

On Exercising Epistemic Agency

Mona Simion

Mona.simion@philosophy.ox.ac.uk

Penultimate Version

Forthcoming in Criminal Law and Philosophy

1. Introduction¹

Scepticism about epistemic agency is widely spread. Agency is about making our own decisions in an importantly autonomous fashion. We can't form beliefs at will, however, so the epistemic is simply not the proper locus of agency. Or so the sceptic thought goes.²

In contrast, optimists about epistemic agency point out that the sceptic is making a mistake in thinking that agency requires a level of control that eludes belief formation. Agency, they argue, simply requires the capacity to act for reasons.³ *Mutatis mutandis*, the thought goes, epistemic agency is cheap: insofar as your belief formation is evidence responsive, you've got the capacity to act for epistemic reasons. You're fine.

This paper is an exercise in optimism. However, on my preferred account, and contra traditional optimism, exercising epistemic agency is not a matter of being evidence responsive, but rather simply a matter of having an unobstructed capacity to respond to evidence. I show that this minimalist account is strongly superior to traditionalist optimism in explaining the contrast between cases of subpar exercises of epistemic agency on one hand, and cases of agential epistemic injustice (Lackey 2023).

Here's the game plan: I start with briefly discussing and dismissing scepticism about epistemic agency. I will then move on to optimist, reasons-responsiveness-based accounts, and show how they fare when it comes to explaining cases in which our agency is, respectively, bypassed, exploited, or subverted (Lackey 2023). Finally, I will offer some problem cases for the optimists' way of understanding exercises of agency, and put forth my preferred account.

2. Against Agency Scepticism

Are we epistemic agents? The question has considerable importance across epistemology, philosophy of action, ethics, and the law. If the answer is yes, then we will be responsible for our beliefs in ways that ground proper criticism, praise, blame, and legal accountability; epistemic norms might be understood as genuine person-level requirements – requiring things of us as agents, rather than merely governing our epistemic processes; and the normative evaluation of epistemic conduct becomes continuous with the broader normative evaluation of conduct in general.

The dominant view in the literature is that believing is not an exercise of agency, or at least not in any robust sense. The arguments are varied and well-known: believing cannot be done at will (Alston 1988); there is no phenomenology of doxastic choice (Williams 1973); deliberating about whether to believe is transparent to deliberating about truth (Shah 2003); belief is a state one is in, not an action one performs – and thus it's more akin to being tall (clearly non-agentive) than to running (clearly agentive) (Shah and Velleman 2005); and so on. These arguments are widely treated as very convincing. The result is a picture according to which epistemic normativity must be explained in ways that do not presuppose genuine agency, nor genuine epistemic responsibility.

The sceptic serves us with a modus tollens: If we bear responsibility for our beliefs, then our beliefs are exercises of our agency; beliefs are not exercises of our agency (since they lack the phenomenology of choice, they cannot be formed at will, they're states rather than actions....); therefore, we don't bear responsibility for our beliefs.

I want to suggest modus ponens-ing on the skeptic's modus tollens. Here it goes: We do bear responsibility for our beliefs (as per robust data supporting appropriateness of epistemic blame, criticism,

¹ This research has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 948356 – KNOWLEDGELAB and grant agreement No 101199255 – KNOW-HOW).

² E.g. Alston 1988, Williams 1973, Shah 2003, Shah and Velleman 2005.

³ E.g. McHugh 2013, Hieronymi 2008, Lackey 2023, Boyle 2009.

and praise);⁴ If we bear responsibility for our beliefs, then our beliefs are exercises of our agency; therefore, beliefs are exercises of our agency.

Whether it is more appropriate to go with *modus ponens* or *tollens* on a conditional “if *p*, then *q*” depends on the comparative evidential strengths of *p* and not-*q*. If we’ve made a really good case that *p*, and the case that not *q* is flimsy, we’d be foolish to give up our commitment to *p* because it’s been anti-correlated with a poorly supported claim.⁵ I submit that the data supporting the claim that we bear responsibility for our beliefs is overwhelming: criticism of bad beliefs is widely spread, as is the practice of discursively justifying our beliefs in order to defend against said criticism (rather than, e.g. dismiss it as inappropriate). If so, we can safely conclude one of two things: either the features the alleged absence of which is cited as evidence for lack of epistemic agency are not essential to epistemic agency after all, or, alternatively, beliefs do have the features in question.

