's arguments for institutionalism. In this section I will explore whether there is a way for institutionalism to accommodate this observation after all.

#### **3.5.1 Routes into normativity for mere institutions.**

In sections 2.3 and 2.4 I introduced two ways that we get genuinely normative reasons to conform to the directives of mere institutions. The first was in the way that the law becomes normative – we do not opt in, we simply have moral or prudential reasons to follow the laws due to

alignment of our incentives. The second was that we have moral or prudential reasons to opt into an institution; reasons to adopt its constitutive goal for the sake of some benefit that it will bring us. This happens most obviously with sports and games where we follow the rules (for the most part) to have fun.

Neither of these routes are particularly robust against Ring of Gyges cases. When we lack prudential or moral reasons to follow the laws of a country (for instance when the law is not enforced), we lack genuine reasons. Similarly, when we lack moral or prudential reasons to play a game (for instance because there is a moral emergency on the sidelines), we can opt out and reasonably stop paying attention to the rules.

Therefore, if the institutionalist is to accommodate the epistemic Ring of Gyges cases described in sections 3.3 and 3.4, they will need to call into question their setup. Is it really the case that we have no desires that favour following epistemic norms in these cases?

The alternative would be to deny the intuitions in epistemic Ring of Gyges cases. I have already discussed the costs of this, therefore, I will consider how an error theorist might argue that there are reasons they can accept to follow epistemic norms even if it doesn't look like it at first.

### **3.5.2 The necessary desires strategy**

Recall from section 1.5.3 that internalists think that all our genuine reasons come from our desires or our desires suitably idealised. It might be, therefore, that there are desires we have necessarily, or would have necessarily were we idealised in the right ways. Cowie mentions this

strategy when discussing how an internalist might make sense of the apparently categorical demands of morality (see section 1.5.3). He suggests that it is implausible in the case of morality. He says that the internalism-based error theorist can shift the burden of proof to someone who claims there are such necessary desires in light of the fact it is possible to imagine someone who lacks them (Cowie, p. 14). Perhaps the prospects of the necessary desires strategy for epistemology are better. It could be that certain epistemic desires have a better claim on necessity and thus we have internalist appropriate epistemic reasons even in epistemic Ring of Gyges cases. I will explore the prospects of this in the remainder of this section but argue it is ultimately unsuccessful.

### **3.5.3 Necessary epistemic desires**

A first thought is that we might try to appeal to the agent's desire to pursue true beliefs and avoid false ones. If every agent had this goal, at least for some of their beliefs, then we could appeal to this when locating the source of their genuine epistemic reasons, even in apparent epistemic Ring of Gyges cases.

The problem is that many agents have genuine epistemic reasons and it is implausible that all of them have the requisite desires. Consider the following apparent counterexample:

Populist Politician: Consider a politician who adopts views not because they are true but because they are popular among their base supporters. Such a person does not seem to care about the truth of their beliefs as such. They seem much more interested in

the power that adopting and advancing those beliefs allows them to amass. It seems that such an agent at least has some reasons to believe truly despite lacking desires that their beliefs be true. For example, it seems they are still subject to genuine criticism for ignoring their evidence about man-made climate-change.

Furthermore, it is not clear that everyone with genuine reasons to conform to epistemic norms is *capable* of complex second order desires like the desire that their beliefs be true; For instance, young children appear to have genuine epistemic reasons:

Optimistic 5-year-old: imagine you are stood outside Leite's convention centre with your 5-year-old. You are watching the people coming out with their dogs. You ask them what animal they think that the next person to leave will have with them and they say 'a rhinoceros'. You would be licenced in correcting them. They may not need to believe that it will be a dog (it may be perfectly acceptable for them to say they don't know), but they genuinely ought not think it will be a rhino. A 5-year-old, however, may well lack a firm concept of 'belief', 'truth' and 'evidence' and thus it is implausible to attribute them with desires about these concepts. Nevertheless, it seems they ought<sub>all-things-considered</sub> not believe the next person will have a rhino with them. This judgement, therefore, cannot plausibly be grounded in an actual desire of theirs that their beliefs be true.

It is implausible, therefore, that the source of genuine epistemic criticism can be found in an actual desire of agents to have true beliefs and avoid false ones. Both populist politicians and young children don't seem to have strong desires to have true beliefs, either because they are more focussed on the pursuit of power or because they lack the concepts for such desires.

#### **3.5.4 Idealised epistemic desires**

Perhaps though this is where an internalist might appeal to idealised versions of an agent's desires. While populist politicians and young children don't seem to have strong desires to have true beliefs, either because they are more focussed on the pursuit of power or because they lack the concepts, perhaps if they were idealised then they *would* have these desires, therefore, the internalist can still grant them genuine reasons.

This is an interesting proposal. Epistemology may stand in better stead than morality on this count. Many internalists, as part of their idealisation, are willing to grant agents perfect information about all the relevant facts. Joyce, for instance, speaks favourably of a view of normative reasons which he attributes to Smith (Smith, 1994) according to which:

an agent has reason to  $\phi$  if a fully rational version of that agent (with all and only relevant true beliefs) would desire that the actual agent  $\phi$ s.

(Joyce, 2001, p. 82)

According to this view, the idealised version of an agent has true beliefs in epistemic Ring of Gyges cases: the idealised version of the Populist Politician believes in man-made climate change and the idealised version of the Optimistic 5-year-old does not believe the next person to leave will be accompanied by a rhino. This is a good start, however, Smith does not say that you have reasons to do as your idealised self would do; he says that you have reasons to do as your idealised self would *want you as you actually are* to do.<sup>5</sup>

It is not clear that the idealised Populist Politician would want the actual Populist Politician to have true beliefs: while the idealised populist has true beliefs themselves, it is still not necessary for them to care about true beliefs as such. The idealised populist may also want the actual populist to achieve power at the expense of the climate and see that that goal will be better served by the populist continuing to deceive themselves; their unjustified belief about climate change is serving them incredibly well in their pursuit of power. In other words, it doesn't look like this idealisation will help in securing a necessary desire to pursue true beliefs and avoid false ones.

### **3.5.5 Necessary features of non-epistemic desires**

Perhaps then inescapable epistemic reasons come from a less explicitly epistemic desire. In making his argument from ordinary explanation, Cowie argues that complying with epistemic norms is a good strategy for

---

<sup>5</sup> It is worth noting that Smith himself does not think that his internalist view leads to moral error theory. Joyce, however, thinks Smith provides a plausible account of normative reasons but that it *does* lead to a moral error theory. It is Joyce's view that is relevant to my enquiry.



achieving almost any goal (Cowie, p. 62): complying with epistemic norms tends to maximise the accuracy of your beliefs, and this is instrumentally useful for achieving almost anything you could want to achieve. In this case, we do not need to find a single goal that conforming to epistemic norms will serve. It is just that confirmation to epistemic norms serves so many goals that agents almost always have reasons to comply with them.

