

For *Synthese* Topical Collection: *Epistemic Vices: From the Individual to the Collective* (eds.) I. Kidd & R. McKenna..

Epistemic Vice and Action

Abstract. How do individual-level epistemic vices contribute to bringing about the epistemic vices of groups? This article explores a largely overlooked pathway, where the former influences the latter through the intermediary of individual-level *action* that is itself compromised by epistemically defective individual-level thinking. This pathway is first explained in some detail (with reference to vice-compromised *know-how* and then vice-compromised *intentional action*) and then illustrated concretely via two case studies: Charity Navigator and Enron.

1. INTRODUCTION

I want to begin with a few comments about how this contribution will address the special issue's theme of epistemic vice from the individual to the collective; these will be followed by a brief road map for what will follow.

My starting point is a simple observation about the relationship between individual epistemology and collective epistemology: it is natural to expect that — since groups are collections of members — collective epistemic traits (including epistemic vices) are going to be determined principally by individual-level epistemic traits. Summativists and non-summativists¹ approaches to the epistemology of groups will differ about the kind of determination (e.g., whether it is aggregate or something more complex).² Regardless, focusing broadly on epistemic traits at the the individual level in order to theorise about epistemic traits at the collective level leaves it easy to overlook (in the case of non-summativist views that are of special philosophical interest) a kind of intermediary way that collective epistemic traits might be determined by individual

¹See, e.g., Bird (2010), Gilbert (1987, 2013), and Tuomela (1992, 1995).

²For discussion, see Broncano-Berrocal and Carter (2021).

level epistemic traits — viz., through the intermediary of *individual level action that is itself compromised by defective individual-level thinking*.

I aim to contribute to the special issue’s theme by exploring and sharpening this easily overlooked path from individual to collective epistemic vice. In a bit more detail, I will focus specifically on how individual-level epistemic vices can wreck (or at least downgrade the quality of) both individual level (i) *know-how* and (ii) *intentional action*, and in such a way that will (in a group setting) have a negative bearing on group-level intellectual traits. The result will hopefully be a clearer sense of the dynamics between bad individual thinking and bad group thinking.

Here is the plan. §2 focuses on how individual level epistemic vices can undermine individual-level know-how and the abilities that are necessary for possessing and exercising it. §3 focuses on how individual-level vices in thinking can downgrade or undermine one’s individual-level capacity to act intentionally. §4 then sharpens the intermediate pathway explored (from individual to collective vice) by showing (through two case studies) how individual level vices (particularly, over- and under-confidence in our abilities to act) are capable of wrecking – via the intermediary of *action* – our collective thinking traits.

2. EPISTEMIC VICE AND KNOW-HOW

Know-how of the sort that has been of interest since Ryle (1946, 1949) has been linked implicitly with intelligent action. There might be a loose sense in which we say that you ‘know how’ to digest food simply because you reliably *do* digest food: it is a ‘doing’ attributable to you, but it’s not an action. The process by which you do it is reliable but not manifestation of any kind of intelligence, much less an intelligent action.

The epistemology of know-how has accordingly restricted the kind of know-how constructions of central interest to be “know-how to ϕ ” constructions, where ϕ -ing ranges over *actions*, broadly understood as at least intelligent doings — doings that (unlike, e.g., digestion) are responsive to, and reflective of, our beliefs, desires, and intentions.³

Although know-how is intimately linked with successful intelligent actions (hereafter, actions), virtually no one thinks that, for some action, ϕ , successfully ϕ -ing is sufficient for either manifesting one’s know-how to ϕ in action, or for being (or having ever been) in a state of knowing how to ϕ . A novice locksmith might desire to pick a very complicated warded lock, and through an inflated sense of confidence *believe* they can do so by simply jiggling a wire in such-and-such a way — and suppose doing this actually *works* on this particular occasion. The novice was *able* to pick the lock,

³For discussion on this point, see Bengson and Moffett (2011) and Carter and Poston (2018).

but their picking it did not manifest know-how, and nor did they possess any such know-how in the first place; they just got lucky.

Such cases might look initially like a problem for *anti-intellectualist* views of know-how, which (in a tradition also due to Ryle) identify know-how with *ability possession* rather than (as per intellectualism) with knowing facts. But this isn't so; granted, we sometimes capture the modality of abilities in natural language with "S can ϕ " (example: I can't hit fastballs, but Derek Jeter can!);⁴ however, "S can ϕ " (and *not* ability attributions), qua possibility attributions, also admit of readings that come out true so long as there is some world where (holding fixed some relevant facts about me) I ϕ . For example: an infant can't pick a warded lock but our amateur locksmith (even when just jiggling) can. What anti-intellectualists want to say, in identifying know-how with ability so as to rule out cases like that of our locksmith as cases of know-how, is that know-how requires (something like) an ability to *reliably* pick a lock⁵; this is something our jiggling locksmith lacks.

Whilst 'reliable ability' is meant by the anti-intellectualist to do the work of separating out cases of mere successful action from cases where one manifests know-how⁶, the intellectualist would distinguish these cases by appealing to propositional knowledge. For the intellectualist, a subject, S, knows how to pick a lock iff (to a first approximation) S knows, for some way of picking a lock, w , that w is a way for S to pick a lock.⁷ Whereas exercising know-how for the anti-intellectualist is exercising an ability that is identified with know-how, exercising know-how for the intellectualist is manifesting one's propositional knowledge identified with know-how in action.

Whereas the intellectualist/anti-intellectualist divide disputes what it is that grounds know-how (dispositions versus factual knowledge), broadly the same kinds of *epistemic vices* are capable of undermining that know-how on both sides of this divide. (Or more weakly – since know-how is gradable⁸ – intellectual vices are, on both sides of the intellectualist/anti-intellectualist divide, capable of undermining the degree to which one knows how to do something, even when not undermining that know-how simpliciter.)