We could stop our anti-scepticism case here. But we don’t have to: let’s look at the first disjunct above: one thing that immediately strikes me as rather peculiar in the sceptic’s toolbox is the overwhelmingly dominant focus on belief: why think that belief is the only locus of epistemic agency? After all, epistemic norms, literature suggests, govern all kinds of other mental and non-mental phenomena: for instance, assertions⁶ - and speech acts more generally - are said to be governed by epistemic norms, and the main debate is on the content of these norms, far less so on whether they are there to begin with. Surely my assertions are exercises of my epistemic agency. Judging, reasoning, wondering, inquiring are phenomena that are governed by epistemic norms and share in the level of agency of any garden variety action. If so, the sceptic’s argument, far from supporting scepticism about epistemic agency, merely raises some questions about doxastic agency. Epistemic agency is safe.

Second, it will be useful to note that all of the worries above share a common structure. Each identifies a feature *F* that beliefs allegedly lack and argues that since *F* is necessary for agency, belief is not agentive. The arguments differ in which feature *F* they target, but they all turn on the same conditional: if *x* is agentive, then *x* has *F*. In turn, the plausibility of these conditionals derives from reflection on paradigm cases of agency: intentional actions, deliberate choices, voluntary performances. From these paradigm cases a conception of agency is built up, and the conditionals are read off that conception. Why so? The Standard View (e.g. Mele 2003) of the nature of agency unpacks agency as the capacity to act.⁷ Folks disagree about how to unpack this further, but tend to be in broad agreement about the minimal Standard View. If, however, the thought would go, agency is the capacity to act, and belief is a state rather than an action, we should not be terribly surprised there’s no such thing as doxastic agency. And one popular thought is, once more, that belief is, in virtue of being a doxastic state, more akin to being tall (which is clearly non-agentive) than to running (which is clearly agentive).

Consider, however: sitting in a chair, holding a glass of wine, wearing a hat. These are all states. They are not actions in any standard sense: I gather, at a minimum, that sitting is a paradigmatic case of a state rather than an action. And yet these states are patently not devoid of agentive involvement in the way that being tall is devoid of agentive involvement. There is something deeply different between my sitting in this chair—a state that falls squarely within the purview of my agency—and my being of the height I happen to be—a state that has nothing to do with my agency at all. Sitting, wearing a hat, holding a glass, and, yes, believing, are agentive states. Since it should not come as a surprise that some properties of garden-variety actions are not shared by agentive states (after all, they’re quite distinct metaphysical beasts), the argument from ‘paradigmatically, actions have property *P*’, to ‘since belief doesn’t have it, it is not an exercise of agency’ should not move us much. Certainly not as much as the overwhelming epistemic responsibility data. We should go with the *modus ponens*.

It is worth noting also that the category of agentive states is not identical with the category of states that were caused by agentive activity. Being tanned is a state caused by the agentive activity of sunbathing. But once formed, the tan persists independently of the agent’s agency. Conversely, some agentive states may not have been initiated by a clear agentive act. One might be sitting in a particular posture because

⁴ See e.g. Brown 2020, Boulton 2021, Kauppinen 2018, Kelp and Simion 2017.

⁵ I follow Cappelen and Dever forthcoming who make this point against AI cognition scepticism.

⁶ See e.g. Lackey 2007, Williamson 2000, Goldberg 2015, Kelp & Simion 2021. For a full epistemic normative account of constative speech see Simion 2021.

⁷ Several philosophers in this tradition like to think of agency as the capacity to act intentionally. This view has lost popularity in more recent times due to dissatisfaction with its focus on the paradigmatic cases of agency. More minimal accounts of agency have been developed, accounts that are better at explaining both cases of non-human agency, and cases of human agency that do not plausibly come with intentional import. See Schlosser 2019 for a very useful overview.

that's how they happened to fall on the chair — yet the state is now an exercise of one's agency. The agentiveness of the state is not determined solely by its causal history but by its ongoing character.

If so, the Standard Account – of agency as the capacity to act – requires a broad understanding of action that includes agentive states. Doxastic agency is safe.

