

How and When to be a Constitutivist: Why Constitutivism is a better fit for Epistemic Norms than for Moral Norms

Louise Hanson

Introduction

Constitutivism has been widely discussed in both metaethics and epistemology.¹ In this paper I argue that it has been misunderstood in a number of key ways. I argue that a hallmark of constitutive norms is a particular kind of escapability - what I will call *Jurisdictional Escapability*. I argue that jurisdictional escapability is the key to whether constitutivism is a good fit or a potentially awkward one, in a given domain.

Jurisdictional escapability is often run together with a different kind of escapability - which I will call *Normative Escapability*. But clearly distinguishing these, has significant upshots both for the normative landscape generally, and for constitutivism. First, it reveals that constitutivism is a particularly good fit for epistemic norms, and not such a good fit for moral ones. Second, it reveals that constitutivism is not, as is often thought, an antirealist position, and is in fact entirely compatible with robust realism. Third, it introduces a new dimension of precision into the normative landscape.

In §1 I present a new argument for epistemic constitutivism, that hinges on jurisdictional escapability. In §2 I show how the normative/jurisdictional

¹ See, for example, Korsgaard 1996: *The Sources of Normativity* (Cambridge: Cambridge University Press). Korsgaard 2009: *Self-Constitution: Agency, Identity, and Integrity* (Oxford: Oxford University Press), Velleman, 1989: *Practical Reflection* (Princeton, NJ: Princeton University Press). Velleman 1996: "The Possibility of Practical Reason," *Ethics* (106) pp. 694–726, Katsafanas, Paul 2013: *Agency and the Foundations of Ethics: Nietzschean Constitutivism* (Oxford: Oxford University Press), Katsafanas 2018 "Constitutivism about Practical Reasons" in Daniel Star (ed.) *Oxford Handbook of Reasons and Rationality* (Oxford: Oxford University Press).

escapability distinction illuminates and precisifies the normative landscape. And that it reveals that constitutivism is compatible with robust realism. In §3 I argue that the normative/jurisdictional escapability distinction reveals that an influential objection to constitutivism - David Enoch's shmagency objection - misses the point. In the epistemic case it misses its target entirely, while in the moral case it is significantly deflated.

§1. The asymmetry argument for Epistemic Constitutivism

1.1 What is constitutivism?

Some activities are constitutively governed by certain norms, in the sense that it's *part of what it is* to be doing that activity, that certain norms apply to you. It is part of what it is to be playing chess not only that certain descriptive things are true of you (you are moving objects on a chequered board), but also that certain *normative* things are true of you (you should not miss an opportunity to checkmate your opponent, you should avoid getting into a position where your opponent can easily checkmate you). It is part of what it is to be building a house not only that certain descriptive things are true of you (you are creating some kind of structure), but also that certain *normative* things are true of you (you should be trying to make that structure stable).

Constitutivism about the norms in a given domain is the view that those norms are constitutive norms of a given activity or practice that people are engaging in. That they apply to people in virtue of their engaging in that activity or practice. For example, constitutivism about moral norms is the view that moral norms are constitutive norms of some activity that people are doing - many constitutivists take the relevant activity to be the activity of *being an agent*.² According to them, just as there are norms that apply to you in virtue of the fact you are playing

² Perhaps it's a bit of a stretch to think of being an agent as an *activity* in the usual sense of the word, but if that's the case, then I take it that the point is just that agency is *like* activities in certain relevant respects - namely in that it is something that we do, and that is governed by certain constitutive norms, and further that it is *because we do it* that we are bound by those constitutive norms.

chess or building a house, there are also norms that apply to you in virtue of the fact you are *being an agent*. And, according to them, *this is what moral norms are*.

It is widely accepted that morality isn't the only domain in which there are norms, and that another such domain is epistemology. Epistemic norms might include *shoulds* and *shouldn'ts* like the following:

- (i) Suppose you're wondering whether P, and then you get overwhelming evidence that P. You should believe P
- (ii) Suppose you believe all the premises of a valid argument whose conclusion is Q, and you are considering the argument and know it to be valid. You should believe Q
- (iii) You shouldn't believe that P on the basis of wishful thinking
- (iv) You shouldn't believe both P and not-P at the same time
- (v) Suppose you get overwhelming evidence that not-P. You should not believe that P

Epistemic constitutivists hold that these *shoulds* and *shouldn'ts* are constitutive norms of, for example, the activity of inquiry, belief-formation, belief-maintenance, etc.

In what follows I present a new argument for epistemic constitutivism. I then draw some deeper lessons about constitutivism and about the normative landscape more generally.

There are different types of constitutive norms, and philosophers who want to be constitutivists about a given normative domain are free to choose which types of constitutive norms they will take as their model. It's popular to use the example of chess to illustrate what constitutive norms are like and how they behave. But even here there are two types of constitutive norm that one could take as one's template for the norms in a given domain:

Norms that if you didn't follow them you wouldn't count as playing chess *at all*

Norms that if you didn't follow them you wouldn't be playing chess *well*

Both are constitutive norms of chess in the sense that part of what it is to be playing chess is that these norms *apply to you*. In both cases, if the norms didn't apply to you, you wouldn't be playing chess. In the first case, however, you'd also not count as playing chess if *you didn't follow them*, while in the second case, you could still count as playing chess - you just wouldn't count as playing chess *well*.