⁴For a recent discussion of this kind of proposal, see Schwarz (2020). For an overview of possibilist approaches to abilities and critical discussion of these approaches, see Jaster (2020).

⁵There will be further caveats here — for example, the reliability here will need to be indexed to circumstances appropriate for exercising the ability. A chef still knows how to cook a risotto even when underwater and deprived of ingredients. For related discussion on this point, see, e.g., Carter and Poston (2018) and, in more detail, Sosa (2010, 2015).

⁶Though for some critical pushback against the sufficiency of 'reliable ability' possession for know-how possession, see Bengson and Moffett's (2011, 172–73) 'Kyttoon' case.

⁷See, e.g., Stanley (2011) and Williamson and Stanley (2001).

⁸For discussion on this point, see Pavese (2017). See also Manley and Wasserman (2008) for more general discussion of the gradability of dispositions.

For simplicity, I will focus on two intellectual vices: *overconfidence* and *underconfidence* in one's abilities and/or aptitude. There are perhaps generalised versions of these vices; (some people might be overconfident in all of their abilities). We needn't assume this, though; the level of specificity I'm interested in here just supposes something like a confidence-distinguished vice that manifests in one's being over/under confident in one's would-be implementation of some means they take themselves as having through which to successfully ϕ . Such a vice might lead an overconfident airline pilot to get in the plane with only a cursory glance at the flight plan, for example. Or (in the underconfidence case) might lead a skilled tenor to sing in an inhibited way that compromises the quality of the song's performance.⁹

In the above case of overconfidence, the pilot's know-how (to fly a plane to the relevant destination) is downgraded in quality even if not vanquished by their overconfidence. Without this inflated mismeasure of the pilot's own capacity to implement the means they would implement to fly the plane, the pilot would know-how to fly (to that destination) *better*. This is so even if we concede that the pilot, despite being overconfident, still knows how to, e.g., fly to Boston from Miami, and knows how to do this better than, say, a pilot trainee.

Note that an explanation for *why* overconfidence can downgrade the quality of know-how is going to get different glosses across the intellectualist/anti-intellectualist divide. For the anti-intellectualist, the pilot's vicious overconfident downgrades their know-how by limiting the range and/or reliability of the ability to make the flight successfully. (On one way of thinking of this, the overconfidence might limit the circumstances under which they'd get there successfully to broadly normal circumstances; suppose one who had not so overestimated their abilities and had done the due diligence – e.g., careful study of flight plan and protocols – would perform better if encountering unusual weather patterns, equipment malfunction, etc.). For the intellectualist, the explanation looks a bit different: the pilot's overconfidence downgrades their know-how not by limiting the range/reliability of their ability, *per se*, but by lessening the degree of justification they have for the proposition that their *way* (which includes the non-consulting of the flight plan and protocols) of flying from Boston to Miami is a way for them to get there successfully.

In some cases, presumably, vicious overconfidence in one's (would-be) implementation of one's means of ϕ -ing will undermine one's know-how simpliciter. The chef's overconfidence (perhaps, due to ego-inflating Michelin stars over many years of run-

⁹What I'm calling overconfidence and underconfidence line up closely (even if not perfectly) with an established distinction in vice epistemology — due largely to recent work by Alessandra Tanesini (2021) — between *arrogance* and its contrast point of diffidence or timidity. However, I'll avoid using 'arrogance' to avoid confusion: Tanesini's theory of the psychology of arrogance is rich, and it builds more in to the psychological profiles of this trait as a vice than I need for the present purposes to make the kind of point I'm after.

ning a restaurant) might lead them to not study or practice making *coq au vin*, and then as a result become incapable of executing something that counts as *coq au vin* when the time comes.¹⁰

What goes for overconfidence goes, *mutatis mutandis*, for underconfidence. Returning to our case of the underconfident tenor: the underconfidence might either downgrade the quality of, or wreck completely, their know-how (to sing a particular complex piece on a given occasion). For the anti-intellectualist, this will be a matter of the tenor's underconfidence restricting the range/reliability of her ability to sing the relevant song (and when, doing so beyond some threshold, this will not merely downgrade the quality of the tenor's know-how but undermine it simpliciter).¹¹ For the intellectualist, this will be a matter of the tenor's underconfidence decreasing the level of her justification for the proposition that their way, *w*, is way for them to sing the song successfully – or, in more extreme cases, defeating the justification for that proposition in a way that would defeat the relevant propositional knowledge.¹²

¹⁰Note that chef-style cases are familiar in the literature due to Snowdon (2004), who questions whether a chef who has lost their arms continues to possess the relevant know-how. This case raises complications, and the version of the chef case I'm using here is not meant to suggest the overconfidence described in the case would be analogous to physical impairment when it comes to the downgrading of know-how. For critical discussion of this case and how to think about it within an anti-intellectualist epistemology of know-how, see, e.g., Carter and Pritchard (2015) and Carter and Navarro (2017).

¹¹For a more detailed view of the defeasibility of know-how on an anti-intellectualist framework, see Carter and Navarro (2017).

¹²What I describe here is what I take to be the most salient way that an intellectualist will explain how such overconfidence will have a deleterious effect on one's know-how. However, I don't mean to imply that this would be exhaustive of how the intellectualist might conceptualise this. Here is, potentially, another avenue: intellectualists (Stanley 2011; Stanley and Williamson 2017; Williamson and Stanley 2001; Pavese 2021) take it that the kind of propositional knowledge that should be identified with know-how is propositional knowledge apprehended under what Williamson and Stanley initially described as a *practical mode of presentation* – which is (roughly) a way of thinking of one's way of doing something, ϕ , as a way for her to ϕ . The key idea here is that, just as thinking of yourself under a *de se* mode of thinking might entail certain dispositions (e.g., thinking of the man whose trousers are on fire as *you*) might entail a disposition to try to put the fire out; likewise, thinking of, e.g., a hammer as a way for you to hit a nail (as opposed to, as merely as a way for *one* to hit a nail) also presumably will entail certain dispositions consisting in ways of interacting with the hammer, how to hold it, etc. I mention practical modes of presentation just to point out that an intellectualist might coherently (from within her own theory) suggest that the kind of over/under confidence cases described might not exclusively undermine one's know how by undermining one's epistemic justification for the target proposition (that some way is a way for them to ϕ), but perhaps also by restricting the range or reliability of whatever dispositions would be entailed by knowing the target proposition under a practical mode of presentation. For some useful discussion of practical modes of presentation and the kinds of dispositions they might entail, see Glick (2015) and Pavese (2015).