3. Agency, Responsiveness to Reasons, and Injustice

Shall we then say that an X is an agent simply if they have the capacity to act in this broader sense, including agentive states? If so, how should we characterise a particular type of agency – such as moral, or epistemic agency? An attractive proposal that many optimists about epistemic agency like (McHugh 2013, Hieronymi 2008, Lackey 2023) goes along the following lines: to be an agent of type T is to have the capacity to act in ways that are (at least minimally) responsive to T-type reasons. The view is illuminating and helpful: it neatly explains why one might think that I am a moral agent but my cat is not – I have the capacity to act in ways that are responsive to moral reasons, my cat less so (at least for all we know about cats – we could be wrong). It also neatly explains why one's T-agency may come in degrees: I'm more of a moral agent, one may think, than my 3-year-old kid, simply because my capacity to act in ways that are responsive to moral reasons is more robust. In turn, this also neatly explains the different levels of moral responsibility involved in my, and respectively, my kid's actions. Finally, it gives us a straightforward account of epistemic agency: X is an epistemic agent insofar as they have the capacity to act in ways that are (at least minimally) responsive to epistemic reasons (McHugh 2013). Optimism about epistemic agency is elegant and explanatorily powerful.

Jennifer Lackey is, in the sense used here, an optimist about epistemic agency, and she follows tradition in unpacking it in terms of capacity for reasons responsiveness. In her recent wonderful book 'Criminal Epistemic Injustice' (2023), Lackey proposes that the account also allows us to identify a novel variety of epistemic injustice (Fricker 2007): agentive epistemic injustice (AEI). Lackey is interested in the epistemic injustice involved in forced confessions. She argues that interrogation tactics aimed at extracting testimony undermine epistemic agency via undermining epistemic reasons responsiveness in three distinct ways: by bypassing it, exploiting it, or subverting it entirely. When this happens, according to Lackey, the subjects are victims of agential epistemic injustice. Here is Lackey:

Epistemic agency is often understood as being grounded in a subject's responsiveness to reasons or evidence, and so epistemic agency is exercised with respect to a subject's testimony when that testimony is responsive to reasons or evidence. In the interpersonal case—where the testimony of one person is brought about by another — *rational persuasion* is a paradigmatic example of obtaining testimony while going *through* a speaker's epistemic agency. If, say, I present compelling data on behalf of there being anthropogenic causes of climate change, and you come to say that this is so on the basis of the force of the proffered evidence, I have elicited testimony from you by appealing to your rational capacities. [...] Eliciting testimony by going through your rational capacities involves the straightforward exercising of your epistemic agency. In contrast, epistemic agency is bypassed, exploited, or subverted when a subject's ability to be responsive to reasons is circumvented, abused, or undermined (Lackey 2023, Ch. 2).

On the topic of bypassing epistemic agency, Lackey gives the example of police interrogation tactics that don't engage a suspect's rational faculties at all, but instead work around them. Tactics like isolation, sleep deprivation, extended questioning, and food deprivation are pressures that increase the risk of a false confession, but also, importantly, deliberate strategies interrogators use to wear down epistemic agency. Because these tactics don't offer suspects genuine epistemic reasons to confess, but instead maneuver them into a psychological state where confession becomes inevitable, Lackey classifies this as manipulation—a mode of influence that circumvents epistemic agency rather than engaging it.

A second way in which interrogation tactics undermine epistemic agency, according to Lackey, involves outright deception, which Lackey distinguishes from manipulation because it takes advantage of the victim's reasons responsiveness rather than bypassing it. Cases of lies calculated to exploit justified trust in authority figures aim to give the subject of the interrogation (misleading) reasons to confess to a crime that they did not commit, despite having no recollection of doing so. In cases like these, interrogators count on the suspects to reason well from false premises to a false conclusion—using their capacity for reason responsiveness against them, rather than sidestepping those capacities altogether.

The third and most severe category is coercion, where suspects are left with no genuine alternative but to comply. Threatened with execution, the subjects see confession as his only viable path. Lackey compares this to a robbery at gunpoint—technically the victim "chooses" to hand over money, but only one real option exists. Because no reasonable alternative to falsely confessing remains available, the suspect's capacity to exercise their own epistemic agency - to only make epistemic-reasons-responsive assertions - is subverted.