The type of constitutivism that I think is a good fit for epistemic norms is a type according to which epistemic norms are norms of some epistemic activity, that tell us how to do that activity well. In other words, norms that, if you don't follow them, you're doing the activity badly.³

The argument I will make is the following:

The Asymmetry Argument for Epistemic Constitutivism

1. There is a particular kind of asymmetry between positive and negative epistemic norms (Premise)
 2. The best explanation for this asymmetry is that epistemic norms are constitutive norms (Premise)
-
3. Epistemic norms are constitutive norms (by Inference to the best explanation)

In the rest of this section I make the argument in more detail. 1.2 presents my case for Premise 1, in 1.3-1.4 I argue for Premise 2.

³ Others who pursue a similar strategy include Horst (2022) 'In Defense of Constitutivism About Epistemic Normativity' *Pacific Philosophical Quarterly*, footnote 5, and (2018) 'Constitutivism about practical reasons' in Daniel Star (ed.), *The Oxford Handbook of Reasons and Normativity*. New York, NY, United States of America: Oxford University Press. pp. 367-394, especially pp. 2-4.

1.2 First Premise: there is an asymmetry between positive and negative epistemic duties

Consider examples (i) - (v) again. (i) and (ii) are what are sometimes called *positive epistemic duties*. They are duties to form a certain belief. (iii)-(v) are what are sometimes called *negative epistemic duties*. They are duties to refrain from forming a particular belief or set of beliefs.

I will argue that there is a striking asymmetry between positive and negative epistemic duties. The positive ones look like they are in some sense escapable - it's possible to opt out of them, while the negative ones are not.

Why think that positive epistemic duties are escapable? Consider this case from Adam Leite:

TRIVIAL TRUTHS

I'm standing next to a door at a convention center. 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. Perhaps I'm busy thinking about things of greater interest or importance to me. 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.⁴

Rather, it looks as though, if you're not at all interested in whether P, then you have no obligation/duty to form the belief that P, even in the face of very good evidence that P.

⁴ Leite 2007 'Epistemic Instrumentalism and Reasons for Belief' *Philosophy and Phenomenological Research*, Vol. 75, No. 2, pp. 456-464, p. 458

Negative epistemic duties, on the other hand, don't seem to behave this way. There is no opting out of the requirement to not believe a contradiction, to not reason invalidly, and so on.

It will be helpful to introduce a new term to capture this asymmetry:

Counterinstantiating: you count as counter-instantiating a particular duty/norm if and only if you're doing the opposite of what it says.

Counterinstantiating a norm covers two kinds of case:

- doing the thing it says not to do, or
- failing to do the thing it says to do

So counterinstantiating is something that is possible with both positive and negative epistemic norms. To counterinstantiate a positive epistemic norm, would be to fail to do the thing it says to do, and to counterinstantiate a negative epistemic norm, would be to do the thing it says not to do.

We're now in a position to state the kind of asymmetry that epistemic norms appear to exhibit:

Positive epistemic norms, but not negative ones, can be *faultlessly counter-instantiated*. You can counterinstantiate a positive epistemic norm without counting as doing anything wrong, whereas you can't counterinstantiate a negative epistemic norm without counting as doing something wrong

For example, the person in the TRIVIAL TRUTHS case counterinstantiates the epistemic norm 'believe that P when you have excellent evidence that P', but they are not doing anything wrong in counterinstantiating it. In contrast, it doesn't look possible to faultlessly counterinstantiate a negative epistemic norm, such as 'don't believe a contradiction'. Counterinstantiating that always looks like a case of going wrong somehow.

So, positive epistemic norms admit of faultless counterinstantiation, while negative ones do not.

This concludes my case for Premise 1. The next two sections present the argument for accepting Premise 2.

1.3. Second Premise: The best explanation for the asymmetry is that epistemic shoulds are constitutive shoulds

What could explain the asymmetry? In this section, I will talk through three candidate explanations people have offered:

- (1) Epistemic norms lack normative force
- (2) There are no positive epistemic duties, only negative ones
- (3) Epistemic duties are restricted to propositions the agent is considering/are important/etc.

I will argue that none of these offers a satisfactory explanation. Then in 1.4 I will show how constitutivism about epistemic norms, by contrast, would offer a highly satisfactory explanation.

(1) Candidate Explanation 1: Epistemic norms lack normative force⁵

What does this mean?

Some norms do not have normative force in their own right. As Philippa Foot puts in, in the case of norms of etiquette:

[...] one may reasonably ask why anyone should bother about what [...]

⁵ See Cowie, C. (2019) *Morality and Epistemic Judgement: The Argument from Analogy*, OUP

should-from-the-point-of-view-of-etiquette be done, and that such considerations deserve no notice unless reason is shown. So although people give as their reason for doing something the fact that it is required by etiquette, we do not take this consideration as in itself giving us reason to act. Considerations of etiquette **do not have any automatic reason-giving force.**⁶

And as others have noted, there are many systems of norms that are like this:

Plausibly, there is **not automatically a good reason to conform to** norms of fashion, gender stereotypes, tradition, religious sects, the Sicilian Mafia code of conduct, and so on. The point is that most norms **do not seem to necessarily entail or provide good reasons to do as they say.**⁷

Norms that lack normative force would be faultlessly counterinstantiable. Take norms of etiquette as an example of norms that lack normative force. If that's right, you can counterstantiate them without doing anything wrong. Doing something wrong *by the norms of etiquette* would not entail doing anything wrong in any meaningful sense.

So one tempting suggestion for explaining the faultless counterinstantiability we've seen with epistemic norms, is that these are like the norms of etiquette, fashion, and the like, in that they lack normative force.