3. EPISTEMIC VICE AND INTENTIONAL ACTION

Acting intentionally is in some way epistemically constrained. I might desire to hit a golf ball in the hole from 250 yards away (suppose, on a long par-3), take a swing, and by luck the ball goes in. The action of hitting the ball is attributed to me; so is the hole-in-one. But hitting the hole in one isn't something I was able to, or could ever, do intentionally, any more than I could ever win the lottery intentionally.

Acting intentionally accordingly seems to be (in some sense) antithetical to acting accidentally.¹³ The non-accidentality of intentional action is something Anscombe (1957) famously attempted to capture via a kind of practical knowledge condition on intentional action according to which, necessarily, if an agent is ϕ -ing (intentionally and under that description), they know that they are ϕ -ing (intentionally and under that description).¹⁴ Although a condition like Anscombe's can easily explain our initial hole-in-one case (viz., it's false that I *know* that I am hitting a hole one as I am hitting it), it remains contentious whether Anscombe's condition is too strong.¹⁵

For the present purposes, we needn't take a stand; I'll rely just on an idea much weaker than Anscombe's, which is that acting intentionally at the very least requires some kind of accident-precluding control over what one is doing when one acts intentionally; the kind of control required for intentional action needn't be assumed to be as intolerant to risk of error as propositional knowledge (e.g., via a safety condition¹⁶) is intolerant to risk of error; we'll assume, less contentiously, that the tolerance of error here is simply much less than that tolerated by performing accidentally.

What I want to register now is that any kind of accident-precluding control will presumably be subject to being thwarted (or minimised in degree) by epistemic vices such as over- and under-confidence; and so to the extent that acting intentionally requires at least some minimal kind of control over what one is doing when one is acting intentionally, epistemic vice to at least that extent stands to undermine (or minimise) such control.

¹³For some recent discussions on this point, see Shepherd and Carter (Forthcoming).

¹⁴This formulation is due to Piñeros Glasscock (2020, 1238).

¹⁵This point was highlighted notably by Davidson's carbon copier case: 'in writing heavily on this page I may be intending to produce ten legible carbon copies. I do not know, or believe with any confidence, that I am succeeding. But if I am producing ten legible carbon copies, I am certainly doing it intentionally' (Davidson 1970, 92). Granted, it's not clear that a friend of Anscombe's principle couldn't simply reject the intuition here: perhaps Davidson isn't intentionally producing all ten of the carbon copies. Even so, Davidson seems to be on to something: as Shepherd and Carter (Forthcoming) have argued, regardless of what we say about Davidson's carbon copier, it seems that athletes can and often do perform intentionally even when succeeding at rates of success lower than whatever level of reliability would be required by a plausible safety condition on propositional knowledge. See also Carter and Shepherd (2022). Cf., for resistance to this idea, Beddor and Pavese (2021).

¹⁶For discussion, see, e.g., Pritchard (2005, 2007).

Just as one can't hit a hole in one intentionally, we should expect that one can't hit an archery target intentionally *if* one's form stands to the (target-aimed) archery shot (*vis-à-vis* control) in the way that a golfer's shot stands to the hole from 250 yards out. By contrast, whereas even the best golfers can't hit a hole-in-one *intentionally* (even if they can sink two-foot putt intentionally), a skilled archer (say, from 30 yards away) *can* hit a target intentionally (and even if the form exercised in doing so isn't infallible). That said, epistemic vices can downgrade the *reliability of* the archer's form, while at the same time raising the extent to which the relevant success would be lucky. To make this point in the case of epistemically vicious overconfidence, just suppose our archer overestimates the limits of their archery ability, by assuming (mistakenly) that their reliability extends to nighttime and windy conditions, and that they assume this despite never having practised under such conditions. Now, finding themselves in just these conditions, they take a shot when forbearing would have been wiser. Under *these* conditions, the shot — even were it to have hit the target — would lack the kind of accident-precluding control that we may suppose is necessary for hitting the target intentionally. The wider lesson here is as follows: epistemically vicious overestimation of one's abilities can contribute to action that lacks the kind of control that acting intentionally requires *even if* acting intentionally doesn't require (as per Anscombe) practical knowledge.

Epistemically vicious *underconfidence* can undermine intentional action perhaps even more directly than can epistemically vicious overconfidence. This is because at least *belief* (even if not knowledge) that one is doing what one is doing is plausibly in at least some cases¹⁷ necessary for one's doing so intentionally; and even more weakly, believing that one is *not* doing something ϕ (and not merely: not believing that one is doing ϕ) precludes doing ϕ intentionally. With that in mind, consider how epistemically vicious underconfidence can straightforwardly undermine such a belief condition.

For a simple example, take an actor who has disregarded credible evidence that suggests they are in fact very strong. Underestimating their strength, they participate in a fight scene in a manner that results in hurting someone. The disregarding of the evidence of their strength is a manifestation of epistemically vicious underconfidence, and it prevents them from having and sustaining any outright belief that they are capable of easily hurting a colleague in the fight scene. When they do hurt the colleague, it is not *intentional* despite being a very easy possibility, and despite the evidence possessed that would propositionally justify them in acknowledging their strength.

¹⁷It might be that belief is always required. Would-be counterexamples to this suggestion (e.g., intentional improvisational jazz performances) can be addressed by appeal to tacit belief. For discussion of this kind of point – viz., that one's action can be guided by tacit belief that are not consciously entertained – see, e.g., Stanley (2011, Ch. 1).