Let's take stock: we have seen that scepticism about epistemic agency remains unmotivated in more than one way. In contrast, epistemic agency optimism comes with a plethora of theoretical virtues. First, it is extensionally adequate: it is able to explain why I am a genuine epistemic agent, while my thermometer is not. I, but not my thermometer, have the capacity to engage in epistemic-reasons-responsive action. It is also able to explain overwhelming data concerning epistemic responsibility attribution: since I have the capacity to engage in epistemic-reasons-responsive action, I am the proper subject of epistemic criticism, blame, and praise. It also predicts that the more robust a capacity for epistemic-reasons-responsiveness one hosts, the higher one's degree of epistemic agency, and thereby the higher one's level of epistemic responsibility (we don't blame three-year-olds for their false or unjustified assertions in the way we blame adults). Finally, and crucially, it gives us the resources to theorise about an important way in which we can be the victims of injustice: when our capacity for epistemic reasons-responsiveness is bypassed, exploited, or subverted, we are the victims of agential epistemic injustice.

4. Expressions of Epistemic Agency

When is an act an expression of one's epistemic agency? We have seen that Jennifer Lackey takes epistemic agency to be exercised with respect to a subject's testimony when that testimony is responsive to reasons or evidence. On the normatively generalised version of the view, then, and where T is the relevant type of normativity, my phi-ing is the expression of my T-agency insofar as it is T-reasons-responsive. For the case of epistemic agency, Lackey's cases are a clear illustration of what the view predicts. Cases in which the subject testifies as the direct result of extensive torture and disorientation will, on this account, be cases in which the assertion in question is not an exercise of their epistemic agency in virtue of not being epistemic-reasons-responsive. In this sense, their epistemic agency is bypassed. Similarly, cases of subversion will also be cases in which the testimony is not entirely epistemic-reasons-responsive – because overwhelmingly practical-reasons-responsive (the subject testifies because doing so is the practically dominant option), so, once more, the account will predict it not to be a full expression of one's epistemic agency, since the latter is subverted by the forced choice. In contrast, cases of exploitation will be, indeed, cases in which one's agency is involved in testifying (the testimony responds to misleading testimonial evidence aimed to defeat the subject's own memorial evidence).

In what follows, however, I will argue that, in spite of its initial appeal, this way to unpack exercises of agency gets things wrong - both generally, but most importantly for our purposes, with regard to agential injustice. Furthermore, I will also look at an alternative, modal way to unpack the view, and argue that it will not make things better. If I am right, the champion of agential epistemic injustice requires a novel account of what makes an act a genuine manifestation of epistemic agency. Here it goes.

The main worry is that Lackey's favourite account of what it means for an act to be an exercise of epistemic agency threatens to fail to distinguish between bad, or subpar exercises of agency, and actions that do not express one's agency. If acts are, as per Lackey's view, exercises of my epistemic agency only insofar as they are epistemic-reasons-responsive, my engaging in actions in resistance to epistemic reasons will be categorised as a lack of manifestation of epistemic agency, rather than a problematic manifestation of epistemic agency.

To see this, let's start with moral agency; consider a case in which I am driving by a lake and see a child drowning. Say that I chose to ignore them and continue driving, excited to come to your excellent party. In this case, an account of agential exercises like Lackey's will predict that my action is not an expression of my moral agency (but rather merely of my prudential agency, perhaps), since it is not moral-reasons-responsive. But clearly that is the wrong result here: my action is an expression of my moral agency: a bad one! In spite of the fact that I have the capacity to engage in moral-reasons-responsive action (save the kid), I chose to go to the party. Moral agency can manifest in failures of reasons responsiveness just as much as in successes. That's why I can be a better or worse moral agent. And moral responsibility will follow suit: you will hold me morally responsible for my action because it is an expression of my agency.

Going back to epistemic agency, consider a sexist cognizer – let's call him George: this fellow has the capacity to form beliefs that are responsive to epistemic reasons (Simion 2024a,b), but fails to do so

when women speak. Now consider a case in which George's male colleague Dan says that *p*, where Dan is not only a non-expert but also so well-known a liar on the relevant topic that his testimony has absolutely no evidential weight. At the same time, his female colleague Anna is an expert on the topic and asserts that not-*p*. Say that George forms a belief that *p*. On Lackey's account, this belief is not an expression of George's epistemic agency – since it's not epistemic-reasons-responsive. But George's epistemic failure is clearly his own in an important sense in which Lackey's subject of torture is not.

Things get worse when we contrast the cases of agential epistemic injustice above with simple cases of bad testimony: we certainly don't want to predict that cases of lies and bullshit (which are paradigmatically not epistemic-reasons-responsive but, at best, prudential-reasons-responsive) are agentially equivalent to Lackey's cases of bypassed or subverted epistemic agency – in that in both types of cases the subjects' epistemic agency is not expressed by their testimony. To the contrary, we want to be able to say that the testimony of liars and bullshitters is an awful expression of their subpar epistemic agency.