However, this suggestion fails to explain what needs to be explained, in two ways. First, it arguably doesn't capture the particular kind of faultless counterinstantiation in play here. Let's look a bit more closely at the kind of faultless counterinstantiation you get with norms that lack normative force.

Take norms of etiquette. And suppose that the following is a norm of etiquette:

⁶ Foot 1972, 'Morality as a System of Hypothetical Imperatives', p. 309

⁷ Cote-Bouchard 2016, 'Can the Aim of Belief Ground Epistemic Normativity?' *Philosophical Studies* p. 3183

NO FLIP FLOPS: you should not wear flip flops to a formal dinner

Why is it that you can counterinstantiate NO FLIP FLOPS and not be doing anything wrong? It's because etiquette norms don't really matter. Are you doing something wrong according to etiquette norms if you counterinstantiate NO FLIP FLOPS? Yes. Are you doing something wrong in any meaningful sense? No, because who cares about etiquette?

Compare the NO FLIP FLOPS situation with what's going on in TRIVIAL TRUTHS. In that case there is an epistemic norm that is being faultlessly counterinstantiated:

ES1: you should form the belief that the next person will have a dog

If the reason you can faultlessly counterinstantiate ES1 was that ES1 lacked normative force, then what the details would be analogous to the NO FLIP FLOPS case. It would be that you are doing something wrong epistemically, but that that doesn't matter.

That doesn't look right, however. It doesn't look like you are doing anything wrong here *even epistemically*. To see this, note how odd it would be to say someone is a bad epistemic agent, and to cite in support of this all the times that there was some trivial truth they were uninterested in, that there was lots of evidence in support of but that they didn't form the belief. In contrast, it wouldn't be odd at all to say that someone is bad at etiquette by citing all the times they wore flip flops to formal events. (At least on the assumption that NO FLIP FLOPS is a norm of etiquette).

This suggests that the location of the faultlessness is different in the two cases. In the etiquette case, you're doing something wrong according to the norm in question. But you're faultless because that norm doesn't really matter. But things are quite different in the epistemic case, It's not: you are doing something wrong *epistemically* but that doesn't matter because epistemic norms lack normative force. But rather: you are not even doing anything wrong epistemically.

I'll have more to say on this point in 2.1, but for now, here is a second reason why the normative force explanation won't work: it doesn't explain the asymmetry between positive and negative epistemic norms with respect to counterinstantiating. In fact, it predicts that there would be no asymmetry here.

If epistemic norms lacked normative force, you would expect that both positive and negative ones could be faultlessly counterinstantiated. And yet it's only the positive ones that can.

This would be enough to show that the normative force explanation is not compelling. So the normative force explanation should be rejected.

(2) Candidate Explanation 2: There are no positive epistemic duties, only negative ones⁸

This suggestion also doesn't succeed in explaining the asymmetry between positive and negative epistemic duties. It merely notes that there is an asymmetry, but doesn't explain it. If there are no positive epistemic duties, but there are negative ones, why is this?

More importantly, perhaps, the claim that there are no positive epistemic duties does not look plausible. Consider the following:

TRIVIAL TRUTHS 2

You're really curious about whether the next person to walk by will have a dog, and you're confronted by the evidence just like in TRIVIAL TRUTHS 1. And yet you don't form the belief.

In TRIVIAL TRUTHS 2 it looks like you should form the belief. That you're doing something wrong in not forming the belief - just as wrong as believing a contradiction. reasoning invalidly etc.

⁸ See Nelson, M. (2010) 'There are no positive epistemic duties' *Mind* 119(473):83-102; Whiting, Steiglich-Petersen.

So the *no positive epistemic duties* explanation should be rejected on the grounds that it's both unexplanatory and implausible.

(3) Candidate Explanation 3: Epistemic duties are restricted to propositions the agent is considering/are important/etc.⁹

This third suggestion looks more promising. The restrictivist approach covers a range of ways in which the details could be filled in, of what the relevant restriction is, and some ways of filling in the details are going to work better here than others. The version where the restriction is to *propositions the agent is considering* would deliver the right results on both TRIVIAL TRUTHS cases. It correctly predicts that the counterinstantiation in TRIVIAL TRUTHS 1 is faultless, while the counterinstantiation in TRIVIAL TRUTHS 2 is not. However, it is not yet an *explanation* of the asymmetry between positive and negative epistemic duties. It doesn't - as it stands - explain why positive epistemic duties admit of faultless counterinstantiation while negative ones do not.

In the next section, I will argue that constitutivism accommodates what's right about restrictivism, but also goes further in that it mobilises it into a natural explanation of the asymmetry between positive and negative epistemic duties.

1.4. A better explanation - Epistemic Norms are Constitutive Norms

I've argued that there is an asymmetry between positive and negative epistemic duties in that the former can be faultlessly counterinstantiated while the latter cannot. This asymmetry stands in need of explanation. In the last section, I argued that three existing options for explaining the asymmetry fail. Here I will argue that constitutivism would offer a natural explanation of the asymmetry. The first step in arguing this is to show that constitutive norms exhibit the same

⁹ Wedgwood, Chisholm, Goldman, Alston

kind of asymmetry. They too are such that the positive ones, but not the negative, can be faultlessly counterinstantiated.