4. FROM INDIVIDUAL TO COLLECTIVE EPISTEMIC VICE (VIA ACTION)

We've seen how individual-level epistemic vices (with a focus on vices of overconfidence and underconfidence) can potentially downgrade or undermine simpliciter both (i) individual level know-how; and (ii) individual-level intentional action. Against this background, we're now in a position to bring these results together in order to see how individual-level vices of under- and over-confidence can easily (in a group setting) have a negative bearing on group-level epistemic traits *by way of* undermining the kind of know-how and intentional action that is essential to bringing about (epistemically) good group-level thinking.

In what follows, I will trace out this kind of pathway from individual-to-collective epistemic vice (via the intermediary of action) using Margaret Gilbert's (1987, 2013) well-known joint-commitment approach to collective traits as illustrative.¹⁸

But first, a proviso. My use of 'collective traits' here and elsewhere is meant to refer specifically to *non-summativist* traits; for a given trait F , F is trait of a group, G , in a non-summativist sense if the group's having the particular trait does not require any individual member of G to have F .¹⁹ To use a simple example: a non-summativist characterisation of (the company) Philip Morris's belief (or acceptance²⁰) that smoking is not causally linked to cancer will be compatible (non-summatively) with all of the individual members of the group accepting, believing, or even knowing that this is false. What matters for Philip Morris's possession of this trait – viz., of believing in the non-causal efficacy of smoking for cancer – isn't something we determine on a non-summativist model by asking about what the members individually believe, *per se*, about smoking and cancer, but rather, by asking whether the individual members interact in such a way as to give rise to a group-level belief/acceptance with that proposition as its content. Different non-summativist proposals will characterise the relevant individual-level interactions differently.²¹ On the Gilbert-style model I'll be

¹⁸The gist of the point I will be making with reference to Gilbert's framework will also be applicable, *mutatis mutandis*, to other frameworks where individual-level action plays a role in determining group-level epistemic behaviour – mostly notably, distributed/functionalist frameworks for collective properties such as Alexander Bird's (2010).

¹⁹For an overview of non-summativism in its connection with various types of summativism, see Broncano-Berrocal and Carter (2021).

²⁰Whether groups have full-blown (non-summativist) beliefs or mere non-summativist 'acceptances' is contentious in collective epistemology (see, e.g., Simion et al. 2020); nothing much hangs on this dispute for the present purposes. I'll refer to group beliefs, however, the reader may feel free to substitute 'acceptance' for belief. What will be of interest is group-level (non-summativist) epistemic character as opposed to beliefs specifically.

²¹See, for example, Tollefsen (2002; 2004, 2015), Bird (2010), Palermos (2016, 2020), Carter (2020, 2015), Tuomela (1992, 1995), and De Ridder (2014).

using as illustrative, what is relevant here is whether the individual members of the group have made a *joint commitment* with respect to smoking and cancer – viz., a commitment to take a stand as a group on the matter of whether the former is a cause of the latter.

Of course, a group can't do anything (including make a commitment, as a group) without individual members acting in some way. This general principle is sometimes called the 'group/member action principle': For every group, G , and act, ϕ , G performs ϕ only if at least one member of G performs some act or other that causally contributes to ϕ .²² On Gilbert's view, the joint commitments that ground group-level traits – commitments to act "in such a way as to emulate as best they can a single body with the goal in question" (Gilbert 2017, 33) – are brought about by individuals committing to act in this way provided others do. They are thus individual-level conditional commitments. In the case of a group belief, a Gilbert-style non-summative account will maintain that whatever the group believes emerges as a result of the individual members jointly committing to *act as if the belief is true* conditioned on others acting this way. Note, of course, that one can act as if a given proposition, $\langle p \rangle$, is true (and thus, discharge one's individual-level commitment as a member of a group that believes that $\langle p \rangle$) even if privately harbouring doubts about whether $\langle p \rangle$ is true.²³

The possibility of group-level beliefs emergent from individual joint commitments gives rise naturally to the possibility of group level justification and *knowledge* that might arise from such group beliefs.²⁴ Philip Morris clearly doesn't *know* smoking doesn't cause cancer because the target proposition in this case is false. For a more epistemically interesting case, consider whether the CDC knows that the wearing a mask curtails the spread of Covid-19. That proposition is true. If the CDC believes it (in virtue of a Gilbert-style joint commitment) it seems very plausibly that they also *know* it. The CDC, as the thought goes, is – epistemically – responsible and reliable in the way that it (as a group) goes about forming its group-level beliefs, and in a way that might seem analogous to the kind of responsibility and reliability we'd attribute to an individual-level scientist who comes to the same view. Further, just as individual level scientists also can have individual level epistemic virtues, so presumably would the CDC at a group level have such virtues; and by contrast, Philip Morris might have group-level epistemic vices that seem analogous to the individual-level epistemic vices of deniers of a smoking-cancer link, where such individuals form their beliefs (with this content) in irresponsible and unreliable ways.

Miranda Fricker (2010) suggests how we might construe a collective epistemic character trait – including both epistemic virtues and vices – within a Gilbert-style

²²See Lackey (2014); cf., Carter (2015).

²³See Mathiesen (2006) and Carter (2015). for discussion.

²⁴For a helpful discussion on this point, see Mathiesen (2006)

joint commitment framework. If we focus on collective epistemic virtue, the idea is as follows: a group *G* has a collective epistemic virtue when *G*'s members jointly commit to an (epistemically) good motive, and do so because the motive is epistemically good. A jury on this view is open-minded at the group-level iff its members jointly commit to behaving in characteristically open-minded ways as a body and because open minded inquiry is truth-conductive (as opposed to, say, because the jury is threatened to jointly commit this way).²⁵ This joint commitment (as all joint-commitments, on a Gilbert-style framework) implicates individual-level commitments: regardless of whether the jurors are in fact open-minded in an individual capacity – when they are acting *as* jurors, they must (in the sense that are normatively obliged to) contribute to deliberation in the open-minded ways committed to.