What to do? The prime source of inspiration for traditional optimism about epistemic agency (e.g. McHugh 2013) is Fischer and Ravizza's (1998) reasons-responsiveness-based account of moral agency. If so, Fischer and Ravizza's account of what makes an act an expression of one's agency might be thought to help here. According to them, my *phi*-ing expresses my agency just in case I exercise what they call 'guidance control' over it. In turn, "an agent exercises guidance control of an action if and only if that action issues from the agent's own moderately reasons-responsive mechanism.". The two crucial components of guidance control, then, are ownership and responsiveness to reasons. I'll set ownership aside for our purposes here, it won't matter much. Crucially for our purposes, Fischer and Ravizza think a mechanism is moderately reasons-responsive if and only if (1) it is at least regularly *receptive* to reasons, some of which are moral reasons, and (2) it is at least weakly *reactive* to reasons (but not necessarily moral reasons). Fischer and Ravizza's receptivity is a modal epistemic property. For it to be instantiated, it needs to be that (1) there are possible scenarios in which (i) there is sufficient reason to do otherwise, (ii) the same kind of mechanism is operative, and (iii) the agent recognizes the sufficient reason to do otherwise, and (2) that (1) constitutes what they call "an understandable pattern of reasons-recognition", in that it's not too random). Reactivity, in turn, is a modal property has to do with a 'capacity to translate reasons into choices (and then subsequent behaviour)'. An actually operative kind of mechanism is weakly reactive to reasons if and only if there is some possible scenario in which (1) there is sufficient reason to do otherwise; (2) the same kind of mechanism operates; (3) the agent recognizes the sufficient reason to do otherwise; (4) the agent thus chooses and does otherwise for that reason.

The fine details of the account will not matter much for our purposes. What we care about here is that Fischer and Ravizza offer a modal (sensitivity-type) account of what makes an action an exercise of one's agency: instead of requiring actual reasons responsiveness, this view looks at the dispositional profile of the mechanism that produced the action in question in order to settle the question as to whether it is an expression of the subject's agency. Can this help? Let's substitute the moral reasons in Fischer and Ravizza's account for epistemic reasons and find out.

The view may indeed help with paradigmatic cases of epistemically bad testimony: in paradigmatic cases of lying and bullshitting, the thought goes, these speech acts are, as a matter of fact, the product of mechanisms that meet Fischer and Ravizza's conditions. Presumably, most testifiers employ assertion-generating mechanisms (or processes, or capacities – the reader is free to choose in accordance with their epistemological preferences here) that are such that they are regularly receptive to reasons, some of which are epistemic reasons, and also regularly reactive to reasons (in the relevant sense). Liars are usually pretty good at figuring out when there is epistemic reason to assert otherwise, as well as to translate reasons into choices (indeed, liars usually lie for prudential reasons).

There are two problems with this route, however: first, the account threatens to predict that Lackey's cases of bypassed and subverted epistemic agency are, contra Lackey, cases of epistemic-agency-expressing testimony. After all (perhaps depending on how one individuates the relevant 'mechanism'), the subjects of interrogation's testimony is sourced in a mechanism that does well, counterfactually: the subjects' assertion generating mechanisms are plausibly pretty good at recognising when there is epistemic reason to assert otherwise in counterfactual situations, as well as to translate reasons into choices (indeed, liars usually lie for prudential reasons). Or it's not clear why we should believe otherwise. Furthermore, we can just go ahead and stipulate that the liar and the victim of torture, while finding themselves in very different situations in the actual world, share the exact same modal profile with regard to their testimony-generating mechanisms. If so, what is the difference between them with regard to agential exercises? Perhaps a way out here would be to blame the result on a generality problem: individuating the testimony-

producing mechanism in the case of bypassed agency should be such, perhaps, that the torturous interrogation tactics used to extract the testimony are part of the mechanism of producing it. But even so, one might think, the case of subverted agency now looks like a garden variety case of full agential action, on a par with liars who act for prudential reasons rather than epistemic reasons.