This strategy doesn't seem to work in either the case of the Populist Politician or the Optimistic 5-year-old though. There is apparently nothing that the populist wants that would be served by a true belief about climate change, at least nothing that isn't outweighed by their desire for power. It is also not clear what the 5-year-old has to gain from not believing there will be a rhino. Nor does it seem they would care about these things were they idealised.

### **3.5.6 General benefits of epistemology.**

Thus, it looks difficult to find a desire that favours following epistemic norms in particular cases. Perhaps then we should be looking for some general benefit to following epistemic norms, even when the truth of particular beliefs doesn't matter to us.

While the Optimistic 5-year-old has no desires impacted by their outlandish belief about a rhino, here is a plausible explanation for why we are licenced in criticising them in this case: their belief about the rhino may not matter as such, it is just that it exercises an unreliable

inference. A kind of inference that, if we encourage it, will lead the child to form bad beliefs in cases where it does matter.

This suggests a strategy along the lines of our reasons to opt into sports and games discussed in section 2.4. Perhaps opting into epistemology as a whole offers us benefits that could not be had otherwise; perhaps being an epistemically responsible believer, even in trivial cases, is necessary to give you reliable beliefs in non-trivial cases.

To see why this is an unsatisfactory explanation of epistemic Ring of Gyges cases, recall the parallel with rule utilitarianism I introduced in section 2.4. Rule utilitarians hoped to overcome some of the problems of act utilitarianism by positing rules which, if followed, lead to greater overall good than reassessing our reasons in each case. A major problem for rule utilitarians is what they can say in individual cases where it is clear that breaking a rule will produce more good than following it; in these cases it seems that rule utilitarianism collapses back into act utilitarianism.

A similar problem arises here. The populist politician finds themselves in a situation where it is relatively clear that following the rules of epistemology and believing in man-made climate change will do very little for their goals and self-deception will do much better. Even if their reasons to follow epistemic norms come from the general benefits of playing epistemology, the specific benefits of violating epistemic norms in individual cases will regularly outweigh these benefits.

Recall from section 2.4 that when we lack reasons to play a game or sport we can opt out; sometimes the benefit of playing the game is not enough to make it worth our while continuing to play. This means that the Populist Politician should be able to opt out of playing the epistemic game when it no longer suits them. This is the wrong result.

### **3.5.7 Takeaways**

Just as it seemed possible to imagine someone who lacked the desires required to give them internalist moral reasons, it is possible to imagine someone who lacks desires required to give them internalist reasons to follow epistemic norms in epistemic Ring of Gyges cases. I have considered a number of responses the institutionalist could make to this but found them all unsatisfactory. It does not seem that internalism-based error theorists can find genuine reasons to respect our intuitions in cases of trivial falsehoods (3.3) and man-made climate change (3.4).

## **3.6 Summary**

In this chapter I have provided two cases that I take to be epistemic analogues of the Ring of Gyges. In section 3.3 I considered trivial falsehoods like ‘all of the kings of Assyria were called ‘Nigel’ which, I claim, even when there are no prudential benefits, we should not believe. In section 3.4 I considered mentally troubling truths like those around man-made climate change. In these cases, prudence favours self-deception and yet it still seems wrong to ignore our evidence. I argued that institutionalism is not equipped to explain these intuitions. In section 3.5 I considered the possibility of explaining them via necessary desires but found these proposals to be unconvincing. The takeaway is

that if the error theorist is to avoid epistemic error theory by adopting epistemic institutionalism, they will have to bite some bullets. They must accept that self-deception is regularly the rational thing to do when we are confronted by troubling truths and that there are no reasons not to do substantial damage to your belief system if given the chance to believe trivial falsehoods.

A notable feature of epistemic institutionalism (which, in light of the arguments of this chapter, is a substantial strike against it) is that it does not allow agents to appeal to the importance of truth simpliciter in justifying or criticising belief practices. To be genuinely subject to an epistemic reason an agent needs to have desires that are serviced by following epistemic norms. In this chapter I have given cases where agents lack these desires and I have argued that unintuitive results follow.

## **4 Opting out of epistemology**

In the previous chapter I argued that there are certain normative intuitions which the error theorist can't explain via epistemic institutionalism. In this chapter I will consider a different angle of attack. I will explore a collection of conceptual problems based on the fact that agents can't opt out of epistemology in the way they can opt out of sports and games.

### **4.1 The Epistemic Sceptic**

When you are playing a game and the rules tell you to do things you don't like, you can step back and consider whether it is worth carrying on playing the game. Following the rules, at least to an extent, is a necessary means to playing, however, playing the game at all is optional. Thus, when you don't like what the rules tell you, you can seriously consider whether to opt out of the game and its rules. Similarly, before you play a game you can look at the rules and decide whether to submit to them or not; you get to decide whether this is a game that is worth playing.

When you are engaging in these processes you are treating the game as a mere institution. During the deliberative process, the fact that the rules tell you to do something does not count as a reason to do it. While engaging in this kind of deliberation, you treat the rules merely as a list of prohibitions and recommendations and decide whether, on the balance of your genuine reasons, it is worth following them. In this section I will consider the epistemic analogue of this. When epistemic norms tell you to do things you don't like, can you step back from them

and reconsider whether to play the epistemic game? I shall argue that you cannot. To successfully explain the normativity of an epistemic judgement it must be that the parties to the explanation are already committed to some epistemic reasons. This, I shall say, is evidence that epistemic judgements are not merely institutional.

#### **4.1.1 Gary, the epistemic sceptic**

Let us return to Gary, the epistemic sceptic who we first met in section 1.7.1. Gary asked us the following:

[P]hilosophers seem to agree that there are certain standards of belief. . . They spend all their time disagreeing about what exactly these standards are, but they just seem to assume that we'll want to follow them. Suppose I don't? What can they say to me?

(Railton, 1997, p. 294)

Gary sets up a game for us; he claims to be suspicious of epistemic reasons and wants us to convince him to follow them. The way he sets up his game, he suggests he is looking for desires of his that will be satisfied by following the rules of epistemology.

At first glance, this doesn't seem too hard: warranted beliefs are pretty useful after all. Surely Gary has some desires that warranted beliefs will help him with. Thus, Cowie is right that it seems natural to respond as his teacher does:

Well, you understand why you should eat your vegetables? You may not care about these epistemic standards as such, but you do

care a lot about other things. And you're more likely to get what you want if you have warranted beliefs.