Collective epistemic vices are somewhat more complicated for the above sort of proposal which grounds collective character in joint commitments. This is because the proposal can't, as would seem initially tempting, simply 'invert' its motivational clause to model collective epistemic vice without requiring some additional caveats. Complexities arise once we consider a potential disanalogy with individual epistemic vice: at the individual level, we have plenty examples of vices (e.g., epistemic carelessness, etc.) that are knowledge-obstructing but not in virtue of any positive commitment at the individual level to a bad epistemic end *as such*.²⁶ Presumably, and absent any good reason to the contrary, we should expect the same to be the case at the collective level. In addition, even in those cases where individual-level vices do seem to involve commitments to bad epistemic ends, we should perhaps relax the requirement on epistemic vice (at least, in some cases) that the epistemically bad end be desired *because* it is bad. For example, at the individual level, desiring to 'cut corners' in a research project seems epistemically vicious even if one doesn't desire this *because* cutting corners is epistemically bad, but rather, because the assignment is due earlier than they expected.

Here is not the place to defend any particular substantive view of collective epistemic vice within a Gilbert-style framework. Rather, it suffices to note that even if

²⁵The idea that the motivational profile of a virtue lines up with desire to bring about some end, *E*, because *E* is good, is found in individual-level virtue epistemology, e.g., Zagzebski (1996) and Baehr (2011). This idea raises potential problems; for example, as Baehr (2011) notes, it implies that a scientist will lack epistemic virtue if the scientist is motivated to seek the truth but has *this* motivation only out of a more general motivation to win a Nobel prize. Whilst Baehr finds the result of denying epistemic virtue to the scientist acceptable, it will perhaps strike some as too strong; at least, it does not seem infelicitous to attribute epistemic virtue in such a case, especially when we make salient the intellectual difficulty of the achievement. For present purposes, it suffices to register that if Baehr's and Zagzebski's characterisation of the motivational component of an epistemic virtue is too strong at the individual level, Fricker's characterisation will be likewise too strong given that she inherits this aspect of their individual-level view at the collective level.

²⁶For a helpful discussion of this kind of point at the individual level, see Crerar (2018).

Fricker’s proposal is too strong as a view of collective epistemic vice for the kinds of reasons noted above, something in the neighbourhood will look plausible enough to make sense of paradigmatic cases of collective epistemic vice, such as our guiding example case of Philip Morris. To make sense of such a case, we will need to just substitute a *de re* for a *de dicto* reading of Fricker’s motivational clause, which weakens the proposal in the following way: sufficient for (at least some) collective epistemic vices is that a group’s members jointly commit to an (epistemically) bad motive *de re*. A motivation to suppress evidence against smoking is an epistemically bad motive *de re*, regardless of whether the group commits to this motive *because* they identify it as epistemically bad (as opposed to, e.g., committing to it to maximise profits). Further, the revised *de re* clause is one that may be read least contentiously as representing a sufficiency condition rather than a necessity condition – viz., such a joint commitment to an epistemically bad motive *de re* suffices for the group to have an epistemic vice. This qualified spin on Fricker’s proposal is further going to be compatible with countenancing that the satisfaction of other (perhaps even less stringent) conditions will sometimes suffice for a group’s possessing a collective epistemic vice (e.g., in cases like collective carelessness).

On the above working assumption, then, a joint commitment to a *de re* epistemically bad motive will suffice, at in some paradigmatic cases, for group epistemic vice on a broadly Gilbert-style joint commitment account of group level traits. With this assumption in hand, we’re now in a position to appreciate how individual level vices of the sort discussed in §§2-3 might contribute to group-level epistemic vice (on a joint commitment account) via the intermediaries of individual-level know-how (§4.1) and intentional action (§4.2).

4.1. *via Know-how*

Let’s begin with know-how. A useful case study for investigating this path – viz., from individual level epistemic vice to collective epistemic vice *by way of* undermining individual-level *know-how* – is found in C. Thi Nguyen’s (2021) recent paper “Transparency is Surveillance”. Here Nguyen focuses (albeit, to his own argumentative ends²⁷) on the case of the [Charity Navigator](#) organisation, which aims to help prospective donors make informed choices about which are the most efficient and least corrupt charities to which to donate.

Charity Navigator, however, is not a story of epistemic success, at least not prior to radically changing its early methodology.²⁸ In its early years, the employees of

²⁷Nguyen uses the Charity Navigator example in order to show how policies that prioritise transparency can be epistemically self-defeating. Or, more weakly, he purports to show how a commitment to transparency can lead causally to negative epistemic (as well as practical) consequences.

²⁸The most recent methodology (post 2020) differs signifi-

Charity Navigator – while jointly committing in their charter to the above *prima facie* epistemically (and morally) good motive – failed to bring about either end (and thus failed to do what they jointly committed *de dicto* to do). For our purposes, what is of interest is that what they in fact committed to doing in the name of facilitating epistemically good, informed choices – was a policy that inadvertently misled donors. The underlying problem was that the individuals within the Charity Navigator (who were tasked with various statistical work and data collection), relied uncritically in the service of executing an ostensibly good epistemic motive, on what was later shown to be an unreliable metric for indicating – as they purported to – efficiency and low corruption²⁹. This was the simple simple metric of “low overhead”. As Nguyen notes, low overhead “seems good on quick inspection” as a metric for ranking charities in terms of efficiency and low corruption. It was taken for granted to be so by the individual researchers and auditors at Charity Navigator (for roughly the first seven years of the Charity’s existence³⁰) to be a reliable way for them to keep donors well-informed about efficiency and low-corruption; but this metric (and ranking relative to this metric) effectively misled donors to the detriment of Charity Navigator and the donors. They didn’t, at least initially (until changing methodology) execute an epistemically good motive, but committed instead to something *de re* epistemically risky.