The second, much more serious problem is that, on closer inspection, the view mishandles agency attributions in the cases of liars and sexists as well. To see this, we just need to dial up the epistemic badness of these cases. Take someone who has a very robust disposition to lie for no good reason. This fellow will fail the reactivity condition in Fischer and Ravizza's account, so their lies will come out as independent of their epistemic agency, instead of a bad manifestation thereof. Surely, this doesn't make this fellow's lies less his own, nor is he less epistemically responsible for them. Dialling up lying can't make things epistemically better.

If all this is so, it is overdetermined that we should not prefer a Fischer and Ravizza-inspired account of what renders an action an expression of one's epistemic agency over the simple account that Lackey employs.

5. Exercise Minimalism

We have seen that traditional optimism about epistemic agency, spelled out in terms of hosting a capacity for epistemic reasons responsiveness, has several attractive theoretical virtues. We have also seen that the literature struggles to identify an associated account of what makes an action a genuine exercise of one's epistemic agency. The nut that's hard to crack is distinguishing genuine cases in which epistemic agency is not manifest, from cases in which it simply manifests in subpar ways. This section offers an alternative way to think about exercises of epistemic agency.

To start us off, it will be useful to distinguish here between having the capacity to ϕ and the ability to ϕ (others (Honoré 1964, Mele 2003) in work on agency put this distinction in terms of general versus specific, or local ability – but the point is the same). An ability to ϕ can be absent in cases in which the capacity to ϕ persists; Place a piano virtuoso under water, and they won't be able to play the piano. Compatibly, they preserve their capacity to do so. Similarly, if I am locked in my car and I see a kid drowning, I am not able to save the child (and thus exercise my capacity to respond to moral reasons), while preserving the capacity for moral-reasons-responsiveness. Finally, if I'm in Fake Barn County, I have no ability to come to know, while my knowledge-generating capacities remain intact.

Abilities can be lost in unfriendly situations through no fault of my own – like above. But it can also be that I lose my abilities by altering my own environment for the worse. When that happens, my preventative actions are bad exercises of my relevant agency, even though the actions they prevent are not. I lock myself in the car as soon as I see a kid drowning, I will not be able to save them. Compatibly, I preserve my capacity to do so. Because my capacity is intact, and my ability is lost by self-sabotage – i.e. by an act that is a bad exercise of my capacity to respond to moral reasons – , the fact that I don't save the child is a stain on my moral agency, not a symptom of its absence. Moral responsibility will follow suit. So will epistemic responsibility: if I form beliefs based on testimony in environments that are notably highly polluted with misinformation, my ability to know is impaired through my own fault. Once more, though, my epistemic capacities stay put, and I am the proper subject to epistemic blame for my bad exercise of my capacity to respond to epistemic reasons that is involved in choosing to inquire in a known epistemically polluted environment.

It will be useful to think of these cases in light of Ernie Sosa's classic SSS account of competence (2015): the seat of my epistemic competence is what I call epistemic capacity: it's that feature of my cognitive architecture that makes it such that I can respond to evidence. The third S stands for situation: cases above are cases of bad situations blocking exercises of my relevant capacity.

There is yet another way in which my capacity for reasons responsiveness can fail to manifest: it can happen in spite of being in a favourable environmental situation, due to being in bad shape. If I broke my leg, I preserve my relevant moral capacity when it comes to saving the child from drowning (my capacity for moral reasons responsiveness) but it's blocked from manifesting in saving the child by my bad shape. Similarly, if you put a pill in my drink that makes me hallucinate pink elephants, I preserve my epistemic capacities, but they are blocked from manifesting by my bad shape.

Just like bad situations, bad shape can also come about through no fault of my own (like above), or through my own agency. If I pour a bottle of wine down my throat as soon as I see the drowning child, such that now I'm very drunk and unable to save the child, I preserve my relevant capacity for moral reasons responsiveness, but it's blocked from manifesting in saving the child by my own (moral-agency-involving)

problematic actions. Moral responsibility will follow suit: I will be blameworthy for not saving the child. Similarly, sexists are not able to believe a word women say due to their sexist bias. Compatibly, they have the capacity to do so: the type of cognizer that they are, with their cognitive architecture, is well equipped to properly respond to testimonial evidence (Simion 2024). This capacity, however, is blocked through their (epistemic (and moral)-agency-involving) acquisition of an epistemic (and moral, in this case) vice. Once more, epistemic responsibility follows suit.