In section 3.4.1 I argued that playing Gary's game is actually harder than it looks: we cannot always find desires that will be advanced by following epistemic norms when it seems there are reasons to do so. However, even if we *can* find the right desires, should Gary be satisfied with this? One way of understanding his question is as paralleling the challenge of convincing the amoralist to be moral. He is proposing to adopt a position where he doesn't see the force of epistemic reasons and have us talk him out of it. However, if this is the game then we are certain to lose: we have no chance of talking someone out of a position like this – certainly not via reasonable argument. We can present Gary with all the arguments we like that he ought to care about epistemic considerations, however, since he is an epistemic sceptic, he does not yet recognise a normatively significant difference between a good argument and a bad argument. Thus, however good an argument we give him, by his own lights, he should not accept it; we haven't told him yet why he should accept the conclusions of good arguments.

This is unlike the challenge of the amoralist. In that case at least we could imagine how the conversation goes: we say you shouldn't kill people, the amoralist says 'why not?', we say 'because their lives matter', they say 'not to me' and so on. With the epistemic sceptic a conversation cannot even get started. We say 'you should believe what your evidence supports' they say 'why?'. Suppose we respond 'because you want lots of things and having warranted beliefs will help you get them'. We might

imagine them considering this and deciding whether to accept it or whether to object. However, this is already going too far. Let's break down what we have said to them. In more detail we might have said:

- (1) You want X.
- (2) Warranted beliefs are a necessary means to getting X.
- (3) (1) and (2) jointly imply you have reason to have warranted beliefs.

Therefore,

- (4) You have reason to have warranted beliefs.

(3) is the kind of principle that a simple reasons-internalist would accept. According to them your desires are the source of your reasons, and it is a plausible general principle that if you have reasons to X and Y is a necessary means to X then you have reasons to Y.

This looks like a good argument; perhaps many who aren't simple internalists will deny (3) but suppose for now Gary accepts it. Suppose he also accepts (1) and (2).

Even then, forming the belief that (4) on the basis of this argument relies on recognising the normative force of modus ponens; that believing P and believing  $P \rightarrow Q$  gives you defeasible reason to believe that Q. Gary does not see the force of this; this is part of what we are supposed to be convincing him of. By his own lights he should not form any beliefs on the basis of modus ponens until you convince him he has reasons to do so.

It is worth noting the close link here with Lewis Carroll's dialogue between Achilles and the Tortoise (Carroll, 1995). In that story, the tortoise asks Achilles to treat him as someone who denies the principle



of modus ponens and to logically compel him to accept the force of a logically valid argument. The tortoise drives Achilles into an infinite regress by continually refusing to accept that Q follows from  $P \rightarrow Q$  and P. Achilles, therefore, fails to convince the tortoise. The tortoise's challenge is very similar to Gary's. Like Achilles, if we try to engage with Gary in this way, we will inevitably fail to convince him.

Thus, when Gary asks us what we could say to him if he didn't care about epistemic standards, we should actually respond, 'there is nothing we could say to convince you then Gary. Such an agent would not be responsive to argument'.

What happens when we respond as the teacher does is we engage with real-life believer Gary who *does* recognise the force of epistemic reasons (else he wouldn't have made it as far as the classroom). This is not what he challenged us to do though. He wanted to know what we could say to someone who was sceptical of epistemic reasons and was looking to be convinced of them.

This demonstrates that there is something wrong with being outside epistemology looking in, deciding whether or not to play the game. Someone outside of epistemology entirely does not make decisions in the way that we do; they cannot work out what reasons they have and weigh them against each other. The discourse with Gary is purported to be the natural way to convince someone to play epistemology, but in fact it would fail against someone who was not already playing.

### **4.1.2 Impact on epistemic institutionalism**

It is a major disanalogy between epistemology and other mere institutions like sports and games that we cannot opt in and out. It is not possible to debate whether it's worth playing epistemology without playing epistemology, whereas it is possible to debate whether it is worth playing rugby without playing rugby.

Imagine someone hesitant about playing rugby who wants to know why they should care about these peculiar rules about passing the ball backwards. We can consider such a person a rugby sceptic who does not care about the rules of rugby as such and is looking for reasons to play.

We can respond as follows:

You might not care about the rules of rugby as such, but you like to have fun, don't you? Rugby's great fun and following the rules (at least mostly) is a means to that end.

This argument's effectiveness requires that the rugby sceptic is not sceptical of modus ponens. But since we may suppose that they are not, there is nothing wrong, by their own lights, with accepting your argument and deciding to play rugby, or equally, deciding that they would rather do something else.

The discussion of the epistemic sceptic suggests that this scenario does not have an epistemic analogue; to put epistemic norms on hold is to become incapable of this kind of rational decision making.

### 4.1.3 What about morality?

So, we have found a difference between epistemology and other mere institutions; it is coherent to put the norms of other mere institutions on hold and debate whether to follow them or not. Such a discussion is not possible for epistemology. One has to respond to epistemic reasons in order to engage in such a debate.<sup>6</sup>

Is this a normatively significant distinction? After all, it is coherent to question the normativity of moral reasons in this way: a sceptic about moral reasons can coherently put those reasons on hold and challenge us to convince them that they have reasons to be moral. This is what the amoralist does. We might try to respond to them like this:

While you may not care about moral standards as such, you care about lots of other things. You don't want to be socially excluded and you don't want to go to prison. Being moral is a means to this end.

There is nothing wrong, by the lights of the amoralist, with accepting this argument – they are not an epistemic sceptic and thus, by their own lights, they can consistently accept the force of a good argument.

However, the problem with such a response in the moral case is that it fails to account for many of an agent's moral reasons. Cases like the Ring of Gyges (see section 3.1) show that this response fails to give the sceptic

---

<sup>6</sup> It may be possible to put individual epistemic norms on hold and debate whether they are worth following, using other parts of the epistemic game. For instance, we can coherently debate whether the law of excluded middle is a good principle. What is incoherent is debating whether the entire enterprise of epistemology is worthwhile while putting it on hold. We can do this for mere institutions – we can debate from outside of rugby whether rugby is in any way worthwhile.

reasons to be moral when they can instead be immoral and avoid social exclusion and legal punishment some other way, contrary to the intuitions of the moral realist. A fully-fledged realist morality can't be justified in this way to someone who denies its normative force.

Both morality and epistemology cannot be reassessed and justified in the way we can reassess the normativity of mere institutions. In the case of morality, non-moral reasons aren't enough to support many of our moral judgements, while in the case of epistemology we cannot even engage in the debate about our reasons if epistemic reasons are put on hold.