1.4.1 Constitutive Norms Exhibit the Same kind of Asymmetry

Consider chess norms. We need an example of a positive chess norm, and a negative one. So let's take the following:

Positive: you should move your queen to E5

Negative: you should not move your knight

Consider the positive one first, and consider the following case:

PHONE CALL

Suppose it's true of you that you should move your queen to E5. But then suppose you get a phone call that you have to take, and then you never get round to returning to the game, so you never get round to moving your queen.

What's the natural thing to say about your relation to this chess duty here? You wouldn't count as complying with it. But it doesn't look like you would count as transgressing it either, because you left the game. That chess duty only looks like it applies to you when you're in the game. And moreover, only when you're in *that particular* game.

Even if it were true of you when you were playing, that you should move your queen to E5, you haven't done anything wrong by chess standards. After all, the chess *shoulds* don't adjudicate on whether you should be playing or not. Or whether you should continue the game. They just tell you what to do in a particular game when you're in that game.

Positive chess duties appear to be faultlessly counterinstantiable.

Now consider the negative chess shoulds. Suppose, for example, it's true of you that you should not move your knight. Unlike the duty to move your queen, you can't counterinstantiate it by exiting the game. If you exit the game, you aren't moving your knight. So you are not counterinstantiating that chess duty.

1.4.2 What explains this asymmetry?

Chess norms, and more generally, norms that are constitutive norms of some activity, are such that you can step outside of their jurisdiction, by exiting the activity.

This means you have three options, for any norm that's a constitutive norm of some activity:

Comply (would involve staying doing the activity and doing as it says)

Transgress (would involve staying in the activity and counterinstantiating - i.e. doing the opposite of what it says)

Sidestep (would involve exiting the activity, and thereby stepping outside of their jurisdiction - so you are neither complying with them nor transgressing them).

This possibility of sidestepping is a direct upshot of the fact that the norms in question are norms *governing some activity*, and that they have nothing to say to people not doing that activity. This results in the following asymmetry between positive and negative constitutive norms:

Positive constitutive norms

You can counterinstantiate without transgressing. If you exit the activity, and don't do the thing it says, you're not transgressing it: you're sidestepping it. So you're counterinstantiating, but the counterinstantiation

is faultless, in the sense that *even by the lights of the norms in question*, you're not doing anything wrong in counterinstantiating them.

Negative constitutive norms

You cannot counterstantiate without transgressing. You can exit the activity, and that would take you out of the jurisdiction of that norm. But if you are counterinstantiating the norm. You're doing the thing it says not to do. And if you're doing the thing it says not to do, you're in the activity. So it's not possible to counterstantiate and sidestep. All cases of counterinstantiating a negative constitutive norm are cases of transgressing it.

So we've seen that constitutive norms exhibit the same asymmetry that epistemic norms exhibit. The positive ones admit of faultless counterinstantiation, while the negative ones do not. We've also seen *why it is* that constitutive norms exhibit this asymmetry. It's because they are norms governing a particular activity, and so their jurisdiction is limited to people who are doing the activity. Faultless counterinstantiation is possible precisely because *stopping doing the activity* is possible, and takes you out of the jurisdiction of the relevant should. This in turn results in an asymmetry between positive and negative constitutive norms, because the negative ones are such that counterinstantiating them places you within the activity, while the positive ones are not. For the negative ones, there is some move in the activity they are telling you not to make.

So counterinstantiating always involves making that move. And if you're making that move, you're in the activity. For the positive ones, there is some move in the activity they are telling you *to* make. So counterinstantiating is just a matter of *not* making that move. This could be done in two ways - either by being in the activity and not making that move, *or* by not being in the activity. The latter is faultless counterinstantiation.

Given this, epistemic constitutivism offers a compelling explanation of the asymmetry between positive and negative epistemic norms. According to epistemic constitutivism, epistemic norms are constitutive norms of some activity. And that activity is one that it is possible for us to be doing or not be

doing. The cases of faultlessly counterinstantiating a positive epistemic should, are cases where the person is not in the activity that that particular epistemic should governs. And so is outside of its jurisdiction.

Consider TRIVIAL TRUTHS 1 again, and the epistemic norm that is being faultlessly counterinstantiated there.

ES1: you should form the belief that the next person will have a dog

According to constitutivism, ES1 is a norm governing a particular activity. So it's possible to faultlessly counterinstantiate it, just in case it's possible to *not be doing that activity*.

What would the activity be? Here's one answer that would not be plausible. If the activity is *belief-formation generally*, that's not going to work, because presumably in TRIVIAL TRUTHS 1, the bar you have to meet to be able to be able to faultlessly counterinstantiate ES1 is not

to be not forming any beliefs about anything whatsoever.

Intuitively, in TRIVIAL TRUTHS 1, you can faultlessly counterinstantiate ES1 much more easily than that. You can faultlessly counterinstantiate ES1 simply by *not* considering whether the next person will have a dog or not.

This suggests that the activity in question would need to be the activity of

considering whether the next person will have a dog or not

How does this generalise? It looks like if faultless counterinstantiation is always a case of *not doing some activity*. The activity has got to be one of considering the relevant subject matter, or considering whether P, for some proposition P, rather than the activity of forming beliefs generally. And this seems to be mirrored in the chess case too. What gets you out of having to move your queen, is

leaving that particular game of chess

not

refraining from playing any games of chess whatsoever

Suppose someone is playing a few different games with a few different opponents. And suppose that in one game they should move their queen. We've seen that they can faultlessly counterinstantiate that by leaving the game, and leaving the game is *all they need to do*. They don't need to leave any of the other games of chess they are playing.