In light of this discussion, let's go back to the cases we are interested in. Recall that we want, in a first instance, is a view of what it is to exercise epistemic agency that is able to distinguish cases of agential epistemic injustice from mere bad exercises of epistemic agency. Now, consider the following minimalist proposal about epistemic agentic exercises:

Exercise Minimalism: S's action ϕ is an exercise of S's epistemic agency just in case (1) ϕ is governed by epistemic norms and, (2) S has an unobstructed capacity to respond to epistemic reasons generated by these norms.

To see the appeal of the view, consider it, first, in its moral variation: an action will be an exercise of my moral agency just in case it's governed by moral norms that generate reasons that I have an unobstructed capacity to respond to. For instance, my decision to go to the party instead of saving the child from drowning is an exercise of my moral agency because it is governed by moral norms (that make it impermissible) and I have an unobstructed capacity to respond to the reasons generated by these norms. In contrast, when I am locked in my car, I can't exercise my moral agency: my capacity to respond to the moral reasons requiring me to save the child is blocked. If you locked me in the car, the story ends there: the failure to save the child is not my own, my moral agency is not involved in it, and I don't shoulder the corresponding moral responsibility either. If, however, I locked myself in the car, I am the author of my own obstruction of moral agency, and moral responsibility follows suit.

The same will hold for epistemic affairs. To see this, let's revisit the cases we want to explain. Recall that we want to understand the difference between cases of lying and bullshitting on one hand, and cases of false testimony that occurs due to bypassed epistemic agency (like in the Lackey case where the subject's false confession is extracted via prologued psychological torture) on the other. The view neatly predicts that these are very different cases, both agentially and in terms of responsibility facts: the lies and the bullshit are (bad) exercises of the liar's and the bullshitter's epistemic agency, in that their capacity for epistemic reasons responsiveness is not blocked either at situation level or at the level of shape. In contrast, what is happening in the prolonged torture cases, is that the subject's epistemic agency is blocked at shape level: they preserve the capacity to respond to epistemic reasons, but it is blocked by the bad psychological shape they are in due to prolonged torture. The false confession is neither an exercise of their epistemic agency, nor are they epistemically responsible for it.

The case of the sexist believer makes for an interesting contrast: this fellow's capacity for epistemic reasons responsiveness is also blocked at shape level – so his resulting belief will not be an exercise of his (extant) capacity to respond to epistemic reasons. This case, however, is one where the shape damage is the believer's own fault, so epistemic responsibility facts will constitute the difference with the case of extracted testimony.

In contrast, cases of agency exploitation are correctly predicted by the view to be genuine cases of exercises of agency: the subjects' capacity for epistemic reasons responsiveness remains unobstructed, and the false testimony is generated, indeed, by feeding misleading evidence to said capacity.

Finally, cases of subverted agency – in which subjects are presented with a forced choice between confessing to a murder they did not commit and capital punishment – will be correctly predicted as cases in which the subjects' exercise of epistemic agency, while present in the overall decision to confess, is undermined by their exercise of prudential agency. This is a case where the subjects' prudential agency is being exploited, which results in the undermining of their epistemic agency. To see that this is the correct result, note that the subjects are still able to weigh their epistemic reasons prohibiting making false confessions against the much stronger prudential reasons in favour. Indeed, there is a clear sense in which they are still able to confess truly. It's just that their prudential agency takes precedence in their decision to confess. The case bears agential similarities to the case where I decide to go to the party over saving the child: here too prudential agency takes precedence over moral agency. The difference lies with the quality of my agential exercise: the party case, in contrast of the confession case, is one where prudential reasons do not outweigh moral reasons.

Exercise minimalism gives every distinction we wanted for both epistemic, moral, and prudential agency. It is highly extensionally adequate and explanatorily fruitful. It also scores exceptionally well on simplicity. Before I close, it may be worth reflecting on one theoretical virtue some may worry the account is struggling with: generalisability. Is the view generalisable beyond the three normative domains we discussed? One may worry the view is too minimalist to generalise well. For instance, you might think that it is implausible for other types of agential exercises. It is not true, for instance, that just any act of mine (e.g. sitting on a sofa and watching Netflix) is an exercise of my capacity to play the piano simply of the fact that, as I sit there, I preserve an unobstructed capacity to respond to reasons generated by norms governing piano playing. Agreed. My agency as a pianist is not, in any sense, involved in sitting on the sofa. Note, though, that this points to a significant difference between different types of agency rather than to a problem for my minimalism. We can choose to not exercise our capacity to respond to piano-playing-norms-generated reasons. Our agency as pianists can be put on hold, while we chose to exercise our agency as Netflix watchers. Indeed, we fully exercise our general prudential agency when we make that choice. An engagement condition on agential exercise will plausibly be required for cases of agency that allow for opting out. One cannot, however, chose not to exercise unobstructed epistemic, moral, and prudential agency. The epistemic, the moral, and the prudential, as opposed to piano playing, are not games that we can opt out of. Whenever present, and unobstructed, our capacity to respond to epistemic, prudential, and moral reasons manifests in our actions, judgments, and speech.