#### **4.1.4 Mere commitments**

There is a second thing to note about my comments on the epistemic sceptic. This is that they seem to show that in order to engage in deliberation, everyone must be *committed* to some epistemic norms. Agents needing to be committed to epistemic norms in order to deliberate is not the same as their having genuine reasons to follow epistemic norms – for one thing we don't need to deliberate all the time. One way to respond to this is to say that deliberation is an essential part of the rational process and thus epistemic responsiveness is part of what it is to be subject to reasons at all. I will explore more considerations in this vein in section 4.5.

For now, I will simply say that even if my comments on the epistemic sceptic do not immediately entail categorical, genuinely normative reasons to follow epistemic norms, they still count against epistemic institutionalism. The important thing about merely institutional

normativity for error theorists is that judgements do not necessarily merit inclusion in our all-things-considered assessment of actions. This is how they reconcile the claim that institutional judgements are categorical with reasons-internalism. Whether to follow the rules of a game is the kind of thing that a rational agent can consider whether or not to do, independently of that game. The normativity of epistemology cannot be explained in this way. To successfully explain the normativity of an epistemic judgement it must be that the parties to the explanation are already committed to some epistemic reasons. This is analogous to the way in which to successfully explain the normativity of moral judgements it must be that the parties are already committed to some moral reasons.

This gives us another way to respond to Cowie's argument from ordinary explanation; any explanation of why we should follow epistemic norms relies on a commitment to epistemic reasons. This is in contrast with the normativity of mere institutions, the explanations for which can be carried out completely prior to any commitment to the institution under discussion. Morality also requires parties to be committed to moral reasons for explanations of its normativity to be successful. This gives us reason to think that epistemology is not merely institutional.

#### **4.1.5 Summary**

The so-called natural response to the epistemic sceptic that Cowie thinks demonstrates the close relationship between epistemic norms and merely institutional norms is flawed – by their own lights, the sceptic should not accept it. Your reasons to follow epistemic judgements cannot

be explained without relying on epistemic reasons. This is very unlike the normativity of mere institutions.

Indeed, no other kind of normativity can be explained without a commitment by the parties of the explanation to epistemic norms. Any such explanation requires a willingness to make epistemically licenced inferences between different beliefs about our reasons. What this suggests is that epistemology plays an essential role in the rational process. An essential part of being a reasoner is trying to work out which genuine reasons you have. This cannot be done without playing epistemology. The fact that we cannot engage in the rational process while putting epistemology on hold puts even more space between epistemology and mere institutions.

#### **4.2 Trivial Falsehoods and No Ordinary Explanations**

It is worth noting that I have now provided arguments that directly parallel Cowie's two positive arguments for institutionalism – the argument from trivial truths (1.7.2) and the argument from ordinary explanation (1.7.1). The first we might call Trivial Falsehoods argued for in section 3.3:

Trivial falsehoods, like 'all of the kings of Assyria were called 'Nigel'', are such that we have no moral or prudential reasons to disbelieve them and yet it seems we have genuine reasons to disbelieve. Therefore, unlike merely institutional judgements, there are epistemic judgements such that we have no moral or

prudential reasons to follow but we still have genuinely normative reasons to do so.

Now I have given an argument we might call No ordinary explanations:

Epistemic norms are such that we cannot explain to someone without a commitment to epistemic norms the genuinely normative reasons they have to follow them. This is unlike merely institutional judgements which are such that we can always explain to someone uncommitted to the institution what genuinely normative reasons they have to follow them.

### **4.3 Epistemology as constitutive of belief**

In section 4.1 I argued that epistemology is different from mere institutions in that we cannot step back from it and rationally consider whether the whole institution is worth following – we can only evaluate it from inside epistemology. In this section I will consider a related challenge. This is that in order to count as having beliefs one has to be the subject of epistemic norms. This is sometimes expressed as their being a constitutive normative relationship between belief and truth. Since we have beliefs, we are always subject to epistemic norms. A similar thought forms part of an argument against epistemic error theory:

#### **4.3.1 Normativist arguments against epistemic error theory**

Recall from section 1.6.2 that the Epistemic Existence premise of the argument from analogy is that epistemic error theory is not true. When I introduced this premise, I didn't provide the arguments for it, I simply

said that epistemic error theory would be an uncomfortable extra cost for those who accepted the argument for internalism-based moral error theory. There are, however, quite convincing arguments that epistemic error theory is not just uncomfortable but completely untenable. A class of arguments of this sort are what Cowie calls normativist arguments against epistemic error theory.

These are based on the alleged constitutive normative relationship between belief and truth. The thought is that the state of belief requires normative epistemic properties and relations. For anyone to count as having beliefs, they must be subject to epistemic norms. If epistemic error theory is true and no one is subject to epistemic norms, then no one counts as having beliefs. Since people have beliefs, the argument goes, epistemic error theory is not true.

#### **4.3.2 Ruling out belief**

If epistemic norms apply to all believers and we are all believers, it looks like all of us are subject to epistemic norms all the time, since I argued in section 3.5 that there are no desires that would get us epistemic reasons all the time, this could be taken as evidence that epistemic reasons are genuinely normative and categorical.

However, this depends in what sense epistemic norms must apply to believers. There is plausibly a constitutive relationship between the rules of rugby and being a rugby player that parallels the alleged constitutive relationship between epistemic norms and belief. By playing rugby, you become subject to its rules. To think otherwise is to misunderstand what



it is to play rugby. This does not undermine the merely institutional nature of rugby judgements.

To bring out this point I will borrow a thought from Joyce:

#### **4.4 Joyce and forestalling ‘so what?’ responses**

Joyce, whose presentation of moral error theory we have been working with, provides a number of motivations for reasons-internalism (Joyce, 2001). One is based on the idea that a person’s genuinely normative reasons should be capable of engaging the agent to whom they apply:

Normative reasons claims – claims concerning what it is rational for an agent to do – must be something that potentially engage the agent to whom they are applied. This doesn’t mean that the presentation of a true normative reason claim immediately results in the agent being motivated; rather, it means that the agent cannot sensibly both acknowledge that something is a normative reason for him to act and ask “But so what?” Any adequate theory of normative reasons must make out reasons to be precisely those things that forestall a “So what?” response. Some theories of reasons threaten to violate this constraint – to “alienate” an agent from his reasons.