On this kind of picture, the epistemic norms that apply to you, apply to you in virtue of the fact that you are considering particular propositions or subject matters, and you can sidestep those norms by not considering those propositions or subject matters. This brings the possibility of faultless counterinstantiation for positive epistemic norms, because as long as you aren't considering whether P, the norms governing the activity of considering whether P, don't apply to you.

For example, you're not considering whether the next person will have a dog, so the norm ES1 doesn't apply to you. Any more than the norm

'since you'll be checkmated if you don't move your queen, you should move your queen'

applies to you when you're not in that particular game of chess in which that situation arose.

But for *negative* epistemic norms, such as 'don't believe that P and not-P', there is no possibility of faultless counterinstantiation. To counterinstantiate would be to form the belief that 'P and not-P', and if you're forming that belief, you can't be outside of the activity of considering whether P.

This possibility of sidestepping, and the related possibility of faultless counterinstantiation, is a kind of escapability. Norms that have this feature are in some sense escapable., and this is because can step outside of their jurisdiction,

and no longer be governed by them. Call this feature *Jurisdictional Escapability*.

If what I've argued here is correct, jurisdictional escapability is something that epistemic norms have, and this gives us good reason to think they are a kind of constitutive norm, and that attending to the particular details of this jurisdictional escapability in the epistemic case, suggests that they are norms governing the activity of considering whether P, for some proposition P.

In the next section, I highlight two broader upshots of this discussion, one about normativity generally, and one about constitutivism.

§2. Broader upshots

2.1 Upshot 1: Two distinctions in normativity

I've argued that what makes it possible to faultlessly counterinstantiate constitutive norms, is that they have *jurisdictional escapability*. Jurisdictional escapability is importantly different from another, more widely discussed, kind of escapability that norms can have. In 1.3 I discussed the escapability had by norms that lack normative force. I'll call this *Normative Escapability*. Both normatively escapable and jurisdictionally escapable norms are possible to faultlessly counterinstantiate, but the details are importantly different in the two kinds of case.

Take normative escapability first. I argued in 1.3 that if you faultlessly counterinstantiate a norm that's normatively escapable, it's not that the norm doesn't apply to you, or that you are not really transgressing it. It does and you are. The reason you are faultless in doing so is that it doesn't really matter - the norm lacks normative force.

So with the case of counterinstantiating the etiquette norm NO FLIP FLOPS, what would be going on there would be:

- You'd be genuinely transgressing an etiquette norm. You'd be doing something wrong, *by the norms of etiquette*.
- But the reason it's faultless is that the norms of etiquette don't really matter in the right kind of way - they don't have any normative force

With jurisdictional escapability, the details are slightly different. If you faultlessly counterinstantiate a norm that's jurisdictionally escapable, the reason counterinstantiating is faultless, is that you've *stepped outside of its jurisdiction* and so the norm doesn't apply to you. So in contrast with the case of normatively escapable norms, you aren't transgressing the norm in this case; you are sidestepping it. So the situation is not that you're doing something wrong *according to that norm* but it doesn't matter. Rather you aren't doing anything wrong, even according to that norm.

So here's the test for whether a norm N is jurisdictionally escapable or normatively escapable:

Is it possible to faultlessly counterinstantiate N?

If no, then N is neither jurisdictionally escapable nor normatively escapable.

If yes, then:


is it the case that when you faultlessly counterinstantiate it, you are doing something wrong according to N?

If yes, then it's because N lacks normative force - aka N is normatively escapable

And if no, then it's because N is jurisdictionally escapable.

What this means is that the landscape of normativity is in an important respect more fine-grained than has so far been appreciated. It is worth seeing the extra

fine-grainedness that is introduced when we take seriously that there are two distinctions here, not just one. One consequence is that raises the at-least-in-theory possibility that you could have one type of escapability/inescapability without the other. It will be helpful to visualise what the different regions of logical space are with these two distinctions properly separated.



In many cases, it is going to be a matter for further debate what, if anything, goes in each of these boxes. I expect most people would take the norms of chess, sudoku, sports, etc. to be in the bottom right box. i.e. that these are norms that are both normatively and jurisdictionally escapable.

I would be inclined (but will not argue for this here) to place the norms of etiquette and the norms of fashion in the top right box. i.e. to take these to be norms that are normatively escapable but not jurisdictionally escapable. If I'm right about this, then these would be norms that whose jurisdiction one cannot

escape, but which one can faultlessly counterinstantiate since they lack normative force.

I expect that the (or at least a) mainstream view on moral norms would place them in the top right box. They are inescapable both normatively and jurisdictionally. You can't step out of their jurisdiction, and they have normative force in their own right.

I think that epistemic norms are in the bottom left box. Jurisdictionally escapable but normatively inescapable. I've argued for the first part of this, but not the second. I'll have more to say on the second in 2.2. If all of this is correct, the table would look like this:

		normative force	
		not escapable	escapable
jurisdiction	not escapable	moral norms	etiquette fashion
	escapable	epistemic norms	chess sudoku tennis

An interesting aside: there appears to be a parallel distinction concerning categoricity. The kind of categoricity traditionally attributed to Kant:

Kantian Categoricity: A *should* is categorical if it *applies to you* independently of whether you have certain desires or aims. So supposing that there is a categorical requirement to treat people as ends in themselves, for example, what that means is that regardless of your desires and aims, it's true of you that you should treat people as ends in themselves. Your lacking certain desires and aims *doesn't stop it being true of you that you should*.

Kantian categoricity, so understood, is a categoricity of *jurisdiction*.