So where did things go wrong? Of course there are likely multiple explanations and multiple contributing factors.³¹ But here is at least one plausible theoretical lens through which to view the case. Consider that, were the individuals in the Charity Navigator organisation initially less hasty in settling on their *way of bringing about* what they purported to commit to do (i.e., to help donors make informed decisions about charities’ efficiency and comparative corruption), they would have known better *how* to achieve what is, on the face of it, an epistemically worthy group aim.

In sum, without initially more critically examining the metric that looked good on quick inspection, the Charity Navigator team as a result lacked the relevant know-how to bring its aims about. And notice that the lack of know-how here contributed directly to a *de re* epistemically bad motive (even if not an epistemically bad motive *de dicto*): they committed jointly to actions and policies that misled.

The path from individual to collective epistemic vice in Charity Navigator’s early years accordingly does not look directly like this: “individual epistemic vice → collective epistemic vice” but more accurately the path appears to look like: “individual epistemic vice → defective individual know-how → collective epistemic vice”.

cantly from early methodology. <https://www.fastcompany.com/90533562/charity-navigator-is-launching-a-new-rating-system-to-help-better-evaluate-nonprofits>.

²⁹See, for discussion, Meer (2017).

³⁰The period of discussion here is 2001-2008, when revisions to the initial methodology were made. For additional criticisms of the early methodology, see Lowell et al. (2005).

³¹See Lowell et al. (2005).

Three brief clarificatory caveats are in order before moving on: (i) possessing *an* epistemic vice (either individual or collective) needn't implicate epistemic viciousness wholesale. (I am assuming that any so-called 'unity of the virtues/vices' thesis of the sort Plato embraced is mistaken: someone can be generous but not wise, intellectually lazy while not intellectually dishonest, etc.) The above interpretation of the Charity Navigator case obviously accordingly needn't imply that in its early years the organisation or its members were wholly (or even mostly or usually!) epistemically vicious; (ii) it should also be noted that post-2020, Charity Navigator's practices are entirely different than in its early years which were of initial interest to Nguyen as well as to me for the purposes of discussing them critically. The four indicators now (post-2020) used are: (a) finance and accountability; (b) impact and results; (c) leadership and adaptability; (d) culture and community. Notice that none of indicators (b-d) would have been captured straightforwardly by the initial low-overheads indicator; (iii) thirdly, and importantly, it is worth registering that the kind of path "individual epistemic vice \rightarrow defective individual know-how \rightarrow collective epistemic vice" is not plausibly idiosyncratic to early-years Charity Navigator. We find it wherever collectives lack the relevant know-how to bring about good epistemic aims committed to explicitly, and where the lack of know-how is sourced in bad styles or traits of thinking on the part of the group's members.

4.2. *via Intentional action*

Prior to filing for bankruptcy in 2001, Texas energy company Enron was, on the surface of things, synonymous with its commitment to *innovation*. The company was voted "America's Most Innovative Company" for six consecutive years by Fortune Magazine.³² Its company motto was "Ask why".

Enron is now synonymous not with innovation, but with with one of the largest-scale financial frauds in U.S. history.³³ The company and its practices offers a useful case study for how individual-level epistemic vice (in this case, of overconfidence – particularly on the part of individuals occupying high-level roles within the organisation) can contribute to collective epistemic vice via the intermediary of a compromised capacity for *intentional action*; in this particular case, individual-level epistemic hubris led to a loss of the kind of *control* that (as per §3), and – without the capacity act so as to make good on its financial promises – Enron made even larger promises in an effort to keep up the stock price.³⁴

The individual-level overestimation by Enron's employees of their own abilities is best illustrated by the practice of 'mark-to-market' accounting, initiated by former

³²See, Van Niel and Rapoport (2004).

³³See, e.g., Healy and Palepu (2003).

³⁴See Healy and Palepu (2003), Tonge et al. (2003), and Kulik (2005).

CEO Jeffrey Skilling. On this type of accounting, projected future profits are represented on balance sheets as a mechanism of capturing the ‘fair market’ value of a company’s assets. Mark-to-market accounting isn’t inherently misleading. For example, if (say) a drug company has just had a drug approved by the FDA, the fair market value of the company might not be accurately reflected if one were to assess the company’s assets with reference *only* to historical profits. Even so, mark-to-market accounting incentivises optimism on the part of a company³⁵; the incentivisation structure looks like this: the higher the future profits are estimated as being, the better the financial health of the company according to its balance sheets; the better the balance sheets, the higher the stock price; and the higher the stock price, the more wealth is generated for the shareholders.

Much has been written in business ethics literature on the moral failings³⁶ of Enron in the wake of its eventual bankruptcy 2001 and imprisonment of its CEO Skilling and chairman and founder Kenneth Lay. Even abstracting from the moral dimension of the scandal entirely – there is also an important epistemic narrative here. To a first approximation: decisionmakers within Enron, quite literally, overestimated themselves over and over. The more inaccurate the overestimation (as communicated in the balance sheets that reflected projected future profits), the more incentivisation to continue overestimating on future balance sheets. The accumulated overestimation widened the gap between the promise and the capacity to deliver, to the point that “managers lost control over [Enron’s] operations” (Culpan and Trussel 2005, 72). Framed within the context of our discussion from §3, the habitual overestimation reflected on the balance sheets of what the company *would* do, effectively left Enron without the kind of control that is necessary to manage its affairs as a company. The loss of company-level control, though it is sourced partly in moral failings such as greed, is plausibly at least sourced in (and causally determined by) a culture of overconfidence – of epistemically vicious overestimation of what could be delivered at a later date.

The path from individual to collective epistemic vice in Enron’s case is one where loss of control is central to the story. And this loss of control, which is reflected by, and then further perpetuated, Enron’s group-level commitments to (increasingly) misleading representations of its financial situation – looks very much to be traced back to initial individual-level overestimations by (at least some) members of Enron – viz., of what they could do through their own efforts and when they could do it by.