(Joyce, 2001, pp. 80-81)

Joyce argues that externalist theories of reasons do badly on this count. For instance, he criticises Roderick Firth’s theory that our normative reasons are those things that would motivate an ideal observer who is omniscient, omnipercipient, disinterested, dispassionate, consistent, and

otherwise normal (Firth, 1952). Joyce argues that if an agent asked why they should do something and you told them that it is what an ideal observer would do, they could quite reasonably respond, 'So what? Why should I care what a dispassionate observer would do in my place?'. He thinks that the fact that this response looks reasonable is a point against Firth's theory of reasons; it leaves individuals unacceptably alienated from their normative reasons.

I understand the 'so what?' response to be more than a question of 'why should that be decisive?' or 'why does that provide me with all things considered reason?'. Very few of our reasons are decisive in this way. I take 'so what?' to be a question along the lines of 'what has that got to do with *anything?*' or 'why is that a reason *at all?*'. This interpretation allows us to make sense of why Joyce thinks that individual appeals to desires will do a good job at forestalling 'so what?' when quite clearly individual desires can be outweighed by other competing considerations.

The idea behind the 'so what?' point seems to be a kind of open question argument. It is an open question whether an agent has reasons to do as a dispassionate observer would do in their place and thus Firth's account cannot be all there is to say about normative reasons. Joyce claims that internalist theories do particularly well on this count. However, it is not clear to me that appeals to desire cannot be coherently questioned in the same way that appeals to an ideal observer can. It seems to me that under certain circumstances, an agent can quite coherently ask why their desires make any difference to what they ought to do. Consider a religious devotee who has vowed to forgo all self-interest and devote

their life to the service of others. It seems that such a person could quite coherently respond ‘so what?’ when you pointed out to them that their desires would be best serviced by ignoring the struggles of other people. They can quite coherently question why their selfish desires make any difference to what they should do. We might disagree but it doesn’t seem incoherent for them to ask this kind of question.

Similarly, someone might think that their desires had been so warped by external factors like indoctrination or implicit bias and as such were no longer any reliable guide to optimal actions. In this circumstance they could quite coherently ask ‘so what?’ when it was pointed out that they want women to stay out of the workplace, for example. Joyce endorses a more sophisticated form of internalism than a simple desire account that involves a certain amount of idealisation (see section 3.5.4), however, I am not sure that this helps much in avoiding coherent questioning.

Thus, I disagree with Joyce over what exactly forestalls ‘so what?’ responses. However, I am willing to entertain the thought that forestalling ‘so what?’ is a sign of a genuine reason.

If this is right, then we will see in the next section that epistemic criticism provides us with a plausible case of forestalling ‘so what?’ and thus of genuine epistemic reasons. We will see, however, that the situation is more complex than this:

#### **4.4.1 ‘So what?’ in epistemology**

Consider an agent who is presented with an epistemic reason they are in conflict with. Perhaps we show them evidence against one of their beliefs

or we point to two of their beliefs that contradict one another. There are a number of ways we can imagine them responding:

Firstly, they could deny that what they have been presented with is an epistemic reason, as one does when one denies that something counts as evidence, for instance, when one denies that testimony from a conspiracy theorist counts for anything.

Secondly, they may accept there is a conflict but point to other epistemic reasons which outweigh the reason presented. We do this when we point to overwhelming scientific evidence in favour of our beliefs that outweighs the usually reliable testimony we have just received to the contrary.

Thirdly, they may accept the conflict and point to non-epistemic reasons from other sources which outweigh the evidence. This might happen for example when you are challenged with evidence that your son is unreliable and yet moral reasons to trust him outweigh your epistemic reasons (see section 3.2.1).

Any of these responses are intelligible from someone in the business of belief. Suppose, however, that they respond, 'so what? What has that got to do with anything?'. While you may be able to answer in prudential terms (something about all the desires they'll satisfy if they aim for true beliefs), if someone actually responded in this way, you would be licenced in being confused. This would be someone who recognises that a consideration speaks for the truth of a proposition and does not see the

relevance of that to their belief formation or retention. This would be very strange indeed. It looks like they have misunderstood something.

The fact that there is overwhelming evidence for something seems to be always relevant to an agent's belief about a subject. We saw further evidence for this in section 3.3 in trivial falsehood cases where despite nothing turning on an agent's belief about a subject, it still looked incredibly strange to form outlandish beliefs that conflict with your evidence.

Thus, it seems that epistemic reasons forestall 'so what?' responses when targeted at beliefs; someone who persists with them seems to be exhibiting genuine misunderstanding. If Joyce is right that forestalling 'so what?' is a sign of a genuinely normative reason, this would seem to suggest that epistemic reasons have genuinely normative force.

However, the situation is not so simple. We have said that 'so what?' expresses misunderstanding. However, it need not express a misunderstanding of your genuine reasons but simply a misunderstanding of *belief*. I have said above that some people have argued for a constitutive normative relation between belief and truth. One might think that this explains the strangeness of 'so what?' here – belief is a normatively laden concept and thus an agent who has beliefs misunderstands when they do not see the force of epistemic reasons. However, we can tell a similar story in the case of sports and games in a way that is unproblematic for error theorists.

#### 4.4.2 'So what?' and rugby

Suppose you have decided to play rugby, and someone criticises you for passing the ball forwards. You answer, 'so what?'. The challenger tells you that it is against the rules to pass the ball forwards, but you persist; 'so what?' you say, 'what are the rules of rugby to me?'

The challenger may be understandably confused; your response seems to misunderstand the activity you are engaged in. The rules of rugby are relevant for you precisely *because* you are playing rugby. This looks like a case in which the rules of rugby, which are categorical and disconnected from desires, forestall a 'so what?' response. A rugby player who persists with saying 'so what?' looks confused. However, this does not show that when we talk about rugby we are committed to genuinely normative, categorical reasons.

This seems well explained by the constitutive relationship between the rules of rugby and the activity of rugby. Part of what it is to play rugby is to (mostly) follow the rules. Thus, when you are playing rugby and do not see the relevance of the rules to your decision making, you look confused.

Despite this, with a slight shift of context we sometimes *can* make sense of 'so what?' responses to rugby criticism without diagnosing confusion.

This is because people can question their reasons to play rugby:

The Bystander from section 2.1.2 may provide a case like this. There are two ways of understanding the case. One is as a case of outweighing – you recognise the importance of the rules of rugby for you as a rugby player but see those reasons as outweighed by reasons to help the

bystander. The other is as a case of *trumping*; if someone criticises you for helping the bystander on the grounds that you were playing rugby and you broke the rules, you might legitimately say ‘so what? what significance does the fact I was playing rugby have to whether I should help a bystander in an emergency?’. Responding in this way does not look unreasonable or confused. We might say in this case your rugby reasons are *trumped* by the emergency; they no longer count as genuinely normative due to a moral emergency raising the stakes. What explains this is that the emergency means you lose your reasons to play rugby and thus you lose your reasons to follow the rules.