Interestingly, metaethicists sometimes talk about a different kind of categoricity, which is essentially categoricity of normative force.

Categoricity of Normative Force: A *should* has categorical normative force if it *has normative force for you* independently of whether you have certain desires or aims. You have a reason to do as it says *independently of* your desires and aims

As Philippa Foot has argued, it's plausible that some norms have categoricity of jurisdiction without having categorical normative force¹⁰. The norms of etiquette and the norms of fashion might be examples. With these, arguably it's not your desires and aims that make them *apply to you*. They apply to you regardless of your desires and aims. So their jurisdiction is categorical.

However, your desires and aims determine whether they *have normative force for you*. So their normative force is not categorical.¹¹

What is the relevance of this to my argument here? Well, it is a distinction of the same kind as the distinction I've drawn here, just made in the case of categoricity

¹⁰ Foot 1972. How exactly we should interpret Kant is a further exegetical question that I set aside. Foot argues that Kant cannot have just meant (in my terminology) the jurisdictional version, and that he needs the normative force version too (see 1972, p. 308).

¹¹ Others who make this distinction include Joyce 2001, who uses the terms *weak categoricity* and *strong categoricity*; Tubert 2010; Cote-Bouchard 2016.

rather than escapability.¹² It's analogous - same distinction, applied to something slightly different. So it shows, first, that the distinction between jurisdiction and normative force that I've drawn above is in fact one that we already recognise and already have independent reason to recognise. The distinction Foot and Joyce are talking about, *is* a distinction between normative force and jurisdiction, even if they don't articulate it in these terms.

¹² How categoricity and escapability relate to each other is a question for further debate. A natural starting thought to have might be that the categorical-hypothetical distinction and the escapable-inescapable distinction map onto each other, so that all norms that are categorical are inescapable, and all norms that are hypothetical are escapable. However, there are reasons to doubt they map onto each other perfectly. Norms that are independent of the agent's desires could still be escapable. I've argued here that being contingent on the agent's desires and aims might be *one way* a norm can be escapable. But if there are other ways (e.g. by being contingent on whether the agent is doing a particular activity), then a norm can be escapable without being hypothetical. Perhaps there are more ways still.

This helps to illustrate what the distinction amounts to. It also adds a little extra, further motivation for thinking the distinction exists, by pointing out that we already recognise that it exists in a slightly different case.¹³

§2.2 Upshot 2: Constitutivism and Realism

A second upshot of the distinguishing jurisdictional and normative escapability, concerns what kind of metanormative view constitutivism is. In particular, where it sits on the realism-antirealism divide. Constitutivism is usually classified as an antirealist position:

Katsafanas: ‘Constitutivism is broadly anti-realist’ (‘Constitutivism about

¹³ How categoricity and escapability relate to each other is a question for further debate. Perhaps a natural starting thought to have would be that the categorical-hypothetical distinction and the escapable-inescapable distinction map onto each other, so that all norms that are categorical are inescapable, and all norms that are hypothetical are escapable. If I am right that there are two kinds of categoricity and two kinds of escapability (jurisdictional and normative), then a follow up question would be: do *both* versions of the two distinctions map onto each other? That is:

- (i) Does the distinction between categorical and hypothetical jurisdiction map onto the distinction between escapable and inescapable jurisdiction? and
- (ii) Does the distinction between categorical and hypothetical normative force map onto the distinction between escapable and inescapable normative force?

I think there are reasons to doubt they map onto each other perfectly in either case. In both cases, the thought that disrupts the mapping is this: being contingent on the agent’s desires and aims might be *one way* a norm can be escapable. But if there are other ways, then a norm can be escapable without being hypothetical. In which case a norm’s being categorical does not entail its being inescapable.

On the jurisdiction version, I’ve argued here that there are indeed other ways a norm can have escapable *jurisdiction*. One such way, as I’ve focussed on here, is by being contingent on whether the agent is doing a particular activity. The agent can then escape the jurisdiction of the norm by stopping doing the activity in question. Crucially, a norm can have escapable jurisdiction in this way without falling on the hypothetical side of the categorical-hypothetical jurisdiction distinction. If this is correct, then on the jurisdiction version, a norm can be categorical and escapable. So the answer to (i) is no - the distinctions don’t map perfectly, at least when drawn about jurisdiction.

How about the normative force version? Here I am less sure. But as before, the claim that the two distinctions map perfectly here is going to depend on whether something’s normative force being contingent on the agent’s desires and aims, is *the only way* its normative force could be escapable. That *might* be true, but crucially it’s not obvious that it must be. The claim that there are no other ways that normative force can be escapable, is a substantive claim. We’d need to see the arguments for it.

Practical Reasons', p. 28)

Ferrero: 'In all of its forms, constitutivism is supposed to offer an alternative to a realist account of objective authority.' ('Inescapability Revisited', p. 5)

O'Hagan: 'Constitutivists do not endorse the existence of normative truths in this strong realist sense' (Shmagents, Realism and Constitutivism About Rational Norms , p. 24)

de Liège: '[Constitutivists] aspire to show moreover that we can have an objective basis for our normative claims without facing the traditional metaethical obstacles to moral realism.' ('Constitutivism and the source of prescriptive norms')

Tubert: contrasts constitutivism with 'an alternative full-fledged realist view' (Constitutive Arguments)

I suspect that a major reason constitutivism is classified as antirealist is the assumption that constitutive norms must lack normative force. And I suspect that a major reason for thinking this, comes from noticing that constitutive norms have a kind of escapability, and assuming that this must be normative escapability.