6. Conclusion

We are agents. Agential injustice is a crucial way in which we can be victimised. Our capacity to respond to all kinds of reasons generated by all kinds of norms makes us interestingly multifaceted agents, but also opens us up to agential abuse of various kinds. This is one of the most crucial contributions of Lackey's excellent book: When our agency is bypassed, exploited, or subverted, we suffer from agential injustice. When what is at stake is our epistemic agency, we suffer from agential epistemic injustice. To fully account for this important phenomenon, however, we need to correct view of what it is to exercise epistemic agency. I have proposed that Exercise Minimalism is it.

References

- Alston, W. P., 1988, The Deontological Conception of Epistemic Justification, *Philosophical Perspectives*, 2: 257–299.
- Boult, C. 2021. Epistemic blame. *Philosophy Compass* 16 (8):e12762.
- Boyle, M. 2009. Active belief. *Canadian Journal of Philosophy Supplementary* 35 (S1):119–147.
- Brown, J. 2020. What is Epistemic Blame? *Noûs*, 54(2), 389–407. <https://www.jstor.org/stable/48574042>
- Cappelen, H. and Dever, J. Forthcoming. *Going Whole Hog: A Philosophical Defence of AI Cognition*. Oxford: Oxford University Press.
- Fischer, J.M. & Ravizza, M. 1998. *Responsibility and Control: A Theory of Moral Responsibility*. New York: Cambridge University Press. Edited by Mark Ravizza.
- Fricker, M. 2007. *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford: Oxford University Press.
- Goldberg, S. 2015. *Assertion: On the Philosophical Significance of Assertoric Speech*. New York, NY: Oxford University Press.
- Hieronymi, P. 2008. Responsibility for Believing. *Synthese*. 161: 357–73.
- Honoré, A.M., 1964. Can and Can't, *Mind*, 73: 463–479.
- Kauppinen, A. 2018. Epistemic Norms and Epistemic Accountability, *Philosophers' Imprint* 18.
- Kelp, C. and Simion, M. 2017. Criticism and Blame in Action and Assertion. *Journal of Philosophy* 114 (2):76–93.
- Kelp, C. and Simion, M. 2021. *Sharing Knowledge: A Functionalist Account of Assertion*. Cambridge: Cambridge University Press.
- Lackey, J. 2007. *Norms of Assertion*. *Noûs*, 41: 594–626.
- Lackey, J. 2023. *Criminal Epistemic Injustice*. Oxford: Oxford University Press.
- McHugh, C. 2013. 'Epistemic Responsibility and Doxastic Agency.' *Philosophical Issues* 23: 132–57.
- Mele, A. R. 2003, *Motivation and Agency*, Oxford: Oxford University Press.

- Schlosser, M. 2019. "Agency", *The Stanford Encyclopedia of Philosophy* (Winter 2019 Edition), Edward N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/win2019/entries/agency/>>.
- Shah, N. 2003, "How Truth Governs Belief", *The Philosophical Review*, 112(4): 447–482. doi:10.1215/00318108-112-4-447
- Shah, N. and Velleman, J. D. 2005, "Doxastic Deliberation", *The Philosophical Review*, 114(4): 497–534. doi:10.1215/00318108-114-4-497
- Simion, M. (2021). *Shifty Speech and Independent Thought: Epistemic Normativity in Context*. Oxford: Oxford University Press.
- Simion, M. 2004a. *Resistance to Evidence*. Cambridge: Cambridge University Press.
- Simion, M. 2024b. Resistance to evidence and the duty to believe. *Philosophy and Phenomenological Research*, 108, 203–216.
- Sosa, E. 2015. *Judgment and Agency*. Oxford: Oxford University Press.
- Williams, B. 1973, *Problems of the Self: Philosophical Papers 1956–1972*, Cambridge: Cambridge University Press.