Another case occurs when you are planning out future actions. If someone points out that you are playing rugby so you better not pass the ball forwards, you cannot simply say ‘so what?’ – this misunderstands what it is to play rugby, however, you can question whether or not to carry on playing rugby or to opt out – to stop playing rugby and play something else instead. If you decide to stop playing rugby (and you are not acting against your reasons in doing so), this makes a ‘so what?’ response intelligible. A non-rugby player can without confusion respond ‘so what?’ when faced with a rugby criticism. We can understand them as denying that they have any reasons to play rugby.

Thus, ‘so what?’ responses to rugby criticism make sense precisely when you lack reasons (or are looking for reasons) to play rugby. I think that what explains your fellow player’s confusion is that on the rugby pitch it is usually taken for granted that you have reasons to play. We ignore special circumstances where someone finds themselves on the rugby

pitch without reasons to play; moral emergencies, general lack of interest in the sport etc. When you make the special circumstances clear – tell them that you have no interest in playing rugby, say, any confusion is resolved.

Therefore, the fact that categorical rugby rules sometimes forestall ‘so what?’ responses to rugby-based criticism is not a problem for the error theorist. This is because this only happens when we take for granted, as we often do on the rugby pitch, your reasons to play rugby. Error theorists can understand your reasons to play rugby in terms of your desires. Furthermore, part of what it is to play rugby is to (mostly) follow the rules. Your genuine reasons to follow the rules of rugby derive from your genuine reasons to play rugby (as was discussed in section 2.4). This is all tolerable to error theorists because all these reasons can be explained in terms of desires.

#### **4.4.3 Returning to epistemology**

Perhaps a similar explanation is possible for epistemology. Perhaps ‘so what?’ only looks strange because a necessary part of believing is a commitment to epistemic norms and we usually take for granted that we have reasons to believe.

Whether this is tenable depends on whether there is an epistemic parallel to trumping and opting out cases, in other words, do ‘so what?’ responses to epistemic criticism make sense in circumstances where you lack reasons to be a believer?



Here we run into the problem from section 4.1. You cannot deliberate without some commitment to epistemic norms thus there is no way to coherently ask ‘so what?’ in response to epistemic criticism. This is to request an argument and when one is questioning whether it’s worth paying attention to arguments. There is nowhere totally outside of epistemology where you can ask this kind of question.

This is not to say, however, that you can’t stop believing. Just because you couldn’t converse with us if you stopped believing does not mean it is unintelligible for us to consider whether it is worth you doing so. There are some spaces outside of belief entirely. These include dreamless sleep, various drug induced states and death. Sometimes it is rational to enter these states. Thus, sometimes it is rational to stop believing.

However, this kind of opting out does not represent a win for the institutionalist. When we opt out of games and sports we can go about our lives as normal. Continuing to engage in reason governed activity. If we opt out of epistemology in these ways, we cease to be the subject of reasons at all. Would it make sense, for example, to criticise someone who was asleep or dead of not maximising their desires? The thought is that this kind of total opting out of belief is to cease to be in the normative sphere. This *is* possible and sometimes rational but no less so for reasons based on desire than it is for epistemic reasons.

#### **4.5 Beliefs, shmeliiefs**

There may, however, be a more sophisticated view on which can opt out of belief by playing a similar game to epistemology with some important

differences. The thought is something along the lines of Enoch's Agency Shmagency paper (Enoch, 2006). Enoch argues that morality will not come from what is constitutive of agency. This is contrary to constructivist views like that of Korgaard (Korsgaard, 1996). Enoch argues that this would not succeed in securing us inescapable moral reasons because one can always question whether you have reasons to be an agent. This is the move we saw when we said that you can question whether you have reasons to play rugby. Enoch thinks this is a problem with any constitutive argument for grounding a kind of reason; one can always question your reasons to be or do the thing that is constituted.

Enoch argues that someone can always ask why they should be an agent rather than a shmagent. Shmagency is a role which has everything in common with agency apart from the commitment to moral norms. Without reasons to be an agent rather than a shmagent, the thought is, the fact that moral norms are constitutive of being an agent does not explain the source of our genuine moral reasons. This suggests the more sophisticated view of what it would take to opt out of epistemology. One might think, without reasons to be a believer we could cease to be believers and be shmeliervers instead. This is just like being a believer; one has just the same mental states except that they do not have to be directed at the truth – perhaps instead they can just be directed towards creating a picture of the world that makes the agent happy. To us, shmeliervers looks like wishful thinking, but to the shmeliervers such states are doing exactly what they are supposed to be doing. Furthermore, shmeliervers is close enough to belief that such an agent can communicate with us

relatively effectively and can be seen as trying to achieve their desires. One might think that a shmeliiever is not outside the normative sphere entirely. We can meaningfully criticise them for failing to achieve their goals. We just can't criticise them for failing to have their shmeliiefs correspond to truth and evidence. This is not what shmeliiefs are for. The Populist Politician from section 3.5.2, for example, might be much better served by shmeliief rather than belief.

The point of this challenge is that it is not enough to say that we are believers and belief requires normative epistemic relations in order to conclude that epistemic norms are inescapable. Given internalism, many agents appear to have reason to be shmeliievers, certainly with regard to some propositions, rather than believers. The desires of the Populist Politician, for example, will be much better served by shmeliiefs instead of beliefs about climate change. Therefore, we can't simply say that epistemology is inescapable for believers unless belief is normatively inescapable in an important sense.

The Populist Politician provides perhaps the best argument for being a believer rather than a shmeliiever. What the intuitions in that case suggest is that sometimes the truth matters for its own sake. On top of doing all the right actions with regards to climate change, it also seems important to understand it and *believe* truly with regards to it. It is important to update your view in response to evidence and avoid wishful thinking. Since we always have these reasons and as I argued in section 3.5, they are not plausibly explained in terms of desires, we always have reasons to be believers and thus to follow epistemic norms.

What this means is that the constitutive strategy must be supplemented by something like my arguments from chapter 3 if it is to ground inescapable epistemic reasons and count against epistemic institutionalism. Institutionalism makes epistemic norms optional. The constitutive strategy argued that epistemology is essential to belief. What I have argued in this chapter is that being a believer is something we can stop doing. What we need, therefore, if epistemology is to be normatively inescapable, is an argument that we should not stop believing. This is what the arguments in chapter 3 supported.