However, clearly distinguishing jurisdictional escapability from normative escapability reveals that this is a bad argument. Constitutive norms do all have jurisdictional escapability, but it doesn't follow from this that they all have normative escapability. This latter claim would require a separate argument.

What arguments could there be for thinking that constitutive norms must have normative escapability? Here are three that might be underlying this thought. I'll argue that none of them succeeds.

(1) Argument from chess norms

It is well documented that *some* constitutive norms are normatively escapable.

David Enoch on chess norms:

Suppose you somehow find yourself playing chess [...], but you do not care about the game and about who wins, nor do you have any reason so to care. It seems rather clear to me that **you have no reason whatsoever to attempt to checkmate your opponent**. And if a metanormative (or metachess) theorist then comes along, explaining to you that attempting to checkmate your opponent is constitutive of the game of chess [...] it seems to me you are perfectly justified in treating this information as **normatively irrelevant**.¹⁴

Enoch is persuasive here, I think, that chess norms are normatively escapable. But this would not establish that *all* constitutive norms must be normatively escapable. From the fact that some of them are normatively escapable, it doesn't follow that they all are. The argument does nothing to rule out the possibility that there could be constitutive norms that have normative force in their own right.

(2) No brute normativity argument

Here's one argument that someone could make for why constitutive norms can't have normative force in their own right. They can't because that would require them to have normative force *as a matter of brute fact*.

The right response to this, I think, would be: why couldn't some constitutive norms have normative force as a matter of brute fact?

To be sure, some philosophers might not like the idea of having to claim that the normativity is brute, and might think this is a theoretical cost. But even if it is, it's

¹⁴ Agency, Shmagency, p. 186

not clear why this would be any more of a cost if the norms you're claiming this about are constitutive norms of some activity, than if they are norms that are not constitutive norms of some activity. Standard normative realists (e.g. moral realists) take certain *non-constitutive* norms to be brutally normative. The Constitutivist realist position I'm outlining, would involve taking certain *constitutive* norms to be brutally normative. it's not clear why that should be considered in any way more costly.

(3) Undercutting defeater argument

An alternative worry might be that it undermines the appeal of constitutivism. You might think that its appeal was precisely that there is the promise of being able to reduce and explain the normativity of the relevant norms, and so if you're going to give that up and just say it's brute and can't be explained, what was the point in saying they're constitutive norms of some activity?

This objection would also miss the point. And the distinction between jurisdictional and normative escapability helps to make this clear. One reason you could have to think the norms in a given domain are constitutive norms of some activity, is that those norms have jurisdictional escapability. This is a reason that has nothing to do with the possibility of being able to reduce and explain their normativity. So if this is your reason for being a constitutivist about a given domain, that reason is not undermined by being a realist and taking them to have normative force as a matter of brute fact. If, for example, you are a constitutivist about epistemic norms on the basis of the asymmetry argument presented in §1, then claiming that epistemic norms have normative force as a matter of brute fact, doesn't in any way undermine your reasons for being a constitutivist here.

Some constitutivists have ambitions of trying to reduce the normative to something non-normative. This is not, I take it, an essential feature of constitutivism, even though it is a feature that many philosophers have found appealing about certain types of constitutivist theory.

If I'm right, then this is an overlooked portion of logical space. The received view that constitutivism is antirealist is mistaken. Constitutivism is silent on the realism question. It's possible to be a constitutivist and be antirealist, but it is also possible to be a constitutivist and a robust realist. What these views all have in common, that makes them types of constitutivism is that they take the relevant norms to be constitutive norms of some activity.

§3. The Shmagency Objection

The distinction between jurisdictional and normative escapability helps to shed light on an influential objection to constitutivism, David Enoch's 'shmagency' objection¹⁵. The objection is presented in terms of moral constitutivism, but it is supposed to be a worry for epistemic constitutivism too. I'll talk through how the objection would go in both cases. There is a fair bit of unclarity about what exactly the shmagency objection amounts to. At its heart is a worry about escapability - that constitutive norms are escapable in a way that would make them a bad fit for moral norms. However, the nature of the escapability in question is left a bit unclear. At times the objection seems to be about jurisdictional escapability:

constitutivism seems to be subject to a powerful objection. For agents need not care about their qualifications as agents, or whether some of their bodily movements count as actions. They can, it seems, be perfectly happy being shmagents – non-agent things that lack the thing purportedly constitutive of agency, but that are as similar to agents as is otherwise possible – or perhaps being something else altogether.¹⁶

But at other times it sounds like Enoch's concern is really normative escapability:

¹⁵ See Enoch 'Agency, Shmagency' *The Philosophical Review*, 2006. See also Enoch 'Shmagency Revisited'

¹⁶ Enoch 'Shmagency Revisited', pp. 2-3

Suppose you somehow find yourself playing chess [...], but you do not care about the game and about who wins, nor do you have any reason so to care. It seems rather clear to me that **you have no reason whatsoever to attempt to checkmate your opponent.** And if a metanormative (or metachess) theorist then comes along, explaining to you that attempting to checkmate your opponent is constitutive of the game of chess [...] it seems to me you are perfectly justified in treating this information as **normatively irrelevant.**¹⁷

even if you find yourself engaging in a kind of an activity, and indeed even if you find yourself *inescapably* engaging in it [...] and even if that activity is constitutively governed by some norm or is constitutively directed at some aim, **this does not suffice for you to have a reason to obey that norm or aim at that aim.**¹⁸

¹⁷ Agency, Shmagency, p. 186

¹⁸ Enoch, Shmagency Revisited, pp. 5-6

Jurisdictional Escapability Reading

If the worry is about jurisdictional escapability, it tells us something interesting about why constitutivism is more plausible in the epistemic case than the moral case.