³⁵See <https://corporatefinanceinstitute.com/resources/knowledge/other/enron-scandal/>.

³⁶See, e.g., Appel et al. (2014), Nolen (2007), Cohan (2002), Kulik (2005), and Premeaux (2009).

5. CONCLUDING REMARKS

By looking in a bit of detail at two case studies – i.e., Charity Navigator and Enron – we are able to see clearly two pathways from individual to collective vice through the intermediary of action. In the Charity Navigator case, the intermediary was through action by way of know-how. In sum, individual-level overconfidence in a way of bringing about an epistemically worthy objective ultimately undermined their (would-be) know-how to attain that very objective, at least during its early years; it led instead to the group committing to a *de facto* epistemically problematic end: namely, the communication of misleading metrics for individuals to rely on. By contrast, in the Enron case, the relevant intermediary between individual and collective epistemic vice is better understood as intentional action. In this case, individual-level epistemically vicious overconfidence (in particular, by individuals in senior management) led the company to lose the kind of *control* over its operations and their legality, which perpetuated the dissolution of Enron as a collective with agency.

Note that, in both cases, the path to collective epistemic vice from individual epistemic vice is not straightforward; to appreciate the latter's effects on the former – and the mechanisms through which it brings about these effects – we need to turn our attention to individual-level action, how such action is influenced by bad individual thinking, and – when it is – how such defects in individuals capacities to act can play a role in defective group-level epistemic commitments.

REFERENCES

- 10 Anscombe, G. 1957. *Intention*. Harvard University Press.
- Appel, Randy, Jae Fratzl, Ruth B. McKay, and Carey Stevens. 2014. “Enron and the 12 Steps of White-Collar Crime.” *International Journal of Business Governance and Ethics* 9 (4): 381. <https://doi.org/10.1504/IJBGE.2014.066277>.
- Baehr, Jason. 2011. *The Inquiring Mind: On Intellectual Virtues and Virtue Epistemology*. OUP Oxford.
- Beddor, Bob, and Carlotta Pavese. 2021. “Practical Knowledge Without Luminosity.” *Mind*, October. <https://doi.org/10.1093/mind/fzab041>.
- Bengson, John, and Marc A. Moffett. 2011. “Two Conceptions of Mind and Action: Knowledge How and the Philosophical Theory of Intelligence.”
- Bird, Alexander. 2010. “Social Knowing: The Social Sense of Scientific Knowledge.” *Philosophical Perspectives* 24: 23–56.
- Broncano-Berrocal, Fernando, and J. Adam Carter. 2021. *The Philosophy of Group Polarization: Epistemology, Metaphysics, Psychology*. Routledge.

- Carter, J Adam, and Jesús Navarro. 2017. "The Defeasibility of Knowledge-How." *Philosophy and Phenomenological Research* 95 (3): 662–85. <https://doi.org/10.1111/phpr.12441>.
- Carter, J Adam, and Duncan Pritchard. 2015. "Knowledge-How and Epistemic Value." *Australasian Journal of Philosophy* 93 (4): 799–816.
- Carter, J Adam, and Joshua Shepherd. 2022. "Intentional Action and Knowledge-Centered Theories of Control." *Manuscript*.
- Carter, J. Adam. 2015. "Group Knowledge and Epistemic Defeat." *Ergo, an Open Access Journal of Philosophy* 2. <https://doi.org/https://doi.org/10.3998/ergo.12405314.0002.028>.
- . 2020. "Collective (Telic) Virtue Epistemology." In *Social Virtue Epistemology*, edited by Mark Alfano, Jeroen de Ridder, and Colin Klein. London: Routledge.
- Carter, J. Adam, and Ted Poston. 2018. *A Critical Introduction to Knowledge-How*. Bloomsbury Publishing.
- Cohan, John Alan. 2002. "‘I Didn’t Know’ and ‘I Was Only Doing My Job’: Has Corporate Governance Careened Out of Control? A Case Study of Enron’s Information Myopia." *Journal of Business Ethics* 40 (3): 275–99. <https://doi.org/10.1023/A:1020506501398>.
- Crerar, Charlie. 2018. "Motivational Approaches to Intellectual Vice." *Australasian Journal of Philosophy* 96 (4): 753–66. <https://doi.org/10.1080/00048402.2017.1394334>.
- Culpan, Refik, and John Trussel. 2005. "Applying the Agency and Stakeholder Theories to the Enron Debacle: An Ethical Perspective." *Business and Society Review* 110 (1): 59–76. <https://doi.org/10.1111/j.0045-3609.2005.00004.x>.
- Davidson, Donald. 1970. *Essays on Actions and Events: Philosophical Essays Volume 1*. Clarendon Press.
- Fricker, Miranda. 2010. "Can There Be Institutional Virtues?" In *Oxford Studies in Epistemology*, edited by T. Szabo Gendler and J. Hawthorne, 3–235. Oxford University Press.
- Gilbert, Margaret. 1987. "Modelling Collective Belief." *Synthese* 73 (1): 185–204.
- . 2013. *Joint Commitment: How We Make the Social World*. Oxford University Press.
- Glick, Ephraim. 2015. "Practical Modes of Presentation." *Noûs* 49 (3): 538–59.
- Healy, Paul M., and Krishna G. Palepu. 2003. "The Fall of Enron." *Journal of Economic Perspectives* 17 (2): 3–26. <https://doi.org/10.1257/089533003765888403>.
- Jaster, Romy. 2020. *Agents’ Abilities*. Berlin, New York: De Gruyter.