#### **4.5.1 Can you criticise a shmeliiever?**

Besides this, though, we might be able to argue that a shmeliiever is not really a reasoner after all, even though they sometimes look like one. Let me emphasise that this argument is not as strong as the previous simple argument based on our intuitions in favour of the importance of believing well. It is, however, interesting to consider:

The advancement of artificial intelligence gives us an idea of what a shmliiever might look like. We can imagine an AI that has an internal picture of the world that it uses to make its plans but it does not update this picture with the goal of making it reflect the world as it is, it simply updates the picture according to its reward function. Large language models, for instance, do not aim to model the world as it is, they just aim to give plausible responses to prompts. Sometimes this overlaps with the truth (enough for us to engage it in conversation) but often it does not. Sometimes giving a false answer confidently does better according to the reward function than saying it doesn't know something. Such a model

therefore has more of an incentive to have a complete picture of the world than it does to have an accurate one.

It is an interesting question whether a model like could count as reasoning and, therefore, whether it is subject to the same kind of normative criticism that we are. We can, as a matter of fact, assess it according to any test we like; we might test how good it is as a lawyer, for example. However, it is not clear that there is any sense in which it *ought* to be a good lawyer, or indeed that it *ought* to do anything at all. Of course, there is lots of logic going on under the hood, but we can't really isolate any desires of such a thing that might generate reasons for it of the kind that we have. It certainly couldn't give us a reliable account of them. If we asked it, it would just tell us what it thinks we want to hear. A model like this, despite appearances, is something very unlike us and it is not clear that it is right to apply ordinary normative criticism to it.

We might think the same about a shmeliiever. When someone has mental states not targeted towards truth, when they willingly ignore evidence when picturing the world, when they shrug in the face of epistemic criticism, they cease to be intelligible to us as a reasoner. As such, they are no longer the subject of genuinely normative criticism. Therefore, opting out of epistemology in this way is like going to sleep or committing suicide. It is perhaps possible and maybe even reasonable, but once it has been done, you thereby escape all kinds of normative criticism. This is a far cry from the normativity of mere institutions.

## 4.6 Summary

In this chapter I have considered two arguments that suggest epistemology is not normative in the same way as mere institutions. The first is summarised by the argument from No ordinary explanations:

Epistemic norms are such that we cannot explain to someone without a commitment to epistemic norms the genuinely normative reasons they have to follow them. This is unlike merely institutional judgements which are such that we can always explain to someone uncommitted to the institution what genuinely normative reasons they have to follow them.

The second is based on the idea that epistemic norms are constitutive of belief. People have argued that you cannot be a believer without being subject to epistemic criticism. I have argued that this alone does not provide a point of distinction with mere institutions because sometimes it looks like you can't play rugby without being subject to rugby criticism. The apparent inescapability comes from a constitutive relationship between the relevant norms and a certain activity. Error theorists can tolerate this for rugby because you can opt out of playing rugby. Having reasons to follow the rules of rugby ultimately derives from having reasons to play rugby and the error theorist can understand these reasons in terms of desires.

I explored a similar strategy for epistemology, where epistemic norms are constitutive of belief but whether you have reasons to follow them is dependent on your reasons to be a believer. Even though I argued in

section 4.1 that one cannot step back and reconsider your commitment to epistemology in the way you can reconsider your commitment to games and sports, there are certain limited ways you can still opt out of epistemology. Therefore, it is in principle possible to lack reasons to be a believer. However, I have argued in section 4.5 that, chiefly due to my arguments from chapter 3, you always have some genuine reasons to be a believer that cannot be explained in terms of desires and therefore epistemology is not optional like other mere institutions. The point is though that the constitutive strategy alone does not provide evidence against institutionalism – it needed to be supplemented by other arguments.

I also briefly considered in section 4.5.1 the possibility that in any circumstance when you can be said to have opted out of epistemology, you have opted out of reasoning altogether and therefore ceased to be the subject of any normative criticism. This was quite speculative but if it is right then it leaves epistemology a long way away from the mere institutions that the error theorist was comfortable with.

The arguments of this chapter are more technical than those of chapter 3 and rely on less firmly held intuitions, but if they are right, they show that epistemology is very far from the model of other mere institutions, particularly sports and games. This is another strike against epistemic institutionalism.

## 5 Conclusion

I have argued that Cowie is unsuccessful in his argument that epistemic judgements, unlike moral judgements, are normative only in the sense that judgements within etiquette, fashion, sports and games, and the law are normative. I have objected to this claim on two main grounds. The first is that it gives unintuitive answers about our genuine reasons in epistemic Ring of Gyges cases. The second is that we cannot reconsider our commitment to epistemology in the way we can for sports and games.

In chapter 2 I set out Cowie's distinction between merely institutional and genuinely normative judgements and argued that if the error theorist is to be successful in explaining epistemic judgements via epistemic institutionalism, they need to give us a story about how and when we get genuinely normative reasons to follow epistemic norms. I presented two ways that we can acquire genuine reasons to conform to mere institutions in the case of the law and of sports and games and noted some of their limitations. I emphasised that we lose these genuine reasons when we do not have the right desires.

In chapter 3 I attempted to provide an epistemic analogue to Plato's Ring of Gyges. I took this to be the task of finding a case where we have no moral or prudential reasons to follow epistemic norms and yet our intuitions still favour compliance. I introduced two possible cases. The first concerned trivial but ridiculous falsehoods which we nevertheless seem to have reason not to believe. The second concerned facts about



climate change for which there is overwhelming scientific evidence but strong prudential reason to avoid believing. After exploring the possibility that epistemic institutionalism can explain these intuitions via necessary desires, I concluded that the error theorist will have to bite the bullet and accept there are no genuine reasons to believe according to epistemic norms in these cases. Thus, epistemic Ring of Gyges cases represent a genuine extra cost for error theorists. The fact that institutionalism cannot explain these intuitions is a point against Cowie's theory.

In chapter 4 I considered the discourse with the epistemic sceptic. I noted that an epistemic sceptic is in a certain sense unintelligible. We couldn't converse with them and they couldn't reason in the way that we do. I argued that this presented a problem for certain understandings of epistemology as merely institutional. I argued that the unintelligibility of the epistemic sceptic means that you cannot step back and reconsider your commitment to epistemology in the way you can reconsider your commitment to sports and games. One can only assess your commitment from inside epistemology.