On the jurisdictional reading, the objection to moral constitutivism is that moral constitutivism predicts that moral norms will be on the bottom half of the table, but in fact they are on the top half.

	normative force	
	not escapable	escapable
jurisdiction		
escapable		

This looks like a pretty good objection in the moral case. It's plausible that moral norms are one the top half of the table, and yet, as I've argued, constitutivism would predict it's on the bottom half.

It doesn't transfer well onto the epistemic case, however. Here's how it would look in the epistemic case:

		normative force	
		not escapable	escapable
jurisdiction	not escapable		
	escapable		

Epistemic constitutivism predicts that moral norms will be *on the bottom half of the table, but in fact they are on the top half*.

If what I've argued in §1 is correct, this argument is entirely unconvincing, since epistemic norms are exactly where epistemic constitutivism predicts they will be. Epistemic norms are on the bottom half, just as constitutivism predicts. So there is no mismatch here.

What we're seeing here is that jurisdictional escapability holds the key to why constitutivism works better in the epistemic case than the moral case. The very feature that makes constitutivism a potentially awkward fit for moral norms, makes it a strikingly *good* fit for epistemic norms.

But while there is a pretty good objection to moral constitutivism on this reading of the shmagency objection, the distinction between jurisdictional and normative escapability reveals that it's not a knockdown objection, but more of a challenge or an awkwardness that the moral constitutivist would need to fix.

A popular response to the shmagency objection has been to argue that the activity that moral norms govern is one that it is impossible for us to not do¹⁹. So for example, if you think that moral norms are constitutive norms of agency, then if we can show that agency something we cannot help but be doing, then we block this escapability worry, by blocking the escape from moral norms by making the activity they govern one that is impossible for us to stop doing).

This looks like a promising line of response since if the worry is about jurisdictional escapability - if the worry is that constitutivism allows the possibility of escaping the jurisdiction of moral norms by stopping doing the activity that they govern - then if it could be shown that the relevant activity is one that it's not possible to stop doing, that type of escapability would be blocked.

Of course, there's a question of whether this can be done - whether it's plausible that the relevant activity is one we can't stop doing. But interestingly, Enoch's objection to this response is not that this can't be done, but rather that this strategy misses the point. According to him, what's needed is not for agency to be impossible to not do, but rather for us to *have a reason* to be agents.²⁰ But if the worry really is about jurisdictional escapability, then these responses are right on target, and it's Enoch's response to them that misses the point.

¹⁹ See, for example, Ferrero, L. L. (2009). 'Constitutivism and the inescapability of agency,' *Oxford Studies in Metaethics*, 4, pp. 303– 333. Katsafanas, P. (2018) 'Costitutivism about Practical Reasons'

²⁰ See Enoch 'Shmagency Revisited', esp. pp. 5-6: 'One of the points I emphasized in "Agency, Shamgency" (185) was that even if you find yourself engaging in a kind of an activity, and indeed even if you find yourself *inescapably* engaging in it [...], and even if that activity is constitutively governed by some norm or is constitutively directed at some aim, this does not suffice for you to have a reason to obey that norm or aim at that aim. Rather, what is also needed is that you *have a reason* to engage in that activity.'

This suggests that the jurisdictional escapability argument is not the argument Enoch is intending to make. Because if it were, it's hard to see why the strategy of trying to show that the relevant activity is one we can't stop doing, wouldn't be a good one.

And moreover, the fact that Enoch thinks that what would fix the problem is our having a *reason* to be agents, suggests that his worry about constitutivism is a worry about normative escapability. After all, it's normative escapability, not jurisdictional, that would be fixed by having a reason. Think of norms of etiquette and fashion. They have no normative force in their own right, but they would have normative force for a given agent *if that agent had a reason to do as they say*.

Normative Escapability Reading

If the worry is about normative escapability, the objection is mistaken in both the moral and the epistemic case. Here's how the objection would go in the moral case:

Moral constitutivism predicts that moral norms will be on the right hand side of the table, but in fact they are on the left hand side.

And in the epistemic case:

Epistemic constitutivism predicts that moral norms will be on the right hand side of the table, but in fact they are on the left hand side.

		normative force	
		not escapable	escapable
jurisdiction	not escapable		
	escapable		

The objection is mistaken in the same way in both cases. If what I argued in 2.2 is correct, then firstly, there is no good reason to think that constitutive norms must all be on the right hand side. Secondly Enoch's argument for the normative escapability of constitutive norms fails, since it relies on inferring from the normative escapability of chess norms, to the normative escapability of *all* constitutive norms. Until a better argument comes along, there is no reason for the constitutivist to be worried here. The normative inescapability of moral and epistemic norms, poses no tension with constitutivism about these norms.

Conclusion

I've argued that a hallmark of constitutive norms is jurisdictional escapability, and that distinguishing this clearly from normative escapability reveals a number of important and so-far-unnoticed truths about constitutivism, and about normativity generally. It forms the basis of a new argument for epistemic constitutivism. It reveals that, and why, constitutivism is a much better fit for epistemic norms than it is for moral norms. And finally, it reveals a new dimension to the normative landscape, and highlights that there is further work to be done on locating the various different norms in the relevant regions of logical space.