- Kulik, Brian W. 2005. "Agency Theory, Reasoning and Culture at Enron: In Search of a Solution." *Journal of Business Ethics* 59 (4): 347–60. <https://doi.org/10.1007/s10551-004-7308-2>.
- Lackey, Jennifer. 2014. "Socially Extended Knowledge." *Philosophical Issues* 24 (1): 282–98. <https://doi.org/10.1111/phis.12034>.
- Manley, David, and Ryan Wasserman. 2008. "On Linking Dispositions and Conditionals." *Mind* 117 (465): 59–84. <https://doi.org/10.1093/mind/fzn003>.
- Mathiesen, Kay. 2006. "The Epistemic Features of Group Belief." *Episteme* 2 (3): 161–75.
- Meer, Jonathan. 2017. "Are Overhead Costs a Good Guide for Charitable Giving?" *IZA World of Labor*, 329–29. <https://ideas.repec.org/a/iza/izawol/journly2017n329.html>.
- Nguyen, C. Thi. 2021. "Transparency Is Surveillance." *Philosophy and Phenomenological Research*. <https://doi.org/10.1111/phpr.12823>.
- Nolen, Donald M. 2007. "Business Ethics After Enron." *The Proceedings of the Twenty-First World Congress of Philosophy* 1: 281–84. <https://doi.org/10.5840/wcp2120071305>.
- Palermos, S. Orestis. 2016. "The Dynamics of Group Cognition." *Minds and Machines* 26 (4): 409–40.
- . 2020. "Epistemic Collaborations: Distributed Cognition and Virtue Reliabilism." *Erkenntnis*, 1–20.
- Pavese, Carlotta. 2015. "Practical Senses." *Philosophers' Imprint* 15.
- . 2017. "Know-How and Gradability." *The Philosophical Review* 126 (3): 345–83.
- . 2021. "Know-How, Action, and Luck." *Synthese* 198 (7): 1595–1617.
- Piñeros Glasscock, Juan S. 2020. "Practical Knowledge and Luminosity." *Mind* 129 (516): 1237–67.
- Premeaux, Shane. 2009. "The Link Between Management Behavior and Ethical Philosophy in the Wake of the Enron Convictions." *Journal of Business Ethics* 85 (1): 13–25. <https://doi.org/10.1007/s10551-008-9745-9>.
- Pritchard, Duncan. 2005. *Epistemic Luck*. Clarendon Press.
- . 2007. "Anti-Luck Epistemology." *Synthese* 158 (3): 277–97.
- Ridder, Jeroen de. 2014. "Epistemic Dependence and Collective Scientific Knowledge." *Synthese* 191 (1): 1–17. <https://doi.org/10.1007/s11229-013-0283-3>.
- Ryle, Gilbert. 1946. "Knowing How and Knowing That: The Presidential Address." *Proceedings of the Aristotelian Society* 46: 1–16. <https://doi.org/10.1093/aristotelian/>

46.1.1.

- . 1949. *The Concept of Mind*. Hutchinson & Co.
- Schwarz, Wolfgang. 2020. “Ability and Possibility.” *Philosopher’s Imprint* 20 (6).
- Shepherd, Joshua, and J. Adam Carter. Forthcoming. “Knowledge, Practical Knowledge, and Intentional Action.” *Ergo: An Open Access Journal of Philosophy*, Forthcoming.
- Simion, Mona, J. Adam Carter, and Christoph Kelp. 2020. “On Behalf of Knowledge-First Collective Epistemology.” In, edited by Paul Silva and Luis Oliveira. Routledge. <https://eprints.gla.ac.uk/222812/>.
- Snowdon, Paul. 2004. “Knowing How and Knowing That: A Distinction Reconsidered.” *Proceedings of the Aristotelian Society* 104 (1): 1–29. <https://doi.org/10.1111/j.0309-7013.2004.00079.x>.
- Sosa, Ernest. 2010. “How Competence Matters in Epistemology.” *Philosophical Perspectives* 24 (1): 465–75.
- . 2015. *Judgment & Agency*. Oxford University Press UK.
- Stanley, Jason. 2011. *Know How*. OUP Oxford.
- Stanley, Jason, and Timothy Williamson. 2017. “Skill.” *Noûs* 51 (4): 713–26. <https://doi.org/https://doi.org/10.1111/nous.12144>.
- Tanesini, Alessandra. 2021. *The Mismeasure of the Self: A Study in Vice Epistemology*. Oxford University Press.
- Tollefsen, Deborah. 2004. “Collective Epistemic Agency.” *Southwest Philosophy Review* 20 (1): 55–66. <https://doi.org/10.5840/swphilreview20042015>.
- . 2015. *Groups as Agents*. Polity.
- Tollefsen, Deborah Perron. 2002. “Collective Intentionality and the Social Sciences.” *Philosophy of the Social Sciences* 32 (1): 25–50. <https://doi.org/10.1177/004839310203200102>.
- Tonge, Alyson, Lesley Greer, and Alan Lawton. 2003. “The Enron Story: You Can Fool Some of the People Some of the Time ?” *Business Ethics, the Environment and Responsibility* 12 (1): 4–22. <https://doi.org/10.1111/1467-8608.00301>.
- Tuomela, Raimo. 1992. “Group Beliefs.” *Synthese* 91 (3): 285–318. <https://doi.org/10.1007/bf00413570>.
- . 1995. *The Importance of Us: A Philosophical Study of Basic Social Notions*. Stanford University Press.
- Van Niel, Jeffrey D., and Nancy B. Rapoport. 2004. “Dr. Jekyll & Mr. Skilling: How Enron’s Public Image Morphed from the Most Innovative Company in the Fortune 500 to the Most Notorious Company Ever.” SSRN Scholarly Paper ID

505662. Rochester, NY: Social Science Research Network. <https://papers.ssrn.com/abstract=505662>.

Williamson, Timothy, and Jason Stanley. 2001. "Knowing How." *The Journal of Philosophy* 98 (8). <https://ora.ox.ac.uk/objects/uuid:1fcf732b-731b-47e6-943d-aaf34905c02a>.

Zagzebski, Linda T. 1996. *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge*. Cambridge University Press.