I also considered what we might call a constitutive strategy for arguing that we always have genuine epistemic reasons; we have beliefs and we couldn't believe without epistemic reasons. I argued that this argument runs into a problem that there are ways to stop believing. I argued that one cannot get genuine reasons out of a role when you do not have reason to occupy that role. We saw this in the case of sports and games and the same thing applies to belief. I argued that, in light of chapter 3, it

appears we do have reasons to believe that cannot be explained by our desires. I briefly considered the possibility that, in addition, it may be unintelligible to consider normative reasons of any kind for someone who refuses to believe. This suggests that epistemic reasons are on just as firm ground as any reasons rooted in desire. They can be escaped in just the same ways.

It is worth noting that my arguments have broader application than epistemic institutionalism. Cowie spends the latter part of his book considering what happens if his main argument is unsuccessful. He argues that the Parity Premise may even be true and the argument from analogy fails, nonetheless. This is because he thinks that despite the arguments that it is untenable, epistemic error theory may yet be true. He argues for the possibility of what he calls the surrogate strategy:

The surrogate strategy: The parity premise is true. But the epistemic error theory could be true too. Its worst effects are offset by surrogates for epistemic judgements, which are modelled on normative judgements in sports and games, etiquette, and the law.

(Cowie, p. 181)

The idea is that by rebuilding epistemology in the image of mere institutions we could salvage much that was important about epistemology.

Even if Cowie succeeds in mitigating the arguments against epistemic error theory via the surrogate strategy, my arguments will apply to this

rebuilt epistemology too: proponents of the surrogate strategy still have to accept that we have no reasons to follow epistemic norms in epistemic Ring of Gyges cases. This is still an intuitive cost. They also run into the conceptual problems of treating epistemology like a game discussed in section 4.1. We would still be unable to deliberate without epistemology. Perhaps the most substantial cost that I have highlighted is that the truth never matters in itself if the surrogate strategy is adopted. We would never have reason to believe something simply because it is true. The Populist Politician who we first met in section 3.5.2 turns out rational in disbelieving in man-made climate change.

Throughout I have been working with a particular argument for moral error theory. I have argued that epistemic institutionalism fails to protect epistemology from this argument and thus that argument over-generates – the internalism-based argument from analogy succeeds. This doesn't yet consign moral error theory to scrapheap of unacceptable scepticism. There are various other fallbacks that Cowie explores that there has not been space to go into here. Also there may be other quite different arguments for moral error theory that do not apply to epistemology. As I touched on in section 3.2, there are a number of differences between epistemology and morality even if we deny epistemic institutionalism. Any one of them might provide a way in for a moral error theory that does not objectionably spread to epistemology.

However, if this version of the argument from analogy goes through it will be a significant blow to moral error theory. As I mentioned in the introduction, the internalism-based moral error theory has roots in

Mackie and is a particularly influential form of the argument. If even it identifies a feature of normative discourse that spreads far beyond morality it will be a significant result. Furthermore, if my arguments against epistemic institutionalism succeed it will remove a valuable tool from the error theorist's arsenal in responding to arguments from analogy. Cowie thinks that institutionalism will help to avoid not only the internalism-based argument from analogy but the irreducibility-based argument too. I have not discussed that argument here but if institutionalism is false then it cannot be used in that case either.

Arguments from analogy are a very valuable tool against the moral error theorist and against scepticism more generally. What I have said here hopefully adds to argument that the scepticism of moral error theorists cannot be contained. Let this act as a case study for a broader argument that moral error theory should be treated as a spectre – an unacceptable conclusion. Besides this, what I have said about epistemology is important in its own right. In a post-truth world, the value of truth is at risk of getting lost. It is important therefore to see epistemology as more than an instrument for achieving whatever goal the user wishes to put it to but as genuine directives for organising our view of the world. It matters when you deny the evidence of man-made climate change, no matter what you want. Moral error theory and epistemic institutionalism deny that. We should deny them.

## 6 Bibliography

- Carroll, L. (1995). What the Tortoise said to Achilles. *Mind* 104(416), 691-693.
- Cowie, C. (2019). *Morality and Epistemic Judgement: The Argument from Analogy*. Oxford: Oxford University Press.
- Cuneo, T. (2007). *The Normative Web: An Argument for Moral Realism*. Oxford: Oxford University Press.
- Das, R. (2016). Why Companions in Guilt Arguments Still Work: Reply to Cowie. *The Philosophical Quarterly* 66(262), 152-60.
- Das, R. (2017). Bad News for Moral Error Theorists: There Is No Master Argument Against Companions in Guilt Strategies. *Australian Journal of Philosophy* 95(1), 58-69.
- Enoch, D. (2006). Agency Shmagency: Why Normativity Won't Come from What Is Constitutive of Action. *The Philosophical review* 115(2), 169-198.
- Firth, R. (1952). Ethical Absolutism and the Ideal Observer. *Philosophy and Phenomenological Research* 12, 317-45.
- Foot, P. (1972). Morality as a System of Hypothetical Imperatives. *The Philosophical review* 81.3, 305-316.
- Harman, G. (1977). *The Nature of Morality: An Introduction to Ethics*. Oxford: Oxford University Press.

- Joyce, R. (2001). *The Myth of Morality*. Cambridge: Cambridge University Press.
- Korsgaard, C. (1996). *The Sources of Normativity*. Cambridge: Cambridge University Press.
- Leite, A. (2007). Epistemic Instrumentalism and Reasons for Belief: A Reply to Tom Kelly's 'Epistemic Rationalist as Instrumental Rationality: A Critique'. *Philosophy and Phenomenological Research* 75(2), 456-64.
- Mackie, J. (1990). *Ethics: Inventing Right and Wrong*. Harmondsworth: Penguin.
- McPherson, T. (2011). Against Quietist Normative Realism. *Philosophical Studies* 154(2), 223-40.
- Olson, J. (2014). *Moral Error Theory: History, Critique, Defence*. Oxford: Oxford University Press.
- Papineau, D. (2013). There Are No Norms on Belief. In T. C. (ed.), *The Aim of Belief*. Oxford: Oxford University Press.
- Parfit, D. (2011). *On What Matters Volume 2*. Oxford: Oxford University Press.
- Plato. (1974). *Plato's Republic / Translated by G.M.A. Grube*. Indianapolis: Hackett.
- Railton, P. (1997). On the Hypothetical and Non-Hypothetical in Reasoning about Belief and Action. In G. C. (eds.), *Ethics and Practical Reason*. Oxford: Oxford University Press.

Rowland, R. (2013). Moral Error Theory and the Argument from Epistemic Reasons. *Journal of Ethics and Social Philosophy* 7(1), 1-24.

Smith, M. (1994). *The Moral Problem*. Oxford: Blackwell.

Williams, B. (1981). Internal and External Reasons. In B. Williams, *Moral Luck: Philosophical Papers: 1973-1980* (pp. 101-113). Cambridge: Cambridge University Press.